
#VOTELEAVE OR #STRONGERIN: Resonance and Rhetoric in the EU Referendum

Thesis submitted in partial fulfilment of the requirements for the degree of
DPhil in Information, Communication and the Social Sciences
in the Oxford Internet Institute at the University of Oxford

Yin Yin Lu
Balliol College



99,159 words

Supervisors: Dr Grant Blank and Dr Jonathan Bright

MAY 2019

#VoteLeave or #StrongerIn: Resonance and Rhetoric in the EU Referendum

Yin Yin Lu

Balliol College

DPhil in Information, Communication, and the Social Sciences

May 2019

ABSTRACT

This thesis explores the persuasive function of tweets in the context of the 2016 European Union membership referendum. *What makes a message resonate with its audience on social media? How can the message be expressed and delivered most effectively? What strategies do individuals and organisations use when they send a tweet and why?*

To explain tweet success, or resonance, an interdisciplinary framework centred around rhetoric was established. It is comprised of nine different features: discoverability, organisational authority, message framing and structure, emotions, topical issues, resolution, questions, and playfulness.

A multi-method design is adopted to address the research questions from micro- and macroscopic perspectives. Approximately 9.8 million referendum-related tweets were collected from Twitter during the referendum period (15 April-22 June 2016). A random sample of 2,976 was hand-coded according to the framework. Semi-supervised algorithms were used to label the features in the full dataset, and multivariate regressions were computed to assess the influence of each feature. This quantitative analysis is complemented by interviews with 22 campaign leaders, journalists, politicians, digital marketing experts, and private individuals who tweeted actively during the referendum. Interviewees included the CEOs of Vote Leave and Stronger In, BuzzFeed's Europe Editor, a Twitter board member, and a former MP and Minister of State for Europe. Finally, exceptional tweets and users were identified from the regression models and interview data and rigorously scrutinised.

The overarching conclusion is that in the fast-paced and high-volume landscape of social media, discoverability is more important for resonance than content or style. Increasing the likelihood that a message is seen matters more than what is said or how it is said, because the vast majority of tweets are not seen. As such, this thesis introduces a revised schema for the resonance of digital cultural objects premised upon algorithmic power, visual media, and organisational authority.

ACKNOWLEDGEMENTS

It doesn't take a village to successfully complete a doctoral thesis. It takes a proper army.

Without the support of the following individuals and organisations, this ambitious project would not have been achievable, or even conceivable.

The generals of my army are undoubtedly Dr Grant Blank and Dr Jonathan Bright, my ever-insightful, ever-supportive, and ever-patient supervisors. They have inspired me to transcend disciplinary and methodological boundaries to a level beyond what I could have imagined. They have also inspired me to defy my own assumptions. I never set out to write a thesis on politics and Twitter. In fact, when I arrived at the Oxford Internet Institute in October 2014, I explicitly declared that I would *not* analyse political content and that I would *not* exclusively focus on tweets. But I am so glad that I did.

My technical expertise was close to zero when I embarked on this project, but I knew that I wanted to analyse language at scale and properly learn Python. I now consider myself conversationally fluent in Python for data science and natural language processing, thanks to the invaluable assistance of Thomas Talbot, Jesse Sigal, Scott Hale, Chico Camargo, and Kelly Littlepage. I am also immensely grateful to SoBigData for the funding support they provided for my visit to the Fraunhofer Institute for Computer Graphics Research IGD in Darmstadt, Germany. It was there that I properly taught myself supervised machine learning for text classification, under the guidance of Dr Thorsten May and Hendrik Lücke-Tieke.

The qualitative dimension of this thesis is as important as the quantitative dimension, and it is largely defined by the extremely rich conversations I had with my 22 interviewees. These conversations exceeded my expectations with regard to length, depth, and—dare I say—resonance, and were a significant motivation during my writing up period. My most heartfelt gratitude goes to all of my interviewees for their generosity with their time, insights, and Twitter data.

I could not have chosen a better home for my DPhil than the Oxford Internet Institute, and indeed, it was the only doctoral programme for which I applied. The students, faculty, and administrative staff are all exceptional. I must thank Bertie Vidgen for coding several hundred tweets to help me establish sufficient validity, Arthur Bullard for letting me repeatedly borrow an exceptional recorder to capture my interviews in sterling quality, Duncan Passey for processing all of my expenses in world-record time, Ulrike Deetjen for her continuous mentorship throughout the past four years, and the DPhil programme directors—Professor Vili Lehdonvirta and Professor Gina Neff—for facilitating such engaging and thought-provoking seminars.

Without the Clarendon Fund, this inimitable journey would never have happened in the first place. I am grateful not only for the full financial support I have received, but also for the lifelong friends I have made through the Scholars community.

Finally, I must thank my parents for always being there for me when I needed it, despite being five timezones and an ocean away. Without my father and his extremely powerful server, my data collection would not have been nearly as efficient as it was. I hope this work will make them proud.

TABLE OF CONTENTS

List of Figures	04
List of Tables	06
1 Introduction and Background	07
1.1 Background of Twitter	10
1.2 Research Questions and Thesis Structure	13
1.3 Significance of Research	15
2 Theory and Literature Review	17
2.1 Tweets as Digital Cultural Objects	18
2.1.1 Definition of a Cultural Object	19
2.1.2 Digital Culture	21
2.2 Definitions of Tweet Success	21
2.2.1 Retweets and Likes	23
2.2.2 Explanations for Resonance	27
2.3 Discoverability	28
2.3.1 Tweet Attributes	29
2.3.2 Organisational Authority	31
2.4 Rhetorical Content	35
2.4.1 Message Framing: Identification vs. Division	36
2.4.2 Emotions: Enthusiasm, Anxiety, and Aversion	41
2.4.3 Relevant Issues	44
2.4.4 Resolution	47
2.5 Rhetorical Style	49
2.5.1 Message Format: Lists and Contrasts	49
2.5.2 Questions	50
2.5.3 Playfulness	52
2.6 Synthesis of Success Factors	55
3 Methodology	57
3.1 Case Study Site Selection: The EU Referendum	58
3.2 Twitter Data Collection, Processing, and Analysis	61
3.2.1 Tweet Collection	61
3.2.2 Twitter Data Limitations	62
3.2.3 Tweet Storage and Processing	64
3.2.4 Measurement of Tweet Resonance	66
3.2.5 Tweet Sampling and Content Analysis	67
3.3 Automated Coding of Full Dataset	72
3.3.1 Text Preprocessing and Vectorisation	73
3.3.2 Algorithm Training and Testing	74
3.3.3 Model Refinement and Selection	76
3.3.4 Dictionary Approach	81
3.4 Interview Strategy and Analysis	82

3.4.1	Interviewee Selection	83
3.4.2	Interview Structure	83
3.4.3	Interview Analysis Process	87
3.5	Ethical Considerations	88
4	Factors that Influence Tweet Resonance	91
4.1	Summary Statistics	92
4.1.1	Dependent Variables	93
4.1.2	Independent Variables	95
4.1.3	Control Variables	106
4.2	Sample-Level Models	112
4.2.1	Regression Diagnostics	113
4.2.2	Control Features	115
4.2.3	Discoverability Features	118
4.2.4	Rhetorical Content	119
4.2.5	Rhetorical Style	120
4.2.6	Full Models	120
4.3	Population-Level Models	122
4.4	Discussion	125
5	Strategies for Tweet Resonance	129
5.1	Interviewees	131
5.1.1	Interviewee Backgrounds and Motives	131
5.1.2	Conceptions of Success and Intended Audiences	133
5.2	Organisational Authority	136
5.2.1	Time on Twitter	137
5.2.2	Orchestrated Success	139
5.2.3	Strategies to Improve Prominence	141
5.3	Discoverability	144
5.3.1	Visual Media	144
5.3.2	Hyperlinks	151
5.3.3	Hashtags	154
5.3.4	Mentions	157
5.3.5	Other Discoverability Strategies	159
5.4	Timing	160
5.4.1	Hour and Day	161
5.4.2	Timeliness and Novelty	164
5.5	Rhetorical Content	168
5.5.1	Emotions	169
5.5.2	Resolution	174
5.5.3	Relevant Issues	176
5.5.4	Message Framing	181
5.5.5	Celebrities	191
5.6	Rhetorical Style	192
5.6.1	Brevity, Clarity, and Simplicity	193
5.6.2	Playfulness	194
5.6.3	Questions	197
5.7	Discussion	198

6	Exceptional Tweets and Users	201
6.1	Most Popular Tweets	202
6.1.1	Tweets from Individuals	203
6.1.2	Tweets from Celebrities and Established Organisations	206
6.2	Outliers (Quantitative Exceptions)	211
6.2.1	Surprisingly Successful Tweets	212
6.2.2	Surprisingly Unsuccessful Tweets	214
6.3	Inexplicable Tweets (Qualitative Exceptions)	219
6.3.1	Unexpectedly Successful Tweets	220
6.3.2	Unexpectedly Unsuccessful Tweets	226
6.4	Discussion	229
7	Discussion and Conclusion	231
7.1	Factors that Influence Tweet Resonance	233
7.1.1	Organisational Power	234
7.1.2	Tweet Attributes (Discoverability)	237
7.1.3	Rhetorical Content	240
7.1.4	Rhetorical Style	244
7.1.5	Timing	247
7.1.6	Revised Cultural Diamond for New Media	249
7.2	Grassroots Success in New Media	251
7.2.1	Playfulness and Negativity	252
7.2.2	The Novelty of Niche Interests	253
7.3	The Measurement of Success	254
7.4	Implications	256
7.4.1	Impact of Twitter on the Referendum Result	256
7.4.2	Politics and Marketing	260
7.4.3	The Regulation of Platform Design	261
7.5	Future Directions	262
7.6	Final Reflections	264
8	Appendices	267
A	List of Hashtags Used in Search Query on 22 June 2016	267
B	Examples of Spam and Unrelated Tweets	268
C	Tweet Coding Scheme	269
D	Interview Questions Template	282
E	Participant Information Sheet & Informed Consent Form	285
F	Bivariate Crosstabulations and Stacked Bar Charts	287
G	Regression Diagnostics (Sample Level)	295
H	Linear Regression (Population Level)	299
9	Bibliography	301
9.1	Books, Book Chapters, and Journal Articles	301
9.2	Website Articles	321

LIST OF FIGURES

FIGURE 1.1: #NuggsforCarter	07
FIGURE 2.1: Cultural Diamond for New Media	20
FIGURE 2.2: Direct Retweet	24
FIGURE 2.3: Quote Retweet	24
FIGURE 3.1: 20 Features with Strongest Coefficients (SVC Issue Classifier)	75
FIGURE 3.2: Boxplots of Classifier F1 Scores for Coded Variables	77-78
FIGURE 4.1: Tweet with Numerous Hashtags	96
FIGURE 4.2: Tweet with Numerous Mentions	97
FIGURE 4.3: Tweet with Multiple Media Items	97
FIGURE 4.4: Example of Division	101
FIGURE 4.5: Example of Identification	101
FIGURE 4.6: Bar Plots of Emotion Frequencies	102
FIGURE 4.7: Example of Enthusiasm and No Theme	102
FIGURE 4.8: Example of Anxiety and Healthcare	103
FIGURE 4.9: Example of Aversion and Immigration	103
FIGURE 4.10: Example of Resolution	104
FIGURE 4.11: Example of Message Format (List)	105
FIGURE 4.12: Example of Message Format (Contrast)	105
FIGURE 4.13: Example of Playfulness and Division	106
FIGURE 4.14: Bar Plot of Mean Retweet Count for Hour of Day	109
FIGURE 4.15: Bar Plot of Binned Hour Frequencies	109
FIGURE 4.16: Bar Plot of Day Frequencies	110
FIGURE 4.17: Bar Plot of Mean Retweet Count for Position (Population)	112
FIGURE 4.18: Bar Plot of Mean Retweet and Like Counts for Position (Sample)	112
FIGURE 5.1: We Are Europe's 'Kiss of Death' Tweet	140
FIGURE 5.2: UKtoStay.EU's Fact-checking Tweet	142
FIGURE 5.3: Scientists for EU's <i>Times</i> Letter Tweet	144
FIGURE 5.4: Andrew's Street in Beverley Tweet	146
FIGURE 5.5: YEM's LGBTQ Rights Tweet	147
FIGURE 5.6: YEM's Gender Violence Tweet	147
FIGURE 5.7: Labour Leave Indiegogo Campaign Tweet (Tony Benn)	148
FIGURE 5.8: Labour Leave Indiegogo Campaign Tweet (Kate Hoey)	148
FIGURE 5.9: Labour Leave Indiegogo Campaign Tweet (Stylised Text)	149
FIGURE 5.10: Labour Leave Indiegogo Campaign Tweet (London)	149
FIGURE 5.11: Harriet's We Are Europe on BBC News Tweet	152
FIGURE 5.12: Ruthi's Molly Scott Cato Tweet	158
FIGURE 5.13: Vote Leave Pledge Card Tweet (Midday)	162
FIGURE 5.14: Vote Leave Pledge Card Tweet (Evening)	163
FIGURE 5.15: Stronger In's Daniel Craig Tweet	165
FIGURE 5.16: Paul's Map of EU Investment Tweet	166
FIGURE 5.17: Most Successful Green Party Tweet	166
FIGURE 5.18: Most Successful Healthier In Tweet	167
FIGURE 5.19: Labour Leave Camden Rally Tweet	168
FIGURE 5.20: Scientists for EU Banner Tweet	170
FIGURE 5.21: Most Successful Labour Leave Tweet	171
FIGURE 5.22: Harriet's Anger-Driven Tweet	172
FIGURE 5.23: Most Successful Better Off Out Tweet	175
FIGURE 5.24: Alberto's Turkey Tweet	177
FIGURE 5.25: UKtoStay.EU's Immigration Tweet	178

FIGURE 5.26: YEM's EU and British Science Tweet	179
FIGURE 5.27: Elise's Scotland Greener In Tweet	180
FIGURE 5.28: We Are Europe's Gove Rebuttal Tweet	183
FIGURE 5.29: Green Party Laughing at Gove Tweet	184
FIGURE 5.30: Green Party Laughing at Gove <i>i</i> Article	185
FIGURE 5.31: Sam's Anti-Carswell Tweet	186
FIGURE 5.32: Brendan's Labour Leave Stoke Rally Tweet	187
FIGURE 5.33: Vote Leave's Message of Hope Tweet	189
FIGURE 5.34: YEM Ryanair Tweet	194
FIGURE 5.35: Scientists for EU War and Bananas Tweet	196
FIGURE 5.36: Robert's NewsThump Tweet	196
FIGURE 5.37: We Are Europe's #POSTIT Campaign Tweet	197
FIGURE 6.1: 'Dear England' Tweet	204
FIGURE 6.2: 'Define #Brexit' Tweet	204
FIGURE 6.3: War and Bananas Tweet	205
FIGURE 6.4: Vote Leave Sign Tweet	205
FIGURE 6.5: Graham Norton Tweet	207
FIGURE 6.6: Billy Bragg Tweet	207
FIGURE 6.7: Taron Egerton Tweet	208
FIGURE 6.8: Richard Branson Tweet	208
FIGURE 6.9: Gary Lineker Tweet	209
FIGURE 6.10: Jo Cox Tweet	209
FIGURE 6.11: Vote Leave £350 Million Tweet	213
FIGURE 6.12: Vote Leave Pledge Tweet	213
FIGURE 6.13: Imagine If We #VoteLeave Tweet	213
FIGURE 6.14: 'White Guys in Ties' Tweet	215
FIGURE 6.15: 'Steal Your Job' Tweet	215
FIGURE 6.16: 'English at Play' Tweet	216
FIGURE 6.17: Welsh Labour Vote Remain Video Tweet I	218
FIGURE 6.18: Welsh Labour Vote Remain Video Tweet II	218
FIGURE 6.19: Andrew's Labour Leave Campaign Tweet	220
FIGURE 6.20: Brendan's Labour Leave Campaign Tweet	221
FIGURE 6.21: Better Off Out's Fishing for Leave Tweet (Cartoon)	222
FIGURE 6.22: Better Off Out's Fishing for Leave Tweet (Map)	223
FIGURE 6.23: We Are Europe's <i>Financial Times</i> Graphic Tweet	224
FIGURE 6.24: YEM Register to Vote Tweet (Successful)	225
FIGURE 6.25: YEM Register to Vote Tweet (Unsuccessful)	225
FIGURE 6.26: Oxfordshire Greens' Monty Python Tweet	226
FIGURE 6.27: Andrew's Steve Turner Quote Tweet	227
FIGURE 6.28: Robert's Lexit Article Tweet (Successful)	228
FIGURE 6.29: Robert's Lexit Article Tweet (Unsuccessful)	228
FIGURE 7.1: Revised Cultural Diamond for New Media	249

LIST OF TABLES

TABLE 2.1: Typology of Emotions	42
TABLE 3.1: Number of Tweets at Each Preprocessing Stage	66
TABLE 3.2: Interrater Reliability Scores	70
TABLE 3.3: Accuracy Measures for Best Classifiers	79
TABLE 3.4: Interviewees and Interview Information	85-86
TABLE 4.1: Descriptive Statistics for Continuous Dependent Variables	93
TABLE 4.2: Frequencies for Hashtag Count	95
TABLE 4.3: Frequencies for Mention Count	96
TABLE 4.4: Frequencies for Visual Media Types	98
TABLE 4.5: Frequencies for Hyperlink Count	98
TABLE 4.6: Frequencies for Message Framing	100
TABLE 4.7: Frequencies for Tweet Issue in Random Sample	104
TABLE 4.8: Frequencies for Style Variables	105
TABLE 4.9: Descriptive Statistics for Continuous Control Variables	107
TABLE 4.10: Frequencies for Position	112
TABLE 4.11: Odds Ratios from Five Models for Was Retweeted (Sample)	116
TABLE 4.12: Odds Ratios from Five Models for Was Liked (Sample)	117
TABLE 4.13: Odds Ratios from Five Models for Was Retweeted (Population)	124
TABLE 5.1: Summary of Pre-Interview Survey Responses	133
TABLE 6.1: Most Followed Users in Population Dataset	202
TABLE 7.1: Hypotheses	232-33

1. Introduction and Background

'To explain culture, then, is to explain why and how some ideas happen to be contagious.'
—Dan Sperber, *Explaining Culture*

Before 6 April 2017, Carter Wilkerson was just an ordinary 16-year-old boy with an extraordinary penchant for Wendy's chicken nuggets. A month later, he had won a year of free nuggets from Wendy's (via \$1,000 in gift cards), appeared on the Ellen DeGeneres Show, attracted 100,000 new Twitter followers (from a starting base of 138), inspired a \$100,000 donation from Wendy's to the Dave Thomas Foundation for Adoption, had a hashtag named after him featuring a custom Wendy's nuggets emoji (#NuggsforCarter), and launched a website branded around this hashtag that sells t-shirts¹ (Beer, 2017). The hashtag has been used by global brands and celebrities, including Amazon, Google, Microsoft, and Aaron Paul.

FIGURE 1.1: #NuggsforCarter



Most popular tweet between April 2017 and January 2019.
Retrieved 5 January 2019 from Twitter. Screenshot by author.

What catapulted Wilkerson into fame? One tweet (Figure 1.1). It attained an astonishing level of popularity overnight, and by 9 May 2017 had become the most shared

¹ <https://nuggsforcarter.com>

Twitter post of all time, with 3.42 million retweets (Koman, 2018; Beer, 2017).² The tweet's remarkable success was not at all planned—Wilkerson did not even expect a response (Koman, 2018). But the story behind its creation, as he explains on his website, reveals that it also was not completely random:

I was laying in my bed scrolling through my twitter timeline when I thought of something. I recalled tweets I have seen where kids will text their parents asking 'How many retweets for a dog?' types of questions. I decided to kick it up a notch and ask my favorite fast food restaurant the same type of question, 'How many retweets for a year of free nuggs?' After I received an outrageous goal of 18 million, I just thought of it as a silly joke to laugh at with my friends, so I replied 'Consider it done.' Then it took off and now I am making twitter history.

This anecdote illuminates that Wilkerson examined what was already popular—a specific type of question—and adapted it to his interests. Indeed, the question is part of the central appeal of the tweet, as it connects retweets directly to a tangible reward (it is cut off in the screenshot in Figure 1.1, which does not display the entire image). Wilkerson, however, does not acknowledge any other factors that might have contributed to the tweet's popularity—he merely observes that 'it took off' and the rest was history. He does not comment on its more overt content ('HELP ME PLEASE. A MAN NEEDS HIS NUGGS') or the image it contains, which focuses on Wendy's seemingly outrageous reply of 18 million and his confident response ('Consider it done'). Rhetorical factors abound in this tweet: playfulness, contrast, a forceful call to action.

These rhetorical factors and many more are at the core of this thesis, which is driven by a desire to answer an age-old question in a new media context: *What makes content work?* Why do people engage with content, and what makes content more or less engaging? The term *rhetoric* is important because it relates to persuasion; content that is merely informational or expressive is not examined. Viz., the main objective of this research is to understand how tweets that are designed to persuade operate, and why they are successful. *Persuasion* is defined within the context of Twitter itself: impacting retweets and likes as opposed to changing opinions or affecting real-life behaviour (although in this particular example, there is a direct link between the two). Retweets and likes are also defined in terms of *resonance*, which is used synonymously with success. The more likes and shares a tweet receives, the more persuasive and resonant it is. As illuminated by the example above, retweets are arguably more important because they increase the spread of a message, but likes are more transparently connected with its appeal. Both of these measures of popularity are encoded into the tweet object itself, and

² As of 6 January 2019, the record has been surpassed by a Japanese tweet (Byford, 2019).

are visible on the platform to users and non-users—which allows them to be collected and studied.

A fuller definition of success will be provided in the next chapter, which introduces the interdisciplinary theoretical framework of this thesis and provides a literature review. It should be noted that two levels of success are explored: obtaining one like or share and obtaining many more. Exemplified by Wilkerson’s tweet, the latter level is often conflated with the term *virality*, but this term will be avoided because it has network structure implications. Goel, Anderson, Hofman, and Watts (2016) propose a formal measure of the structure of virality by ‘interpolat[ing] between two conceptual extremes’: spread from person to person or via a mass broadcast (p. 180). This thesis does not investigate how popular content spreads on Twitter, but is only concerned with how popular the content is, and which features might account for its popularity.

Contested events with high stakes are an excellent arena for persuasive communication. As such, this thesis focuses its examination of rhetoric on tweets pertaining to one such event: the 2016 United Kingdom European Union membership referendum. These tweets are extremely likely to be rhetorical, and high volumes were published throughout the campaigning period. Selecting this case study also allows this thesis to contribute to the political communication literature, which has emphatically recognized social media as being indispensable to 21st-century campaigning. As Messina, Fersini, and Zammit-Lucia (2017) contend, ‘The importance of social media in contemporary politics is beyond doubt’ (p. 5). Donald Trump, the Arab Spring, and the rise of the Five Star Movement in Italy constitute compelling evidence of the effectiveness of digital campaigning. Social media allows for political messages to be dramatically amplified, and for voters to be reached directly (p. 6). Messaging is not only unmediated, but also micro-targeted. Moreover, communities on social media can be interactive, which reflects and heightens political engagement. As such, Usherwood and Wright (2017) argue that an improved understanding of the political use of social media is important for the ‘long-term maintenance of the political system’ (p. 386). They actually focus specifically on the EU referendum on Twitter, and examine the tweeting strategies of both official groups (Stronger In and Vote Leave), as well as of an alternative pro-Brexit group that fought to become the main campaign (Leave.EU). The rhetorical features they examine—message framing and thematic content—are part of the framework of this thesis.

The persuasiveness of tweets can only be understood in the context of Twitter’s design. Likes and shares cannot be examined in isolation from the user interface of the platform. To

contextualise the thesis, an overview of Twitter's affordances is provided in the next section: how do users produce and engage with content? The research questions and thesis structure are presented in Section 1.2, and the chapter ends with an account of the contributions this research has made.

1.1 Background of Twitter

Twitter was chosen as the case study platform for this thesis for two main reasons. The first concerns data access and scale. Launched in March 2006, it quickly became the most popular microblogging service: as of December 2018, there are 326 million active users each month, and 500 million tweets are sent every day (Statista, 2018; Internet Live Stats, 2018). Twitter data is well-suited for research purposes not only because of this immense volume, but also because most of the content is explicitly public and 'encoded in a single format', which facilitates large-scale data collection (Bamman, Eisenstein, & Schnoebelen, 2014, p. 139). The second reason concerns the users of the platform: because almost all politicians and journalists are active on Twitter, it is ideal for the examination of EU referendum rhetoric.

Given this preponderance of politicians and journalists, Twitter is not at all representative of the general population. One of Jungherr's (2016) conclusions from his systematic literature review of Twitter use in election campaigns is that Twitter users are more likely to be politically interested and partisan than the average person. During campaigns, they 'contribute to a collective negotiation of meaning between political elites, journalists, and other Twitter users'; this is best evidenced by 'mediated political events of high public attention', like televised debates (*ibid*). Thus, political tweets are not nearly as frivolous as they might seem—rather, they are an important communicative element of Chadwick's (2013) hybrid media system. Blank (2017) has demonstrated that Twitter users are disproportionately elite in the UK: they are 'younger, wealthier, and better educated than other Internet users, who in turn are younger, wealthier, and better educated than the offline British population' (p. 679). A Pew Research Centre study of social media platforms finds that American Twitter users have similar demographic characteristics (Greenwood, Perrin, & Duggan, 2016). Given how unrepresentative Twitter is, this thesis is not concerned with predicting the results of the EU referendum, but general reflections will be made upon the matter in Chapter 7. Persuasion is defined in the context of the platform, as noted above.

To fully understand rhetoric and resonance on Twitter, its affordances (and their evolution over time) should be explicated. First introduced by Gibson (1979), affordances can

be conceptualised as possibilities for action—we perceive objects in terms of these action potentials; indeed, our perceptions drive our actions. What are the possibilities for action on Twitter, and why are they so compelling?

Until 7 November 2017, Twitter’s defining affordance had been its 140-character limit, which is premised upon its mobile-first approach. Sending a tweet is supposed to be akin to texting, which is fast and effortless. Interestingly, at the time of its launch Twitter could be conceived of as a mobile version of Facebook (which at the time was not mobile-friendly). Like Facebook, it was meant to serve more as a social networking site, an ‘easy way to keep in touch with people in your life’ (Stone, 2009). The initial question posed to users emphasised this personal connection: ‘What are you doing?’

The most important distinction that Twitter introduced at its launch, however, was the potential for *asymmetrical connections*—users can ‘follow’ other users to receive their tweets, but the followed user has no obligation to follow back. Thus, multiple types of communication are possible: one-to-one, one-to-many, many-to-one, and many-to-many (Anstead & O’Loughlin, 2010). Tweets can be broadcast to all of a user’s followers, or directed to a specific account with the @ symbol. This openness of relationships led users to ‘share anything they wanted, completely ignoring the original question’ (Stone, 2009). What differentiates Twitter the most from traditional media is its enablement of many-to-many broadcasting, or interactive multicasting—traditional media fall into the one-to-many category (Murthy, 2013, p. 6). This allows for a far broader range of content producers and rhetorical strategies.

Twitter’s founders deliberately launched the platform as a multipurpose tool; they wanted users to influence its functions so that their needs could be served effectively. For several years, it was labelled as a ‘service in search of a user application’ (Van Dijck, 2013, p. 71). Thus, the technology was socially shaped by its users even as they adapted to its affordances. On 19 November 2009, the question changed from ‘What are you doing?’ to ‘What’s happening?’ (Stone, 2009). This reflects the far broader nature of information posted to the platform, spanning national news, celebrity gossip, and personal updates.

The character limit for tweets, combined with the fact that over three-quarters of Twitter users were on mobile devices even in 2009, creates the perfect condition for frequent updates—especially through the incorporation of links to other media (Lenhart & Fox, 2009; Zappavigna, 2011, p. 790). Instant information sharing is one of the core functions of Twitter. As Murthy (2013) notes, ‘Microblogs, as a medium, depend on the regularity of content contribution’ (p. 10). Updates unfurl in real time, and sometimes faster than those from traditional news sources; some headlines are revealed on Twitter before CNN, confirming

Twitter's role 'as a media for breaking news in a manner close to omnipresent CCTV for collective intelligence' (Kwak, Lee, Park, & Moon, 2010, p. 596). Over 80% of trending topics on the platform can be classified as 'headlines or persistent news' (p. 600). Moreover, it is common journalistic practice to quote tweets in news articles. What distinguishes microblogging from old media is that the flow of messages is a 'live stream of uninhibited, unedited, instant, short, and short-lived reactions' from a multitude of diverse perspectives (Van Dijck, 2013, p. 79). The voices of citizen journalists can mingle with those of professional reporters and casual observers in this ambient news environment, 'a media space where news is always present' (Murthy, 2013, p. 51). Such an environment did not exist before the microblog.

Twitter's ambient news environment is amplified by the possibility for information cascades that the platform allows—a faster and wider dissemination of information than anything traditional media could attain. This primarily happens through the retweet function, which allows users to acquire information from accounts they do not follow. The importance of tweets is determined collectively by users, who choose whether to retweet or not; this is the 'emergence of collective intelligence', as well as an enactment of Twitter's mission to empower individuals to share information instantly, without barriers (p. 598).

The affordances detailed in this section are effectively encapsulated by the message on Twitter's homepage: 'Follow your interests. Hear what people are talking about. Join the conversation'.³ It is quite different from Facebook's main message: 'Facebook helps you connect and share with the people in your life'.⁴ While Facebook's motto emphasizes its role as a social networking site designed to connect people to those they know, Twitter's more expansive description highlights a looser model of social networking (not restricted to friends) and immediate access to interesting information. Such features collectively contribute to its uniqueness as a communicative medium. They also allow its communicative space to be both stable (e.g., follower-followee connections) and highly dynamic (e.g., the tweets that include a trending hashtag) (Schmidt, 2013, p. 6). It must be noted, however, that the designs—and thereby affordances—of all social media platforms are constantly changing, in response to user feedback and other players in the ecosystem. More shall be said on this in the discussion chapter.

³ <https://twitter.com>

⁴ <https://facebook.com>

1.2 Research Questions and Thesis Structure

Given the versatility of tweets, it is no wonder that they have attracted the attention of a diverse range of disciplines, including computer science, linguistics, political science, communications, marketing, and network science. Many of these disciplines are integrated in the theoretical framework of this thesis. Studies focusing exclusively on the tweet began to emerge in 2007, only a year after Twitter was launched, and the number has increased exponentially ever since.

Although this thesis explores the rhetoric of tweets sent in a political context, its main research questions are more general:

1. What makes a message resonate with its audience on social media? How can the message be expressed and delivered most effectively?
2. What strategies do individuals and organisations use when they send a tweet and why?

Despite its exclusive focus on the EU referendum, this thesis is not only about social media and politics. It is also about the broader notion of culture, specifically in a digital landscape; tweets are conceived of as digital cultural objects, which will be defined in the next chapter. Its contributions, therefore, are multidisciplinary: the domains of cultural sociology, political science, and new media communication are all advanced.

To explain tweet success, an interdisciplinary framework centred around rhetoric has been established. It is comprised of nine different features: discoverability (attention-arresting features), organisational authority, message framing and structure, emotions, topical issues, resolution (action-orientation), questions, and playfulness (non-literal use of language, including sarcasm and metaphor). Broadly speaking, apart from discoverability and organisational authority, the features relate to either content or style.

A wide range of theories undergirding each of these features were tested (the features are presented in parentheses): Schudson's (1989) dimensions of the potency of a cultural object (discoverability and resolution), Chaffee and Metzger's (2001) reconceptualization of agenda setting (organizational authority), Atkinson's (1984) findings about the most resonant types and formats of political messages (message framing and structure), Marcus, Neuman, and MacKuen's (2000) concept of *affective intelligence* (emotions), Ansolabehere and Iyengar's (1995) claims about issues-based campaigning (topical issues), Ahluwalia and Burnkrant's (2004) cognitive response approach (questions), and Mulkay's (1988) remarks about humour (playfulness). Related to Atkinson's (1984) findings, the extensive literature surrounding the effectiveness of positive and negative campaigning is inconclusive. This thesis attempts to provide conclusions about the circumstances under which either technique is resonant (both are captured by the message framing feature).

A multi-method design is adopted to address the research questions from micro- and macroscopic perspectives. Approximately 9.8 million tweets related to the EU referendum were collected from Twitter during the referendum period (15 April-22 June 2016). A random sample of 2,976 was hand-coded for the presence of discoverability, content, and style features that might explain tweet resonance, and multivariate regressions were computed to assess the influence of each feature. These features were labelled in the population dataset from which the sample was drawn through semi-supervised algorithms for text classification that were trained and tested on the sample; regressions were subsequently computed with the population dataset. To address the second research question, this data-driven perspective is complemented by interview data from 22 individuals who tweeted actively during the referendum. The interviews focussed on strategies that users consciously employed to amplify the resonance of their messages. The degree to which these strategies are aligned with the results of the empirical analysis was assessed. Exceptional tweets and users derived from both datasets were qualitatively examined.

All in all, this research project took four years to complete. The first year focussed on the determination of the research questions and theoretical framework, which required extensive reading across eight different literatures spanning the humanities and social sciences. The second year was dedicated to data collection, data cleaning, content coding, and preliminary statistical analyses. Over 26 million tweets were gathered from Twitter between March and October 2016, and thousands of sample tweets were scrutinised to determine which hashtags to add to the collection query, the coding scheme, and the filtering criteria. The third year was devoted to interviews and advanced statistical and machine learning analyses. Over a period of six months, over 50 prospective interviewees were contacted, and in-person or video call interviews were conducted with 22 of them. Before each interview, all of the interviewee's tweets that appeared in the dataset were reviewed (this amounted to 10,700 tweets in total), a curated sample was selected for further discussion, and a dashboard of summary statistics and visualisations was created. After the last interview was conducted, a two-month research visit was made to the Fraunhofer Institute for Computer Graphics Research IGD in Darmstadt, Germany. During this visit, machine learning approaches to labelling the tweets for all of the manually coded features were developed, and the approaches were visualised. This allowed for population-level regressions to be conducted. In the fourth and final year, the 66 hours of interview recordings were manually transcribed and coded, the statistical analyses were refined, exceptional tweets and users were scrutinised, and all chapters were written and edited.

The next chapter presents the interdisciplinary theoretical framework of this thesis—the tweet is conceptualised as a cultural object, and factors that might influence its success are embedded in relevant literature. Chapter 3 justifies the multi-strategy methodology that was adopted to answer the research questions, and explains how tweet and interview data were collected, processed, and analysed. Chapters 4, 5, and 6 communicate the main findings. To address the first research question, a macro-level statistical overview of features that influence tweet success is provided in Chapter 4; analyses were conducted on both the random sample and population datasets. Chapter 5 is the micro-level qualitative complement to this, and focuses on the second research question through an analysis of interview data. It illuminates not only features that were not included in the theoretical framework, but also the *context* of every feature. Chapter 6 bridges the quantitative and qualitative analyses in its examination of exceptional tweets and users—those that were the most resonant in the entire dataset, and those that defied the expectations of the statistical models and the interviewees.

Insights from all three results chapters are integrated in Chapter 7 to establish the key factors that influence tweet resonance. The discoverability, content, and style factors derived from the literature reviewed in Chapter 2 will be reassessed. An updated schema for the resonance of digital cultural objects as well as basic strategies for how tweets from ‘ordinary’ users can achieve success will be proposed; the operationalisation of success will be critically revisited. The chapter will also comment broadly upon the extent to which Twitter mirrored or impacted the actual results of the referendum, along with other implications of this thesis for the intersection of politics and marketing and platform design. Chapter 7 concludes with reflections upon further directions that can be taken—how the research can be extended and enhanced through alternate approaches.

1.3 Significance of Research

The contributions of this thesis can be divided into three areas: methodological, theoretical, and practical. The methodological contribution is emblematised by its unique multi-strategy approach to content analysis. ‘Big data’ studies of text from the information science literature define content the way in which this thesis defines context: as, e.g., the number of characters in a tweet or the presence of hashtags, URLs, or @mentions (Ma, Sun, & Cong, 2012). Alternatively, to measure more ‘substantive’ aspects of content such as emotion and cognition,

they employ automatic tools such as Linguistic Inquiry and Word Count⁵ (Tsur & Rappoport, 2012). Large-scale content analyses from the communications literature blend manual coding and supervised machine learning algorithms, but are limited to specific features such as message framing and binary topics (Burscher, Odijk, Vliegenthart, de Rijke, & de Vreese, 2014; van Zoonen & van der Meer, 2016). This thesis considers a far wider range of content and style features, many of which are extremely complex. It demonstrates both the potential and limitations of machine learning classifiers for short digital text—accuracies are sufficient for simplified content features (which can be defined by keywords), but are insufficient for most style features. Moreover, the large-scale content analyses are enriched by insights from the producers of tweets, and close examinations of exceptional tweets.

The theoretical contribution of this thesis is represented by its interdisciplinary framework, which integrates the cultural sociology, political science, and new media communication literatures (linguistics, marketing, psychology, and neuroscience are also drawn upon). Rhetoric and resonance are highly multidimensional concepts, and as such require an interdisciplinary approach. Marketing is an inevitable aspect of any study of social media content; it is becoming inextricably intertwined with political communication. This thesis also contributes to theory building. Empirical evidence is provided for the foundations of a new theory of resonance on social media, and a revised schema for the resonance of digital cultural objects is presented.

The third contribution is a practical one—the findings can be directly applied to any digital marketing campaign to improve its effectiveness. In fact, multiple interviewees remarked that the interview process prompted them to think critically about their strategies, and was thus beneficial in and of itself. Two of their unsolicited follow-up comments highlight the reflexivity that the conversation activated:

I smiled when I wrote my latest tweet, remembering that it ticks all the boxes you identified of a typical tweet from me—positive, call to action ('We need to...') etc. It'll be interesting to see how it does!

The discussion we had was really helpful in letting me assess the evolution in my own Twitter postings and, in a weird way, also cathartic. And also just really interesting! So thank you again.

The practical contribution is also illuminated by the fact that the thesis design represents a meta-strategy that some of the interviewees executed: they scientifically tested their messages in order to determine how to achieve resonance, and examined both successful and unsuccessful tweets to better understand what worked and what did not. This effectively represents the objective of the research questions from an applied perspective.

⁵ <http://liwc.wpengine.com>

2. Theory and Literature Review

‘Culture implies all which gives the mind possession of its own powers, as languages to the critic, telescope to the astronomer. Culture alters the political status of an individual.’

—Ralph Waldo Emerson

Viewing the tweet as a digital cultural object, this chapter will situate Twitter within the framework of reception to propose possible explanations for what makes tweets work. Success in this thesis is defined as *popularity*—the spread and appeal of a tweet. To understand popularity, the concept of *power* is essential: cultural objects have power, or resonance, because they contain meaning for their audiences. The more powerful an object is, the greater the likelihood that it will be popular. The terms *power*, *popularity*, and *resonance* will be used interchangeably in this thesis.

Four factors account for cultural power: the producer of the object, the receiver of the object, the cultural object itself, and the social context in which it is created. In the new media landscape, receivers are producers—the audience of a tweet can tweet themselves; they are active meaning-makers as opposed to passive consumers. This shifts the dynamics of reception and allows for a far more diverse range of cultural objects to be formed. From a media theory perspective, due to digital technologies, communication systems are no longer top-down and focused on the dominant culture promoted by mass media. Rather, they are open-ended, flexible, and diverse. Because individuals are afforded ‘greater initiative’ and ‘recognised for their diversity’, a ‘creative audience’ emerges that remixes ‘the multiplicity of messages and codes it receives with its own codes and communication projects’ (Castells, 2009, p. 136). In this more democratic environment, there might be a wider range of factors that influence popularity.

The potency of a tweet is commensurate to how much it resonates with its creative audience. But what explains resonance? Why do some tweets have more power than others, and thus become more popular?

After defining the concept of culture and situating Twitter within the cultural object framework (Section 2.1), this chapter will define success on Twitter (Section 2.2), and then introduce factors that might influence the success of tweets: discoverability, organisational authority, message framing and structure, emotions, topical issues, resolution (action-orientation), questions, and playfulness (non-literal use of language, including sarcasm and metaphor) (Sections 2.3-2.5). These factors are grouped into three categories: *discoverability*, *content*, and *style*. Because organisational authority contributes to discoverability, it is discussed in the same section (2.3). Message framing, emotions, topical issues, and resolution are aspects

of content, and will be discussed in Section 2.4. Message structure, questions, and playfulness are aspects of style, and will be discussed in Section 2.5. Broadly speaking, many of the factors are components of *rhetoric*, a complex concept with three main dimensions: speaker, message, and delivery. All of the factors are embedded within literature drawn from sociology of culture, political communication, linguistics, cognitive psychology, social psychology, and advertising, and form the basis of hypotheses that are tested in this research. Indeed, the aim of this thesis is to determine their relative importance in the landscape of new media.

2.1 Tweets as Digital Cultural Objects

What is culture? Can a tweet be considered an aspect of culture?

The term *culture* is quite opaque and difficult to precisely define; anthropologists and sociologists have approached it in different ways. Edward Burnett Tylor and Clyde Kluckhohn had a high-level conception of culture as ‘the total way of life of a people... a way of thinking, feeling, and believing’ (Geertz, 1973, p. 4). This holistic definition was impossible to analyse, as it encompassed too many domains—how exactly does one measure the ‘total way of life of a people?’ (ibid). Clifford Geertz narrowed the definition by giving it a semiotic turn: culture is a symbolic form through which *meaning* is expressed (Keesing, 1974). Symbols and meanings are shared, and thereby public; they are codified and can be viewed as a collection of texts to be interpreted. Examples of culture as symbolic ‘text’ include ‘beliefs, rituals, art, and ceremonies as well as informal practices such as language, gossip, and stories’ (Swidler, 1986, p. 273).

As Geertz recognises, the ‘proper object’ of culture is ‘the informal logic of actual life’ (p. 17). Keesing (1974) also emphasizes the importance of exploring the ‘world of the familiar and immediate, the everyday and mundane’ in order to critically understand ‘how humans perceive, understand, and act’ (p. 93). *Mass culture*—represented by television, radio, books, newspapers, magazines, and other forms of traditional media—embodies the familiar and mundane, as it is designed to be accessible to many individuals. More might be learned about human behaviour from the analysis of *Harry Potter* than, say, a Virginia Woolf novel. An important characteristic of mass culture is that it is inherently *top-down*: its production is concentrated in a few individuals and organisations. Griswold (2013) notes that mass culture theory leans toward the strong culture and weak receivers perspective: recipients are passive and have to accept the meanings created by authorities (p. 90).

Social media aggregates and amplifies the familiar and mundane: platforms such as Facebook and Twitter have provided unprecedented access to the everyday thoughts and actions of individuals and groups, and in extremely large volumes. These platforms represent a new form of culture in the age of digital communication, distinguished by brevity of expression as well as the fact that anybody with an internet connection can participate. Meanings are expressed through both textual and visual symbols (e.g., hashtags, GIFs, memes). Because virtually anyone can participate, these meanings can be a reaction against authority (which is expressed through mass culture). Chaffee and Metzger (2001) recognised this even before the rise of social media platforms, and proclaimed that the decentralised internet is ‘demassifying’ mass communication (p. 369). They predicted that the issue would soon be ‘less about what media companies are doing *to* people and more about what people are doing *with* the media’ (p. 370). Social media is the ultimate form of *popular culture*—people, the non-elite majority, are ‘active makers and manipulators of meaning’ (Griswold, 2013, p. 92, p. 90).

Tweets are shared because they have meaning, and have meaning because they are shared. In fact, Twitter is one of the most public social media platforms, as tweets can even be read by non-users unless they are from a private account; less than ten percent of Twitter accounts are private (Moore, 2009). In the context of the EU referendum, tweets with relevant hashtags capture the meanings that individuals create about the EU, the UK, and other Brexit-relevant topics. Because traditional media organisations and organisations in general have a large presence on Twitter, a conflicting range of meanings is created. Twitter thus represents the collision of popular and mass culture, and is an ideal environment in which to determine what power, if any, organisational authority holds in new media—to what extent is mass culture theory applicable? The authority of organisations is elaborated in Section 2.3.

2.1.1 *Definition of a Cultural Object*


If Twitter is a manifestation of digital culture, tweets can be conceived of as cultural objects. Griswold defines a *cultural object* as ‘shared significance embodied in form... an expression of social meanings that is tangible or can be put into words’ (1987b, p. 4). *Significance* refers to the symbolisation of extended meanings that must be shared by multiple people; *embodied* indicates that cultural objects are humanly created. The meanings are manifest in a *formal pattern* that allows them to be understood and communicated by others.

Communication is inherent to the concept of a cultural object. As Berezin (1994) notes in her analysis of theatre in fascist Italy, ‘the meanings embedded in artistic practices and artistic objects are codes that permit communication’ (p. 1243). Tweets contain coded meanings that

are established by communities of users; the meanings are co-constructed and negotiated (memes are an excellent example of this). Griswold stresses that the communication of meaning must be expressible in words; like Geertz, she believes that culture is empirically observable.

Cultural objects are anchored in a particular social, political, and economic context: the social world (Griswold, 2013, p. 14). Griswold proposes the schema of the cultural diamond to depict the interrelationships among a cultural object, its producers and receivers, and the social world; all of these linkages must be considered in the assessment of what makes a cultural object succeed. In the new media age, the diamond must also take *participation* into account—‘agency and interactions become primary’, the line between producers and receivers is blurred, and the focus shifts to the *practise* of using cultural objects to communicate, persuade, or influence (p. 153; Figure 2.1). These actions are particularly relevant in the context of a political campaign, and are associated with the concept of *rhetoric*, which forms the core of the theoretical framework of this thesis. It will be fully explicated in Section 2.4.

FIGURE 2.1: Cultural Diamond for New Media



Schema for understanding the resonance of a cultural object in new media.

Retrieved from Griswold (2013), p. 153.

2.1.2 *Digital Culture*

Although Griswold herself does not explicate what new media objects are and what they look like, some communication theorists have. Manovich (2001) notes that the ‘computerisation of culture’ began in the 1990s, and defines computers as ‘universal culture carriers’ (p. 6). The shift of culture to computers and mobile devices heralded a new media revolution (p. 19). The term *new media object* encompasses a wide range of media types, ranging from digital photographs and film to games, DVDs, websites, and ‘the web as a whole’ (p. 14). Manovich further argues that ‘new media objects are cultural objects’, because digital objects represent and help construct physical referents (p. 15). Given this fundamental shift in the production and representation of culture, it is important to rethink categories and models. It is also important to recognise that new media representations of culture are as biased as traditional representations: one worldview or ‘possible system of categories’ is promoted ‘at the expense of [numerous] others’ (p. 15).

Extending Manovich’s observations to social media, Murthy (2013) contends that Twitter is a digital object that contains digital objects (tweets and retweets). Tweets are *discursive* because they communicate ideas and opinions, and are not restricted to text: they can contain hyperlinks, references to other users (@mentions), hashtags, images, and other media types. The concept of an *object* is useful because it allows for an ‘exploration of the parts that make the object’ and their cultural significance (p. 26). Various ‘parts’ of tweets that might contribute to their social meaning—and thereby success—include Twitter-specific elements (e.g., hashtags and @mentions) and content-based elements (e.g., imperatives, emotions, relevant issues). These elements will be introduced and embedded in theory in Sections 2.3-2.5. The next section will define the dependent variable of success in this study, which is closely connected to the concept of resonance. *Success* and *resonance* will be used synonymously in this thesis.

2.2 Definitions of Tweet Success

What does resonance on Twitter look like? As indicated in the introduction to this chapter, this thesis measures success through the *popularity* of a tweet. Popularity is typically represented in the information science and communication literature as *diffusion*: the spread of an idea from person to person, network to network, place to place (see, e.g., Maybaum, 2013; Vosoughi, Roy, & Aral, 2018; Romero, Meeder, & Kleinberg, 2011). If a cultural object is especially meaningful to an individual, he or she will be compelled to share it with others.

Diffusion has always been a fundamental metric of success for campaigns because it promotes *visibility*. It is important for political campaigners to spread their messages to a vast audience in order to persuade as many people as possible; the practice of *broadcasting* is a key concept in political communication (Lilleker, 2006). Broadcasting originated with mass media and has been amplified through social media, which allows campaigners to reach larger and more targeted audiences in a cost-effective manner. Donald Trump is the archetypal example of this: in his 2016 presidential campaign, he saved millions of dollars in television advertising partly because of his popularity on social media (Confessore & Yourish, 2016). Reaching as large an audience as possible was especially important for the EU referendum. The stakes were higher than they are for a general election—this was literally a once in a lifetime opportunity for citizens to influence the future direction of their country.

Many Twitter studies use diffusion as a metric of success, and focus specifically on predicting hashtag popularity (Yang, Sun, Zhang, & Mei, 2012; Sanli & Lambiotte, 2015; Tsur & Rappoport, 2012; Ma, Sun, & Cong, 2012; Weng, Flammini, Vespignani, & Menczer, 2012; Weng, Menczer, & Ahn, 2013). Although hashtag popularity and tweet popularity are not identical, they are related in that if hashtags are included in popular tweets, they will also become popular. Yang, Sun, Zhang, and Mei (2012) demonstrate that the hashtag's dual roles as topic marker and community symbol influence its adoption. Sanli and Lambiotte (2015) establish that behaviour patterns of hashtags (how often they appear) can indicate their potential virality. Two earlier studies (Tsur & Rappoport, 2012; Ma, Sun, & Cong, 2012) demonstrate that both content and context⁶ features impact hashtag propagation. Two other studies, however, have come to a different conclusion: the popularity and persistence of hashtags derive entirely from competition for users' limited attention and the structure of Twitter (Weng, Flammini, Vespignani, & Menczer, 2012), and future popularity can be predicted from the number of communities penetrated (Weng, Menczer, & Ahn, 2013). Message content does not have to be considered. As noted above, content and context features are both assessed in this thesis; moreover, a richer qualitative approach will be taken for the analysis of content (Chapter 3).⁷

⁶ Time and the topology of the network.

⁷ Tsur and Rappoport (2012) define content as hashtag length and orthography, and automatically classify words in hashtags according to emotional and cognitive dimensions using Linguistic Inquiry and Word Count (LIWC). Ma, Sun, and Cong (2012) measure content as the clarity of hashtags and the topics they can be classified into. Unlike Tsur and Rappoport, Ma et al. examine the tweets in which hashtags are embedded, but only quantitatively—they count retweets, URLs, and the @ sign. For this thesis, all elements of a tweet were closely scrutinised: its text and any visual media or URLs that it contains.

Diffusion is connected to the concept of *influence*, which is how Schudson (1989) defines the success of a cultural object: the *influence* it exerts on ‘what people think and how they act’ (p. 153). Margetts, John, Hale, and Yasseri (2015) argue that while most attempts to mobilise action on social media do not succeed, some attempts are extraordinarily influential and can even trigger revolutions. There are many studies of collective action on Twitter in particular (see, e.g., Yuce, Agarwal, Wigand, Lim, & Robinson, 2014). This thesis, however, is not concerned with determining how tweets change people’s opinions or affect offline action—indeed, as Margetts et al. have demonstrated, the mobilisations that do succeed are highly unpredictable and unsustainable; politics in the social media age is defined by turbulence. Many commentators claim that social media has limited impact on ‘real life’ results; even the most popular hashtags are just fads (Jones, 2017). What is important is that highly visible tweets *do* exert influence online, in that they trigger extensive discussions and are often cited by journalists, especially for breaking news and crisis communication (Kim, Kim, Lee, Oh, & Lee, 2015; Vis, 2013). In fact, Twitter has redefined journalistic practice because it blends official and unofficial sources and is constantly being updated; Hermida (2010) uses the concept of *ambient journalism* to capture the ‘always on’ nature of contemporary reporting. Given how short tweets are, this ‘limited’ influence is significant in and of itself.

2.2.1 *Retweets and Likes*

More specific remarks shall now be made on the Twitter features that represent the dependent variables in this study. New media platforms allow for success to be directly measured—it is visualised on the platform itself. The popularity of a tweet is reflected by the number of retweets and likes that it receives; both are part of the tweet object. Most tweets that receive multiple retweets also receive multiple likes, and vice versa, but there are qualitative differences between these two measures of success. Retweets are shares that disseminate the tweet to the followers of the retweeter, and they are published on the retweeter’s profile (in the main timeline). Thus, they are a direct measure of popularity: the more a tweet is shared, the greater its exposure.

The most comprehensive analysis of retweet motivations was conducted by Metaxas et al. (2015). Through a user survey, scrutiny of user profiles, and a meta-review of over 100 articles, they found that retweeting indicates interest, trust, and agreement for the majority of users. The disclaimer ‘Retweeting is not an endorsement’ in profile descriptions is not very common, and appears mostly in journalist accounts. It is no coincidence that the originator of the phrase ‘RT ≠ endorsement’ is a *New York Times* blog editor; he included it for professional


reasons (Warzel, 2014). Even more significantly, Metaxas et al. (2015) affirm that if hashtags are included in a tweet, the signal of support conveyed by a retweet is strengthened, particularly if the hashtags are political. This is extremely relevant to this study, as all of the tweets in the dataset contain hashtags pertaining to the EU referendum (Chapter 3).

FIGURE 2.2: Direct Retweet



Retrieved 15 December 2018 from Twitter. Screenshot by author.

FIGURE 2.3: Quote Retweet



Retrieved 15 December 2018 from Twitter. Screenshot by author.

Moreover, a distinction must be made between *direct* retweets and *quote* retweets. Introduced in 2015, quote retweets are a more effective vehicle for expressing disagreement, as they are retweets with comments attached. As Madu (2017) contends, they are often used to showcase something that the user believes to be stupid. Figure 2.2 is the pop-up that appears when the retweet button is clicked; Figure 2.3 represents what the pop-up is transformed into

if a comment is typed into the text box. As the ‘Tweet’ button in the lower right of Figure 2.3 illuminates, this comment is published along with the original tweet as a new tweet, as opposed to a retweet. Thus, quote retweets do not increase the retweet count of the shared tweet. Only *direct retweets*, which occur if the ‘Retweet’ button is clicked and no comment is added, increase the retweet count. This is what is captured by the dependent variable in the models, and its measurement will be explained in the next chapter.

A like does not spread the tweet further, but directly indicates support and appreciation due to its representation as a red heart icon (Twitter Help Centre, 2017b). Thus, whereas retweets are public and amplifying, likes are more private. Liked tweets are stored in the ‘Likes’ tab of the user’s profile, as opposed to the main tweet timeline (the default tab). It is interesting to note that the ‘like’ feature was formerly called the ‘favourite’ feature, which was represented by a yellow star icon, as it used to serve more of a bookmarking function (inspired by social sharing sites such as Delicious and Flickr). It was introduced in 2006, around the time Twitter was launched. Within seven years, it had blossomed into a multidimensional tool for a wide range of purposes beyond bookmarking: communicating thanks, ending a conversation (the equivalent of ‘LOL’), passive aggression, and endorsement, among other actions. Moreover, its usage was distinct from that of the retweet given the surge in user growth and explosion of content on Twitter; it became a way to acknowledge a tweet without ‘adding more noise to the chaotic conversation’ (Greenfield, 2013).

In part due to this complexity of usage, and in an attempt to attract new users, Twitter replaced ‘favourites’ with ‘likes’ in November 2015. The latter was becoming too confusing: as a Twitter product manager declared, ‘You might like a lot of things, but not everything can be your *favourite*’ (Kumar, 2015). As the word *favourite* is a superlative, there is an implication that not many tweets can be favoured. There is no limit to the number of tweets one can like, however—and the meaning of a ‘like’ is much more straightforward, as the word itself denotes positive sentiment. This shift in appellation and design encourages users to use the like function more and thereby to engage more with the platform, which is of commercial benefit to Twitter as it provides the company with increased volumes of user data (to be specific, it reveals what users are interested in). Moreover, there is no doubt that Twitter was influenced by like and heart functions on Facebook, Instagram, Tumblr, and other prominent social networks—these functions ‘have become a kind of universal currency of the social web’ (Newton, 2015). On a higher and more societal level, the heart is an expressive ‘universal symbol that resonates across languages, cultures, and time zones’ (Kumar, 2015). Some users

may still ‘like’ a tweet ironically, but the connotation of an endorsement is far stronger than it had been with the ‘favourite’ function.

Despite the fact that likes strongly connote endorsements, it is important to take into account the concept of *paralinguistic digital affordances* (PDAs), which is derived from adaptive structuration theory. Through focus groups and semi-structured interviews, Hayes, Carr, and Wohn (2016) explored likes on Twitter and Facebook as PDAs: cues that ‘facilitate communication and interaction without specific language associated with their messages’ (p. 173). Thus, likes can serve a phatic purpose—users might like content not because they endorse it, but rather to maintain interpersonal relationships with message posters. They can also be a signal of acknowledgement (that the message has been seen) or serve a utilitarian purpose (bookmarking useful content). This utilitarian purpose was common on Twitter, but the literal interpretation of endorsement was more common—likes are indicative of ‘good’ content.⁸ Conversely, liking as an indication of social support was more common on Facebook.

Another reason why retweets and likes are both strong measures of tweet success is because they are visible in the tweet object itself (private measures such as impressions and link clicks, which will be discussed in the next chapter, are not). The phenomenon of *social influence* explains why users tend to share or like messages that have already been shared or liked (Salganik, Dodds, & Watts, 2006). Individuals are influenced by the behaviour of others when they make decisions; the greater the strength of social influence, the more unpredictable decisions become (as the effect of more consistent personal preferences is weakened). Thus, Twitter’s affordances promote the propagation of tweets that have been liked or shared at least once; success begets further success. The effect is more pronounced for retweets: Kwak, Lee, Park, and Moon (2010) have demonstrated that any retweeted tweet reaches ‘an average of 1,000 users’, regardless of the number of followers of the user who sent the original tweet (p. 600).

Alternative measures of success on Twitter are the number of interactions a message triggers and the prominence of users who share or like a tweet (the tweet’s social *network*) (Honeycutt & Herring, 2009; Bruns, 2012). These measures are not as easy to capture methodologically as popularity. Nevertheless, they overlap on some level with popularity in that tweets are often popular because they have been shared or liked by an influential user, and popular tweets are more likely to generate discussions because of their visibility. Thus, they will be indirectly explored in this research.

⁸ Definitions of ‘good’ are extremely subjective, and encompass humour, insight, and irony.

2.2.2 *Explanations for Resonance*

Resonance, or success, is dependent on shared meanings. Viewed as cultural objects, tweets are ‘meaningful and even potentially evocative in the way physical objects are’ (Murthy, p. 27). Which meanings, then, are stronger than others? Which symbols are more powerful?

This is where reception theory comes into play: the success of a cultural object depends on the ‘cultural recipients who make their own meanings from it’ (Griswold, 2013, p. 83). It is important to note that reception—the meaning created from a cultural object—is influenced not only by the object itself or individual idiosyncrasies, but also by the social attributes of the audience. Membership in particular social groups conditions what people like, value, and recognise (p. 84). Based upon previous experiences, groups have particular *horizons of expectations* against which all cultural objects are located (p. 87). In fact, different groups may ‘exhibit systematic differences in their perceptions and interpretations of the same object’ (Griswold, 1987, p. 1081). To enhance the success of their cultural objects, producers should attempt to frame the objects so that they are aligned with the audience’s presuppositions (Griswold, 2013, p. 88).

Griswold’s (1986) cultural metaphor theory is another means through which tweet success can be explained. A cultural object is like a metaphor in that it emphasizes ‘aspects of the social world’ (e.g., shared interests) and suggests ‘an attitude toward them’ (p. 201). Success of the metaphor depends on *cultural elegance*: the ‘striking yet efficient rendering of ideas, conceptions, and concerns’ (p. 205). The metaphor must be topical, archetypal, or social: immediately relevant to contemporary issues, universally appealing, or inclusive of multiple social meanings, respectively. Through elegant presentation, cultural objects can provide ‘memorable, legitimating expression’ to private concerns felt by the audience, thereby affirming their significance and illuminating that they are commonly felt (p. 211). The tweet certainly has potential as a cultural metaphor, especially given the length constraints imposed by Twitter—in a discursive realm where 140 characters is the limit,⁹ striking yet efficient messages are heavily favoured.

The explication of resonance is not one of Griswold’s strongest points, and cultural metaphor theory does not provide a clearly operationalisable framework. Consequently, this thesis establishes a new interdisciplinary framework composed of nine empirically measurable features: discoverability, organisational authority, message type, message format, emotion, topical issues, resolution, questions, and playfulness. These features are drawn from sociology

⁹ Although it has doubled since then, this was the limit during the period of the referendum.

of culture, political science, communication studies, psychology, and linguistics. They relate to both context and content; the greater their presence, the more likely the tweet is to work. Features associated with *discoverability* are first discussed, because a tweet must be seen in order to be shared or liked (Section 2.3). These features relate to both the tweet itself and the user who sends the tweet. *Organisations* are a particular type of user and have inherent discoverability due to their reputation, resources, and power. Discoverability and organisational authority are both context variables and provide the ‘background’ of a tweet, as it were.

Rhetoric is discussed in Section 2.4 because it anchors the content-related features: message framing (*logos*), emotion (*pathos*), topical issues (the counterpart of emotion), and resolution (action-orientation). On a higher level, Section 2.4 is concerned with the tone and substance of a message, and it is followed by a section that explores features related to style: message structure, questions, and playfulness. Concluding remarks that synthesise the different features are made in Section 2.6.

2.3 Discoverability

For a tweet to have a chance of being retweeted or liked, it must first be discovered. Given the fact that over 500 million tweets are sent every day, being discovered is not easy (Internet Live Stats, 2018). The bird metaphor that inspired the brand is apropos: *twitter* denotes the noise of birds, which resembles cacophonous chatter. How can a tweet be heard over the general cacophony?

The answer to this question should be contextualised in the broader attention economy of social media. As many journalists and scholars have noted, the entire purpose of social media is to arrest users’ attention (Alter, 2017; Morgans, 2017; Greenwald, 2014; Bosker, 2016; Manson, 2014). Platforms like Twitter and Facebook are designed to keep users engaged for as long as possible, using techniques derived from the psychology of human behaviour. Twitter and Facebook feeds resemble slot machines, which work because of a principle called *intermittent variable rewards*: addictiveness is maximised if rewards for actions are sometimes given immediately, and sometimes not given at all (Harris, 2016). The next post in the feed might be immensely interesting, or immensely boring—the variable reward keeps the user scrolling in suspense.

On the one hand, it is in Twitter’s commercial interests for users to create and engage with content because attention is ‘the commodity traded on the internet’—the business model of social media is premised upon advertising, which is in turn premised upon maximising

human attention (Morgans, 2017).¹⁰ Twitter can charge more for advertising if more users engage with the platform. On the other hand, it is in the users' own interests to create attention-arresting tweets, as such tweets have a higher chance of improving account prominence.

Attention, however, is a finite resource (Wu, 2016a). The limited ability of humans to process information was recognised in the early 1970s; Simon (1971) declared that 'a wealth of information creates a poverty of attention' and called for the design of organisations to effectively allocate attention (p. 40). He also referenced the phrase *information overload*, which has come to characterise the era of digital communication. As bite-sized as they are, there are only so many tweets that one can compose, read, or engage with in a day.

Schudson (1989) uses the term *retrievability* to denote how available, accessible, and attention-arresting a cultural object is. It is one of his dimensions of potency: factors that make the object work (p. 160). A distinction must be made, however, between discoverability (synonymous with availability and accessibility) and attracting attention; the concept of retrievability conflates the two. Discoverability comes before attention can be arrested, and is exclusively caused by context features. A mixture of context and content features can increase a tweet's potential to attract attention.

Tweet and user attributes are the context features that increase the chances of a tweet being seen. Four Twitter-specific attributes that increase discoverability will be discussed in the first subsection: *hashtags*, *@mentions*, *visual media*, and *hyperlinks*. The second subsection will focus on *organisations*, a subcategory of users that deserve special attention because of their resources, power, and reputation. They are also very active on Twitter, and so are particularly relevant to this study.

2.3.1 *Tweet Attributes*

Four tweet attributes influence discoverability, and reflect the affordances of the platform: hashtags, @mentions, visual media, and hyperlinks. Both hashtags and @mentions expose a tweet to a larger audience. Hashtags are user-defined and represent topics or communities of interest. In the context of the EU referendum, many represented campaigns (e.g., #VoteLeave, #StrongerIn, #LeaveEU), the referendum in general (e.g., #Brexit, #EURef), or referendum-related events (e.g., #MarchForEurope, #EUdebate). This makes it easy for users to follow what is happening. Tweets with hashtags are more public, as they can be viewed by anybody on Twitter who searches for the hashtag (or clicks on it in another tweet). The search results

¹⁰ The commodification of attention—what Wu (2016b) calls the rise of the 'attention industry'—had its origins in newspaper advertising in the nineteenth century, and was accelerated by British propaganda during both world wars (p. 5).

page for a hashtag displays all tweets that contain it in reverse chronological order. Moreover, the most popular and time-sensitive hashtags are often trending topics, and have even more visibility—trending topics are displayed on the right sidebar of the Home timeline.

The more hashtags that are included in a tweet, the more opportunities the tweet has to be seen by different user communities. There is a norm, however, not to oversaturate tweets with hashtags: over-hashtagging is considered a social media crime because it makes the tweet look unprofessional and spamlike; the hashtags take up too much space and deprive the tweet of substantive content (Custom Fit Online, 2014). Professionalism is important because it is related to credibility and expertise, which enhance the appeal of the tweet. Digital marketers claim that best practice on Twitter is to use two hashtags, but analysis has shown that one generates the most engagement (Hutchinson, 2016).

@mentions are references to other users, and facilitate conversations on Twitter (Twitter Help Centre, 2017a). All users who are mentioned in a tweet receive a notification about it, and are thus more likely to see it. @mentions are clickable and searchable: clicking leads to the user profile; a query displays the user profile as well as all tweets that have mentioned the user. A tweet that begins with an @mention is a reply to that user, and is published in the feeds of the user's followers.¹¹

A third tweet attribute, *visual media*, influences how discoverable and attention-arresting the message is. There are three types of media that can be embedded within tweets: images, videos, and animated GIFs. All three improve the ability of a tweet to capture attention. In his study of political campaign advertising on American television, Brader (2006) observes that the use of 'music and images to stir emotions makes ads more "entertaining" or "attention-getting" ' (p. 33). Media items are more attention-arresting because they are emotionally evocative; West (2005) even claims that visuals have a greater impact than words (p. 10). Emotion is considered as a separate feature and is explicated in the next section.

The fourth tweet attribute that influences discoverability is hyperlinks to other websites or social media posts. Many of these URLs generate clickable thumbnails within the tweet that contain both preview text and an image. This is formally called a Twitter Card, and there are two types of Summary Cards for URLs: one with a 'title, description, and thumbnail', and one with an additional prominently-featured image or video (Twitter Developers, 2018). The Twitter Card thus visually expands the content of a tweet, and in a sense can be considered as a media item.

¹¹ Provided the user who receives the reply is following the sender.

In sum, the discoverability literature generates the following hypotheses about tweet attributes:

- HYPOTHESIS 1A** If hashtags are included in a tweet, it will be more successful.
- HYPOTHESIS 1B** If user mentions are included in a tweet, it will be more successful.
- HYPOTHESIS 1C** If a tweet contains visual media, it will be more successful.
- HYPOTHESIS 1D** If a tweet contains hyperlinks, it will be more successful.

2.3.2 *Organisational Authority*

Hodgson (2006) identifies organisations according to three parameters: they have ‘(a) criteria to establish their boundaries and to distinguish their members from non-members, (b) principles of sovereignty concerning who is in charge, and (c) chains of command delineating responsibilities’ (p. 8). Thus, they are both exclusive and hierarchical. Another defining feature that Hodgson recognises is *durability*—they have endured for a long time, and thus have a more powerful social position. Organisations exist in all shapes, forms, and domains: educational (schools, universities), cultural (museums), political (government, legislature), economic (banks), and professional (corporations). Due to their social position and resources, organisations not only have inherent discoverability, but also an authority that increases the impact of their messages in and of itself.

What makes organisations so powerful? The concept of *gatekeeping* helps to explain their authority. In the communications literature, gatekeeping is defined as the selection of information—which messages are featured in media outlets, and which are not? How are the messages framed? The concept derives from Lewin (1947) and was immediately applied by White (1950) to news organisations; White conceived of editors as gatekeepers. Gans (1979) and other media sociologists subsequently embedded gatekeeping in the ‘activities of organisations’ (p. 77). As Shoemaker, Vos, and Reese (2009) contend, ‘gatekeeping has continually reminded us of the importance of institutional, organisational, and professional factors in understanding the media landscape’ (p. 75). Thus, it applies not only to media organisations, but also to other organisations that influence the media (e.g., the government, corporations). Gatekeepers exert significant influence over information that is disseminated to the public, as well as the way in which the information is framed, and thereby condition what citizens view as important or interesting (Entman, 1991).

The rise of what Castells (2009) terms *mass self-communication* might challenge organisational power, and Chaffee and Metzger (2001) also recognise this—they claim that mass communication could be a fleeting 20th-century phenomenon. Traditional mass

communication (represented by radio, television, newspapers, and books) is one-directional; messages are sent from one producer to many receivers. Enabled by the internet and social media, mass self-communication is the capacity of many producers to send messages to many receivers, and the expansion of these producers to include individuals who do not represent organisations. Thus, it ‘multiplies and diversifies the entry points in the communication process’ and grants individuals ‘unprecedented autonomy’ (Castells, 2009, p. 135). Messages can be targeted to niche audiences (Chaffee & Metzger, 2001, p. 369). These audiences can be located anywhere in the world, as mass self-communication transcends national boundaries.

Castells claims that mass self-communication has ‘deeply modified the gatekeeping capacity’ of traditional media and ‘may change power relationships in the communication sphere’ (2009, p. 419, p. 422). In cyberspace, ideas that challenge corporate power, government authority, and socio-cultural norms abound (p. 421). On Twitter, they have been represented by hashtags such as #OccupyWallStreet, a protest against the corruption of financial institutions, and #Oct26Driving, a campaign for female driving rights in Saudi Arabia (Gerbaudo, 2012; Yuce et al., 2014). Nevertheless, mass media continues to exert considerable control over communication—what Castells calls *networking power*—and branding is still important ‘in the source of the message’ (2009, p. 419). There might be unequal competition between media produced by professional organisations with financial resources and the low-quality videos and gossip produced by individuals (p. 422).

Castells’ points reflect the inconclusiveness in the communications literature regarding whether or not organisations still have gatekeeping power in the new media age. Williams and Carpini (2004) declare that the internet has undermined the concept of gatekeeping entirely because of the explosion of information sources it provides; their case study is the Monica Lewinsky scandal, which political and media elites could not control. Cassidy (2016), on the other hand, claims that routine gatekeeping forces ‘exert more influence than individual gatekeeping forces’ in both print and online newspapers; new media has not shifted gatekeeping power to individuals (p. 17). He supports Singer (1997)’s assertion that while the internet is changing gatekeeping dynamics, organisations still retain power—the difference is that the role of gatekeepers has shifted from selecting news stories to interpreting them and improving their quality. Indeed, the recent phenomenon of fake news indicates that there is an explosion of misinformation online, and has impacted democratic elections (Hunt, 2016). In their reevaluation of Chaffee and Metzger’s (2001) claim that new media might herald the end of mass communication, Weimann, Weiss-Blatt, Mengistu, Tregerman, and Oren (2014) take a middle ground. They conclude that the shift of power from elite groups to a ‘greater

proportion of media users' is only partial—producers and receivers are not interchangeable, and problems such as media concentration, hegemony, and lack of democratic access persist (p. 821).

Another important dimension of organisational authority is *trust*, which is closely connected with reputation. Prior to the information revolution, individuals were confident in organisations such as the government and big businesses; the pervasive belief was that they would ensure safety as well as the quality and availability of goods and services. As Botsman (2016) contends, however, the internet has heralded a 'new world of trust', in which individuals—not organisations—are naturally trusted. She notes that 'institutional trust isn't designed for the digital age', which is characterised by collaboration and distributed networks. Distrust of politicians is especially salient; the term *post-truth* has become such a 'mainstay in political commentary' that it was selected by Oxford Dictionaries as the 2016 Word of the Year (Oxford Dictionaries, 2016).

Thus, it is unclear whether organisations still have power in the age of social media; any power they might have seems to be limited. Their financial and social resources, however, allow them to increase the discoverability and attention-arresting features of their messages, thereby granting them an inherent advantage. In an environment where there is a great deal of noise, they might still have the loudest voices. Some brief remarks shall now be made on the two organisations most relevant to this study and the specific authority that they possess.

Broadcast (or mass) media—newspapers, magazines, television, and radio—is relevant because it occupies the same space as social media. As noted in Chapter 1, Twitter represents an 'ambient news environment' that is updated in real time; many articles quote tweets or obtain breaking news updates from Twitter (Murthy, 2013). In fact, Twitter co-founder Jack Dorsey thanked journalists on the company's nine-year anniversary for using the platform to share and discover news, thereby expanding its audience and purpose (Kaufman, 2015). Consequently, broadcast media is at once a contrast to and deeply intertwined with social media.

Napoli (2003) affirms the distinctive authority of media industries: they are different from other organisations partly because of the sheer volume of 'time and money that the public spends interacting with them'—few other industries 'have the ability to affect the attitudes, beliefs, values, and behaviours of citizens across such a wide range of areas' (p. 1). The BBC is a unique exemplar of this power: established in 1922, it is the most prominent media organisation in the UK, as it is publically funded (and thus publicly owned) (Public Media Alliance, 2017). 96% of the British population watches or listens to a BBC programme; over

half of smartphone owners use the BBC news app regularly (National Union of Journalists, 2015). Even if its gatekeeping force has been reduced by new media, it is difficult to conceive of the BBC losing authority entirely, given its size and popularity.

Broadcast media has a particularly significant impact on politics; numerous scholars have discussed its role in structuring political campaigns and discussions (Page, 1996; Iyengar & Reeves, 1997). The influence of mass media on public deliberation is captured by the concept of *agenda setting*. Protess and McCombs (1991) define agenda setting as the ‘influence of the news media on the perceived salience of key political issues’ (p. 2). By making decisions regarding the curation, placement, framing, and length of content, journalists help the public understand what the most important issues are. In a sense, agenda setting is the application of gatekeeping power to news organisations.

Do new media platforms diminish the news media’s agenda setting ability by empowering individuals to increase the salience of minority topics? Chaffee and Metzger (2001) predicted that ‘the key problem for agenda setting theory will change from what issues the media tell people to think about to what issues people tell the media they want to think about’ (p. 375). Sayre, Bode, Shah, Wilcox, and Shah (2010) find only partial evidence of this in their analysis of California Proposition 8 coverage on YouTube and in mainstream media. While YouTube was used to voice unrepresented opinions, these opinions were only salient after the state court decision had already been made. In the period leading up to the decision, mainstream media coverage was dominant (and reflected on YouTube). Sayre et al. suggest that social media platforms provide the *potential* to erode the agenda setting power of traditional media, but that many efforts fail to mobilise in time.

Political organisations are also of relevance to this study because all political parties and interest groups actively campaigned for Leave or Remain. They are represented by Members of Parliament (MPs) and other government officials. Moreover, all of the campaign groups that were established represent organisations: Britain Stronger In Europe and Vote Leave were the lead campaigns selected by the Electoral Commission, but there were numerous other registered campaigns, including British Influence, Leave EU, Labour Leave, and European Movement. They were durable in the context of the referendum period because they were funded and staffed by a wide range of organisations (political, economic, and social). The financial resources of the lead campaigns were particularly large: they were given a grant of up to £600,000 and could spend up to £7m. All campaigns registered with the Electoral Commission could spend up to £700,000. The lead campaigns were also supported by broadcast media: they could disseminate messages on television and radio, and had ‘free mail

shots and free access to meeting rooms’ (BBC News, 2016b). Because of the close connection between Twitter and the news media, there is reason to believe that tweets from campaigns and their representatives had an advantage over tweets from individuals, and that the lead campaigns had a particularly strong advantage.

All in all, although some scholars have argued that the authority of organisations is greatly reduced in the era of mass self-communication, it seems that they might retain some degree of power due to their public prominence. The following hypothesis from the media and communication literature tests the extent to which this is true:

HYPOTHESIS 2 If a tweet is from an organisation, it will be more successful.

2.4 Rhetorical Content

Once a message is discovered, its content can be read. How memorable, moving, or persuasive is the tweet? How well does it capture the imagination?

Because tweets designed to persuade are the focus of this thesis, a rhetorical framework is adopted for the concept of content. Persuasion has been identified as the aim of rhetoric since Aristotle and Cicero’s treatises on the subject (Aristotle, 2007). Recognised by Schudson (1989) as *rhetorical force*, this dimension of potency deserves special attention because it is complex to operationalise, especially in the context of new media discourse. Persuasion is relevant to new media because when the line between producers and receivers is blurred, cultural objects are used to influence others (Griswold, 2013). The greater the presence of rhetorical force (the more persuasive the cultural object is), the more successful the object is.

Aristotle divided rhetoric into three components, or appeals: *logos*, *pathos*, and *ethos*. *Ethos*, the appeal to the character and trustworthiness of a speaker, is related to the user attributes measured in the concept of discoverability (Section 2.3). Users with more influence are, by and large, more credible. Given the prevalence of social bots and fake accounts on Twitter, *ethos* is an extremely important dimension of rhetoric (Ferrara et al., 2016). This section will focus on *logos*, the appeal to logic, and *pathos*, the appeal to emotion. Both relate to message content itself. Two other content strategies that are relevant for political campaigning are also discussed: *relevant issues* and *resolution*, the usage of imperatives to action. Issues are perceived to be the complement to emotions, and all campaigns are action-oriented (the ultimate imperative being one to vote).

2.4.1 *Message Framing: Identification vs. Division*

Given that many different interests were in conflict during the EU referendum campaigning period, Kenneth Burke's approach to rhetoric is relevant: *identification* is the 'key rhetorical act' (Kastely, 2004, p. 227). A speaker may persuade by *identifying common interests* to establish rapport with his audience (Burke, 1969, p. 46). Griswold (2013) also recognises the importance of identity for 'political thought and action'; persuasion is more effective when the message is aligned with the audience's established frame (p. 172).

Implicit in the theory of identification is the concept of *division*: 'identification is affirmed with earnestness precisely because there is division' (p. 22). Shared interests are also contested interests; groups define themselves in direct contrast to other groups. Emphasising distasteful aspects of these other groups is another powerful rhetorical strategy. Burke quotes from *Mein Kampf* to illustrate the power of division: Hitler observed that a great leader should always focus attention on a single enemy (1984, p. 62). This technique effectively 'strengthens the belief in one's own cause and increases one's bitterness against the attacker' (p. 63). As history attests, it proved to be terrifyingly successful for Hitler himself.

Stewart, Smith, and Denton (2007) recognise identification and division¹² as two common language strategies in social movements. Like Burke, they affirm that persuaders attempt to identify with their audiences by establishing common ground (p. 160). They also recognise that identification is inextricably linked to division; the rhetoric of division creates a we-they dichotomy (p. 164). The concept of *naming* is relevant here, as names are often used favourably to describe supporters or unfavourably to describe opponents; they have 'powerful emotional overtones' and are easily remembered (Marlin, 2002, p. 101, p. 102). An example of unfavourable naming is 'Dodgy Dave', a nickname given to David Cameron that trended on Twitter in the wake of the Panama Papers scandal.

Given the highly controversial nature of the EU referendum, Burke's conceptualisations of identification and division are an appropriate framing of rhetorical content for this thesis. However, it is not clear what they look like in discourse, and they are not explicitly connected with success. Atkinson's (1984) study of British political speeches provides a clearer operationalisation of Burkean rhetoric. Atkinson sought to determine the types of political messages that resonate the most with their audiences. His dataset consisted of transcripts and recordings of large-scale political party rallies in the United Kingdom, which were televised and broadcast over radio. The significant advantage of this method is that it

¹² Stewart et al. actually use the term *polarization*, but it is equivalent in meaning.

allowed him to scientifically examine how speakers and their audiences interact, to a much higher level of scrutiny than before. As resonance was measured by the physical sounds of clapping and cheering, Atkinson could focus his analysis on the messages that resulted in this empirical manifestation of success.

Atkinson discovered that displays of audience approval occurred ‘in response to a narrow range of very simple types of political message’: *favourable references to ‘us’* (in the form of praise or boasting) and *unfavourable references to ‘them’* (in the form of attacks or insults) (1984, p. 44). Combinations of the two were also resonant. These two types align well with Burke’s concepts of identification and division, respectively. Atkinson examined not only message *type*, but also message *format*: how are these types packaged, delivered, and structured? What is the syntax of the message? The two most dominant formats in resonant messages are *three-part lists* and *contrastive pairs* (antitheses), which are inherently binary. As message format relates to rhetorical style, it will be discussed in more detail in Section 2.5.

Atkinson’s findings hold for all political parties and speaker ideologies. They also apply to advertising slogans and newspaper editorials—other contexts in which persuasion is the aim. The narrowness of the types and formats of resonant political messages might be surprising, but is reasonable when the size and diversity of the audience is considered. The bigger and more heterogeneous the audience, the more difficult it is to communicate complex ideas and captivate attention. This is most certainly true for Twitter, whose users are from a broad range of demographics (Greenwood, Perrin, & Duggan, 2016).

A number of scholars have tested and expanded Atkinson’s framework and proven its validity in other political contexts. Grady and Potter (1985) compared televised speeches delivered by Michael Foot and Margaret Thatcher during the 1983 general election. They found that ‘outgroup criticism and ingroup praise’ was associated with 60% of applause in Foot’s speeches and 77% of applause in Thatcher’s (p. 176). Heritage and Greatbatch (1986) examined 476 political speeches at British party political conferences in 1981 (from Conservative, Labour, and Liberal parties and totalling 42 hours). They found that seven message types accounted for over 81% of applauded messages (independent of political party or speaker status). External attacks, supportive statements, and combinations of these two categories were the three most prevalent categories of message types. Contrast was the most applauded format (25% of applause events), followed by a combination of contrast and list, followed by the list. Heritage and Greatbatch also found that critical statements are ‘approximately twice as likely to elicit applause than supportive or constructive statements’ (1986, p. 121). This might be because, in the words of Simmel (1902), when there are more

people involved the number of different motives increases, and unification is only possible through negative and destructive content (1986, p. 12). Indeed, Heritage and Greatbatch observe a ‘phenomenon of crowd negativism’ (p. 122). More shall be said on the effectiveness of negative campaigning later in this subsection.

No one has tested Atkinson’s framework in the context of social media. The shift of political content to social media is as revolutionary as its shift was to television in the 1980s: while television exposed us ‘to the sights and sounds of politicians speaking in a much wider range of contexts than ever before’, social media has given us even more direct exposure to how politicians and politically-minded citizens speak, as well as what they think (Atkinson, 1984, p. 165). This thesis will determine whether Atkinson’s findings hold true on Twitter.

Burke and Atkinson’s theories and findings connect well with the political science literature on positive vs. negative campaigning. This literature is quite extensive in volume, but has mainly focused on elections in the United States (Congressional and presidential) and traditional media. The EU referendum is different from a normal election in that multiple politicians are involved in both sides; indeed, traditional party loyalties might be compromised. The campaigning literature also relies predominantly on survey methodology: traditional self-reported evidence. Social media provides directly observable evidence of thoughts, feelings, and behaviour. Moreover, studies on the effect of negative vs. positive campaigning use voter turnout and evaluation of candidates as their dependent variables (measures of success). Retweets and likes are different, as they are directly connected to the resonance of the tweet itself; they are not connected to the user’s voting preferences or real-life actions.

Various definitions of negative and positive campaigns have been proposed. Lau and Rovner (2009) claim that negative campaigns are characterised by attacks on one’s opponent (as a contrast to self-promotion). Ceron and d’Adda (2016) have a similar definition, specifying that the attack can be on their policy platform or on personality traits. Positive campaigns, on the other hand, broadcast the ‘values, virtues, and policies of the party’ (p. 1938). Tinkham and Weaver-Lariscy (1990) define positive and negative campaigning in a broader context, acknowledging messages that do not take a side. They propose four types of message strategies: emphasising one’s own qualities (positive), the opponent’s characteristics (negative), issues as opposed to personal qualities (non-personal), and both personal traits and the opponent’s traits (comparative). Kahn and Kenney (2000) take a different approach; their five categories of negative advertisements—attack, comparative, negative-trait, negative issue, and criticism—blend traits and issues, implying that it is not possible for a political message to be non-personal.

The conventional wisdom expressed by political consultants, candidates, and the media is that negative campaigning (attack politics) works (West, 2005). The success of negative campaigning is perhaps best demonstrated by Lyndon Johnson's 'Daisy' ad in the 1964 US presidential election, which begins with a little girl picking the petals of a daisy while counting them and ends with a nuclear explosion (the transition is via a countdown). It strongly suggests that his opponent, Barry Goldwater, would start a nuclear war, and is credited for triggering an era of negative political advertising in the United States (Mann, 2016). Negativity is especially important in the case of the EU referendum because Remain campaigners were accused of fearmongering (the Leave campaigners termed this strategy 'Project Fear').

To systematically assess the effectiveness of negative political campaigns, Lau, Sigelman, and Rovner (2007) conducted a meta-analysis of 111 studies. They concluded that although the literature reveals that negativity tends to be more memorable and 'stimulate knowledge about the campaign', there is no evidence that it is effective in securing more votes (p. 1176). In fact, negative campaigns sometimes have a backlash effect: 'attacking a political opponent' is more likely to 'result in lower evaluations of that candidate' than in lower 'evaluations of the target of the attacks' (Lau & Rovner, 2009, p. 296). Carraro, Gawronski, and Castelli (2010) confirm Lau et al.'s (2007) findings, and further demonstrate that both the source and the target of negative campaigns 'will experience negative outcomes'—so the effect of a negative campaign might neutralise itself (p. 462). Weaver-Lariscy and Tinkham (1996) examined advertising message strategies in US Congressional campaigns in 1982 and 1990 and concluded that going negative is a risky strategy.

Other academics provide evidence that negative campaigns can indeed be effective in certain contexts. Fridkin and Kenney (2011) analysed survey data relating to US Senate campaigns and found that negative messages that are both uncivil and relevant are the most successful in lowering evaluations of the targeted candidate. Ceron and d'Adda (2016) examined the campaign for the 2013 Italian general election on Twitter, analysing tweets from politicians and voters. They concluded that 'negative campaigns have positive effects and their impact is stronger when the attacker is meanwhile under attack' (p. 1935).

The inconclusiveness of research into the effectiveness of negative campaigns partly derives from the fact that there are different types of negativity: 'negative information does not have a uniform effect' (Kahn & Kenney, 1999, p. 887). According to Martin (2004), studies of Senate campaigns reveal that most campaign negativity fosters participation, with the exception of outright mudslinging, which depresses turnout (Kahn & Kenney, 1999). Lau and Pomper (2001) also found that most levels of negativity stimulate turnout, but that negativity

at extremely high levels seems to do the opposite. Such levels of negativity have distinct effects because citizens ‘distinguish between legitimate and tempered criticisms’ and ‘acrimonious and unjust criticisms’ (Kahn & Kenney, 1999, p. 884). The former type of criticism is useful, whereas the latter type can be alienating (particularly if it is amplified by the news media). Other scholars, however, have argued that extreme negativism can have the opposite effect. Political advertisements that violate norms of decency and good taste are likely to be noticed due to their shock value; this attention eases message comprehension and retention (Fridkin & Kenney, 2011). Negative messages that violated norms of civility in the 2006 US Senate campaigns strongly influenced ‘citizens’ evaluations of candidates’ if they were concerned with relevant topics; irrelevant messages had little impact (p. 322).

Social psychologists have posited various theories about the appeal of negativity. Negative information has ‘more weight than equally extreme or equally likely positive information’ in both cognitive processing tasks and political behaviour (Lau, 1982, p. 353; Lau, 1985). This is derived from the *figure-ground hypothesis*, which posits that negative information is more salient because it is not as frequent as positive information; it ‘stands in contrast to a predominantly positive set of norms and expectations’ (Kanouse & Hanson, 1971, p. 61). Moreover, as the ‘tangibility, importance, and proximity’ of an outcome increase, the ‘salience of its negative attributes’ also increases (p. 59). Another theory for why negativity has more weight than positivity is the *cost-orientation hypothesis*, which stipulates that negativity can be motivational: avoiding costs is preferable to approaching gains (Kanouse & Hanson, 1971). Thus, negativity should be particularly effective when the stakes are high (Lau, 1982, p. 372). On a broader, more biological level, negativity is motivational because it signals ‘the need to change or adjust... current state or activity’ (Pratto & John, 1991, p. 380).

In the context of this thesis, Burke’s terminology—*identification* and *division*—will be used to refer to positive and negative campaigning, respectively. It must be noted that the EU referendum is more complex than a traditional political campaign, which only involves a handful of candidates. The definition of *identification* can thus be expanded to encompass anything positive about anyone or anything related to either the UK (if the tweet is from a Leave campaigner) or the EU (if the tweet is from a Remain campaigner): people, organisations, policies, and beliefs. *Division* encompasses anything negative said about anyone or anything related to the opposite side (e.g., a migrant slur or anti-EU statement from a Leave campaigner; an anti-UKIP or anti-nationalist statement from a Remain campaigner).

The literature on rhetoric and political communication leads to the following hypotheses:

HYPOTHESIS 3A If a tweet contains identification rhetoric, it will be more successful.

HYPOTHESIS 3B If a tweet contains division rhetoric, it will be more successful.

2.4.2 *Emotions: Enthusiasm, Anxiety, and Aversion*

Emotion is an important dimension of rhetorical force. Given the fact that audiences have ‘particular histories and interests’, the appeal to emotion (*pathos*) is critical ‘for an argument to work rhetorically’ (Kastely, 2004, p. 224).

The word *emotion* comes from the Latin *e* and *movere*: ‘to set in motion, to motivate, or to be in a state of agitation or perturbation’ (Marcus, Neuman, & MacKuen, 2000, p. 14). Thus, the idea of stimulation is embedded within its etymology. *Pathos* means ‘suffering’ in Greek, and is the root of the word *passion*, which is associated with political emotions (ibid). These root meanings already reveal why emotions might be so powerful and persuasive.

Despite its significance in both politics and new media, emotion has been viewed as inferior to rationality in the Western tradition of political science. Moreover, emotion was deemed irrelevant to the study of politics: social scientists were concerned with issues ‘predicated on a conception of rationality’ (such as law, the nation-state, and political parties). If emotions were considered at all, they were represented as ‘deviant behaviour patterns’ (Berezin, 2002, p. 34). As Marcus, Neuman, and MacKuen (2000) put it, the exhibition of emotion about politics was commensurate to ‘psychological distraction, distortion, extremity, and unreasonableness... the detriment of sound political judgment’ (p. 2).

Aristotle was the first Western philosopher to conceptualise emotion as complementary to cognition; prior to his writings, emotion was viewed as the antithesis of reason by scholars such as Plato. Aristotle recognised that reasoned persuasion and emotional response were mutually influential (Charteris-Black, 2005, p. 13). In the past two decades, social scientists have begun to accept the importance of emotion. Berezin (2002) declares that emotions ‘are so constitutive of social, and by extension, political life that they cannot be ignored’ (p. 36). Chilton (2004) proposes a neurological connection between political discourse and emotion, which is situated in the brain, and enumerates politically relevant feelings: ‘territorial belonging and identity (“home”), love of family, fear of intruders and unknown people’ (p. 204). These feelings all apply to the EU referendum.

The acceptance—and, indeed, embrace—of the role of emotion in the social sciences has been evidenced by an explosion of studies about its importance for politics, social movements, and diplomacy (Neuman, Marcus, MacKuen, & Crigler, 2007; Papacharissi, 2015; Hall, 2015; Goodwin, Jasper, & Polletta, 2001). Emotional appeals make political messages

resonate more with voters: they increase not only discoverability, but also acceptance of message content (Brader, 2006). Marcus, Neuman, and MacKuen (2000) have established a theory of *affective intelligence* that will be used in the theoretical framework and explained later in this subsection.

Emotion plays an important role not only in politics, but also in new media. Platforms like Twitter, Instagram, and Facebook encourage users to express their feelings and opinions concerning a wide range of subjects; posts are replete with judgments and evaluations. Indeed, sentiment analysis is often performed on large volumes of tweets to determine how the public feels about certain events (Wang, Wei, Liu, Zhou, & Zhang, 2011). Recent scholarship has established the prevalence of emotion in socio-political contexts on Twitter. Through an analysis of #egypt during the Egyptian revolution of 2011, Papacharissi and Oliveira (2012) discovered that tweets combined emotion and opinion in a deeply subjective manner (p. 277). What's more, the affective dimension was functional, for it sustained 'involvement, connection, and cohesion' (p. 279). Papacharissi and Oliveira do not differentiate between different types of emotion, however—are some more resonant than others?

Berger (2013) presents a typology of emotions in his examination of the contagiousness of content: what makes some content more likely to be shared, to 'catch on', than others? Emotion is one of the six principles that he claims can affect the contagiousness of content, as it 'helps deepen our social connection' (p. 105). The concept of contagiousness is measured by retweets, one of the dependent variables in this study, as they literally enable a tweet to spread.

TABLE 2.1: Typology of Emotions

	HIGH AROUSAL	LOW AROUSAL
POSITIVE	Awe, excitement, amusement (humour)	Contentment
NEGATIVE	Anger, anxiety, disgust ¹³	Sadness

Retrieved from Berger (2013), p. 109.

Berger not only claims that content containing emotions is more likely to be popular than content that does not, but also classifies emotions along two dimensions in order to explain which are more likely to be contagious (Table 2.1). The first dimension is *valence*, captured by the traditional continuum of positivity vs. negativity: how pleasant is the emotion? The second dimension is *arousal*, an excitatory state that mobilizes individuals and thus

¹³ Not officially in Berger's table, but mentioned in his chapter on emotions.

increases information sharing and the capturing of attention (Berger, 2011, p. 891). How intense is the emotion? Berger and Milkman (2012) found that although positive content is more viral than negative content, content evoking high-arousal emotions is more viral regardless of its valence. Thus, tweets expressing anger might be more successful than tweets expressing sadness—even though anger and sadness have the same valence (they are both negative), anger has higher arousal.

As logical as it appears, there are three issues with the application of Berger's typology of emotions to Twitter data concerning the EU referendum. Firstly, Berger's framework was tested on an analysis of long-form journalism: *New York Times* articles published over a period of three months. Tweets are significantly shorter, and thus it is significantly more difficult to detect the presence of emotions in them or to differentiate between emotions of a similar valence (e.g., contentment and excitement, anger and disgust). The second issue is that Berger's framework is designed to be applicable to content in general—everything from news to products to fundraising videos (2013, p. 21). It is not tailored to political content in particular; as Berezin (2002) recognised, there are emotions that are more specific to politics. The third and final issue with Berger's typology is that it is too binary, and does not recognise that people can express multiple emotions at once.

Claiming that emotions are centrally important to political behaviour, Marcus, Neuman, and MacKuen (2000) have formulated a theory of *affective intelligence* that is premised upon both neuroscience and politics, and acknowledges the coexistence of emotions (p. 129). They posit that there are two systems of emotional appraisal—*disposition*, reliance on habits, and *surveillance*, explicit consideration of the options. These systems allow us to make political decisions and explain why we make them. The disposition system is driven by *enthusiasm* (which encompasses pride, happiness, and hope), and the surveillance system by *anxiety* (fear, worry, unease). Brader (2006) also recognises anxiety and enthusiasm as the fundamental emotions in political campaigning. It is important to note that they are not oppositional, but rather can coexist. This is why simpler valence models of emotion are not as useful or accurate as 'rich (dimensional or discrete) accounts of emotions' (p. 53).

Marcus et al. (2000) also recognise a third emotional dimension that marks 'strong normative disapproval': *aversion*, a combination of anger, hatred, and disgust (p. 158). Although they are both traditionally classified as 'negative' emotions with regard to valence, aversion is fundamentally distinct from anxiety. It is governed by the disposition system, which is concerned with learned associations and habit: 'who and what we loathe is part of the familiar world' (ibid). Anxiety, on the other hand, is connected to the surveillance system, which

identifies ‘novelty and threat’ as opposed to ‘familiar domesticated stimuli’ (p. 161). Thus, the conception of negativity is not useful in the study of political emotions; anxiety and aversion are as different from each other as each is different from enthusiasm. Aversion is not a primary emotional process, but it is activated when there is a ‘powerful and polarised’ political stimulus (p. 173). Brexit is undoubtedly such a stimulus, and so aversion will be considered in this thesis.

Of all the emotions encompassed by *enthusiasm*, *pride* is the most relevant to the context of the EU referendum—in politics, pride appeals have a special association with ‘group, especially national, identities’ (Brader, 2006, p. 151). It is often treated by appraisal theorists as a distinct emotion because of its ‘retrospective and egocentric’ nature (ibid). However, due to measurement issues that will be explained in the next chapter, *pride* is not considered as a separate variable in this study.

There are theoretical arguments for why enthusiasm should be more effective than anxiety, and for why aversion can be effective, but empirical evidence so far has been inconclusive. Expressive theories claim that people express emotions in order to improve how they are feeling—affect is a coping mechanism as opposed to an irrational influence on behaviour (Roseman, Abelson, & Ewing, 1986). If emotions are oriented toward the goal of feeling better, hopeful appeals should be more successful than fearful appeals. On the other hand, anger (a component of aversion) might be appealing because it is associated with ‘power and moral legitimacy’, cognitions that presage victory (p. 293).

These theories have not been sufficiently tested, and so the relative effectiveness of enthusiasm, anxiety, and aversion is unclear. This thesis will help illuminate the differences (if any) in potency of the three fundamental political emotions. The underlying point, however, is that the presence of *any* emotion is more effective than the absence of emotion, as Berger (2013) contends. Thus, the political science, psychology, and neuroscience literature results in the following hypotheses:

HYPOTHESIS 4A If a tweet contains enthusiasm, it will be more successful.

HYPOTHESIS 4B If a tweet contains anxiety, it will be more successful.

HYPOTHESIS 4C If a tweet contains aversion, it will be more successful.

2.4.3 *Relevant Issues*

In the political science literature, the ‘rational’ counterpart to emotion is issues: substantive themes that relate to citizens’ interests and concerns. Political messages can, of course, contain both themes and emotional appeals; indeed, affect often enhances the power of a substantive message. As Ansolabehere and Iyengar (1995) observe, campaign advertising works to the

extent that it resonates with ‘viewers’ concerns and political beliefs’ (p. 96). In fact, issues distinguish political parties and give them meaning; they are inherently partisan (p. 65).

A more applied construction of the emotion vs. reason dichotomy in the political campaign literature is *traits vs. issues*: does the message focus on personal characteristics (strengths of the candidate, weaknesses of the opponent), or does it focus on salient topics that relate to domestic or international policy? Recognised by Ansolabehere and Iyengar as image advertising vs. issue advertising, this dichotomy has long been explored by political scientists. The literature on positive and negative campaigning reviewed in Subsection 2.4.1 highlights some of the studies that have distinguished between trait and issue-focused strategies (Ceron & d’Adda, 2016; Tinkham & Weaver-Lariscy, 1990; Kahn & Kenney, 2000). Most of these studies examine negative campaigning, given its increased prominence over the past few decades. In his analysis of US presidential election campaigns between 1952 and 2004, West (2005) discovered that substantive appeals often appeared in negative commercials; Fridkin and Kenney (2011) also contend that negative advertisements focus more on issues than on traits.

The correlation between issue-oriented messages and negative campaigning is not surprising, in light of evidence that it works. In fact, West (2005) found that negative advertising is only effective if it focuses on issues—attacking an opponent’s character or background is risky for the attacker, due to the backlash effect defined in Subsection 2.4.1 (Lau & Rovner, 2009). Brooks and Geer (2007) suggest that trait-based negative advertisements that are uncivil are most likely to anger the public. In her study of news coverage of negative campaigns, Min (2004) concluded that policy-based attacks reported in the media ‘slightly stimulate voter turnout’, whereas personality attacks significantly depress turnout (p. 95). One of the only studies that reported the opposite result did not find strong evidence for it: issue negativity has a ‘*slightly* demobilising effect’, while trait negativity ‘appears *slightly* to energize the electorate’ (Finkel & Geer, 1998, p. 590; emphasis added).

Issue-oriented messages are not only effective in negative campaigns, but are also indispensable to overall strategy. As West (2005) notes, one of the ‘two key campaign decisions’ is which topics to mention in advertisements (p. 114). To increase the success of the message, *relevant* topics should be included. Relevance is heavily conditioned by the social world that Griswold represents in her cultural diamond (Figure 2.1)—the audience is shaped by societal factors, which encompass ‘traditions and history’ as well as contemporary ‘interests and concerns’ (Rupke & Blank, 2009, p. 136). In the wake of a recession, economic issues become prominent; during times of war, foreign policy takes centre stage. In the context of this thesis,

issues that are relevant to the EU referendum will be considered: i.e., themes that are most salient due to the history of Britain's relationship with Europe, as well as more recent topics of public interest.

Given the fact that issue-oriented messages represent 'matters that are on the public mind' because they directly impact citizens' lives, it is no surprise that they attract attention and are likely to be remembered—and thus shared (West, 2005, p. 68, p. 66). As Ansolabehere and Iyengar (1995) declare, '*effective advertising speaks directly to the voter's concerns*' (p. 83; emphasis in original). The mere mention of a timely and important issue strengthens partisan support. It must be noted, however, that the relevance of issues is not necessarily universal—some issues appeal more to certain groups than to others; voters have personal predispositions that correlate with the political party that they support. Thus, another tenet of effective advertising is directing voter attention to '*issues that the voter already believes to be priorities*' (p. 88).

Some scholars might argue that substantive themes are extremely difficult, if not impossible, to capture in tweets. Such arguments were made when 30-second campaign commercials were introduced on television—a significant contrast to political speeches or pamphlets. Brevity of form, however, does not prevent 'the transmission of substantive information' (Ansolabehere & Iyengar, 1995, p. 59). In fact, length constraints might *strengthen* the information value of an advertisement, because shorter messages are more attention-arresting (p. 60). Moreover, tweets can contain hyperlinks or multimedia. As noted in Section 2.3, such items increase a tweet's discoverability; they also often contain much of the tweet's thematic content. Through both multimedia and text, the most powerful political messages rely upon 'widely shared cultural symbols and scripts' to remind voters of what they already know about salient issues (p. 60). Such symbols are often in the form of stereotypes, visual or verbal shortcuts that convey oversimplified generalisations about particular groups (West, 2005). These shortcuts have a wealth of social, cultural, and political connotations—they may be short, but they are pregnant with meaning.

What are the prominent themes that relate to the EU referendum? The five included in this study are the economy, immigration, sovereignty, security, and healthcare. Immigration and the economy have been recognised as the most divisive issues; a NatCen Social Research report calls them the 'two poles of the referendum debate' (Curtice, 2016). Remain campaigners were most concerned with economic arguments, whereas Leave campaigners were preoccupied with immigration. Immigration has long been a salient issue in British politics. Moreover, free movement is one of the founding principles of the EU, due to perceived economic benefits—so the two issues are intertwined (Portes, 2016). Sovereignty,

however, was arguably the ‘principle at the heart of the Brexit argument’, and it was expressed by Vote Leave’s main tagline: Take Control (Williams, 2016).

Security became a prominent issue after terrorist attacks in Paris and Brussels in November 2015 and March 2016, respectively; the number of refugees in Europe dramatically increased in 2015 and 2016 due to escalating conflicts in the Middle East. From a more historical perspective, the EU was founded in the aftermath of two world wars, with the goal of establishing and preserving peace (European Union, 2017). Finally, healthcare is a salient theme because the National Health Service (NHS) is one of the most important institutions in the UK, given the range of free services that it has provided since 1948. It is such a fundamental aspect of British society that it is one of the ‘top reasons people feel proud to be British’ (Smith, 2015). Because the NHS is understaffed and underfunded (total annual expenses are £115 billion), Vote Leave made a prominent campaign pledge: ‘Let’s give our NHS the £350 million the EU takes every week’ (Vote Leave, 2016). This became one of the most-quoted slogans during the referendum period.

As noted above, different issues matter more or less to different groups; there is no objective way to rank the relevance of these five themes. It is clear, however, that all of them were extremely timely and important in the context of the EU referendum. Thus, the following hypothesis can be derived from the political science literature:

HYPOTHESIS 5 If a tweet contains relevant issues, it will be more successful.

2.4.4 Resolution

Imperatives are an operationalisation of Schudson’s (1989) concept of *resolution*, which is another dimension of the potency of a cultural object that is applicable to this thesis. How does the tweet relate to action? Does it tell the audience how to respond? As Rupke and Blank (2009) note, ‘symbols are more powerful when they support a popular action or a commonly pursued goal’ (p. 136). Thus, tweets that are action-oriented might be more powerful than tweets that are not. Resolution is especially relevant in the context of campaigns, as their main objective is to persuade citizens to vote. Given the significant implications of the EU referendum for Britain, Europe, and the rest of the world, the ‘once in a generation’ nature of the event, and how uncertain the outcome would be, many tweets called for actions related to voting to Remain or Leave: from watching a video or reading a blog post to attending a rally or completing the voter registration form.

What does the expression of action look like? Linguistically, persuasion is executed through *imperative verbs* such as ‘do’, ‘don’t’, ‘attend’, ‘read’, and ‘vote’. These verbs tell the

audience what to do, leaving little room for questions or discussion (Grammarly, 2017). As Schudson (1989) notes, advertisements are examples of high-resolution texts—they are replete with imperatives, as their primary goal is to persuade audiences to make a purchase. In marketing terms, such imperatives are ‘calls to action’ (CTAs): the part of the advertisement that directly implores users to ‘perform a certain act immediately’ (Marketing Terms, 2017). On websites, CTAs are often placed on buttons; one of the metrics of success in digital marketing is the number of clicks the CTA button receives (Boag, 2017). As the above definition of CTAs indicates, immediacy is essential—time-sensitive words such as ‘must’, ‘please’, ‘ASAP’, and ‘now’ can be used to increase the urgency of a message, and thus the strength of its resolution (Wishpond, 2017).

Digital marketing techniques are relevant not just for marketers, but also for everyone on the internet. Marwick (2013) contends that digital technologies teach users ‘to be good corporate citizens... by harnessing marketing techniques to boost attention and visibility’ (p. 12). The design of social media sites amplifies the effectiveness of calls to action, as ‘share’ and ‘like’ buttons are attached to all posts. On Twitter, the most prominent action is a retweet, which often encourages social action or expresses collective identity; as such, it can be a political act (boyd, Golder, & Lotan, 2010). Enough tweets associated with movements and protests such as #YesAllWomen, #BlackLivesMatter, and #BringBackOurGirls were retweeted to make them trending topics on Twitter. The word ‘retweet’ in and of itself is an effective CTA: social media scientist Dan Zarrella (2013) has demonstrated that tweets containing the phrases ‘please retweet’ or ‘please RT’ are more likely to get retweeted.¹⁴

Critics have commented on the fact that calls for social action online amount to ‘slacktivism’, or armchair activism. Retweeting a post on Twitter or changing one’s profile picture on Facebook to express solidarity are incredibly superficial actions; campaigns might become viral and highly visible, but ultimately they are just transient fads (Jones, 2017). Regardless of whether this might be true, the focus of this thesis is the success of the message itself in the context of social media, as measured by retweets and likes. The effect of a tweet on offline action (in this case, voting to leave or remain in the EU) is not being measured, although broad reflections will be made about this in Chapter 7.

The linguistics and advertising literature leads to the following hypothesis:

HYPOTHESIS 6 If a tweet contains resolution, it will be more successful.

¹⁴ 50% of tweets with ‘please retweet’ are retweeted; 40% of tweets with ‘please RT’ are retweeted.

2.5 Rhetorical Style

The previous section explored aspects of message *content* that might contribute to a tweet's success: the political framing of the message, and whether it contains emotional appeals, relevant issues, or imperatives to action. When persuasion is the main objective, however, 'the manner in which a statement is expressed may be more important than its propositional content' (McQuarrie & Mick, 1996). This section focuses on *rhetorical style*: how is the message expressed? Does the tweet contain a list or a contrast? Does it pose a question? Does it play with language through rhetorical figures? As these questions indicate, three style features are part of the theoretical framework of this thesis: *message format*, *questions*, and *playfulness*.

2.5.1 Message Format: Lists and Contrasts

As mentioned in Section 2.4, Atkinson (1984) examined not only message *type*, but also message *format*: how are these types packaged, delivered, and structured? The two most dominant formats in resonant messages are *three-part lists* and *contrastive pairs*. These formats are particularly effective for speeches because they signal to the audience when the speaker will come to the end of a phrase, so that they know when they can clap. This signalling mechanism is not relevant for tweets because they are not delivered orally, but lists and contrasts have inherent rhetorical force as well. Research into everyday conversations has shown that three-part lists are extremely common. They have 'an air of unity or completeness about them', and are excellent vehicles for praise or criticism because 'listing similar items can work to strengthen, underline, or amplify almost any kind of message' (Atkinson, 1984, p. 57, p. 60). In the context of Twitter, three-part lists might be more common than longer lists because of the 140-character constraint.

Well-known to classical orators, contrasts (also known as antitheses) inherently promote *us vs. them* rhetoric. They are manifest when there is an explicit contrast in words, meaning, or both: the juxtaposition of two words or ideas that are opposites or strikingly dissimilar, usually in parallel form. Contrasts often contain a puzzle or joke and then a punch line. As Atkinson observes, in adversarial contexts they 'provide an economical way of formulating rival positions' (1984, p. 159). The EU referendum was certainly an adversarial context; Remain and Leave campaigners were highly polarised and antagonistic.

Both lists and contrasts are arresting because they 'provide a way of formulating the gist of a message which is both brief and complete in itself' (Atkinson, 1984, p. 158). The arguments and ideas contained in the message are often highly complex, but presented in simplified form through these two formats (a two-word sentence can constitute a contrast if

the words are opposed in meaning). Brevity is very important for rhetoric and politics. Jamieson (1988) contended that in the era of mass communication, ‘concise, memorable phrasemaking’ is demanded of rhetors—only those who can ‘synthesize an issue in a clear, concise, dramatic statement’ in under 35 seconds will actually be heard (pp. 90-91). Her contention is all the more relevant now, in the era of social media. The tweet is the modern day synoptic phrase.

Another important reason why lists and contrasts are resonant is that while they are short and snappy, they also provide their audience with ‘food for thought’: effort has to be exerted to understand them. Speakers who allow their audiences to arrive at their own conclusions demonstrate ‘greater respect for their intelligence and powers of reasoning’ (Atkinson, 1984, p. 160, p. 161).

It must be noted that on Twitter, contrasts are not restricted to words: they can also be the juxtaposition of visual media and text, or of two visual media items. Moreover, many images embedded in tweets contain text that contains contrasts.

Thus, the political communication literature leads to the following hypotheses:

HYPOTHESIS 7A If a tweet contains a list, it will be more successful.

HYPOTHESIS 7B If a tweet contains a contrast, it will be more successful.

2.5.2 Questions

Questions provoke attention by their very nature: they invite the audience to respond, to engage with the communicator. Black (1992) observes that they are socially mandating: ignoring a question ‘is to risk giving offense’ (p. 2). As such, questions have an inherent degree of *resolution*.

There are many different types of questions. Some are adapted for specific contexts, like the interview or the courtroom. Leading questions, for example, strongly suggest a specific response and are deployed during the cross-examination of witnesses during a trial. Probing questions are used for following up on points in interviews, when the respondent has not provided sufficient detail.

The most famous type is the rhetorical question, broadly defined as a question that should not be answered because the answer is implicit. Rhetorical questions are used to convey opinions as opposed to seek new information; these opinions are often challenges to other opinions, yet have the guise of neutrality because of the interrogative form (Koshik, 2005). Since Greek and Roman times, speakers have used this technique to persuade audiences; its

effectiveness is well-documented (Frank, 1990, p. 723). It is especially common in contested situations: debates, campaigns, and polemical interviews, to name a few.¹⁵

The effectiveness of rhetorical questions is embedded within their form: asking a question is persuasive in and of itself, as it ‘solicits assent to a proposition’ from the listener, who has a ‘social obligation to respond’ (Black, 1992, p. 2). Seeking agreement has long been recognised as advantageous in debate (Zillman, 1972). Moreover, rhetorical questions cloak opinion as fact: they appear ‘to be making a truth claim, and implying that that truth is known to the recipient’ (Koshik, 2005, p. 152). Because they require no response, they have more force and are difficult to contradict. In the psychology and advertising literature, the leading explanation for the persuasiveness of rhetorical questions is the cognitive response approach, which stipulates that questions focus attention on message arguments and encourage recipients to reflect and elaborate (Ahluwalia & Burnkrant, 2004, p. 27).

Zillman (1972) directly compares rhetorical questions with statement versions of these questions in an experimental study, and concludes that the questions are more persuasive, especially when the subject was initially opposed to its message. Mothersbaugh, Huhmann, and Franke (2002) demonstrate that rhetorical questions improve the persuasiveness of advertisements by increasing consumer processing of their messages.

It is important to note, however, that rhetorical questions can sometimes backfire. They are only effective when their usage is salient and when evaluation of the source of the question is positive. Source evaluations are derived from knowledge of the speaker and/or strength of message content (Ahluwalia & Burnkrant, 2004). In Aristotle’s framework, the former is captured by *ethos*, the latter by *logos*. Petty, Cacioppo, and Heesacker (1981) contend that relevance to the audience is another factor: rhetorical questions are most effective when messages are of low relevance, and disrupt thought when messages are of high relevance.

Because rhetorical questions are dependent on context, speaker intention, and audience response, they are difficult to recognise. Given their brevity, tweets provide very little context, and speaker intention is almost impossible to detect. Even if more context were provided, identifying a question as rhetorical necessarily involves a subjective judgment. Frank (1990) demonstrates that they have high pragmatic complexity; even linguists cannot provide a precise definition of what makes a question rhetorical (p. 723). Thus, the broad category of *questions* will be explored as a potential factor for tweet success.

The rhetoric and psychology literature generates the following hypothesis:

¹⁵ In his analysis of dialogues drawn from a wide range of contexts, Frank (1990) notes that rhetorical questions featured most frequently in conflicts.

2.5.3 Playfulness

According to Lieberman (1977), playfulness is composed of spontaneity, joy, and sense of humour. It is the dimension of style that is especially relevant to computer-mediated discourse, due to the constraints introduced by internet technology on communication: context cues are diminished, there is a greater distance between speakers and their audience, and the boundaries between private and public are collapsed (Papacharissi, 2012). Other factors that contributed to online playfulness in the early days of the World Wide Web were the interactive nature of websites, hacker culture, and increased anonymity (Danet, 2001). In the broader context of the attention economy of social media, users must add variety to make their messages more noticeable (Su, 2007, p. 67).

Defined by its 140-character message limit, Twitter is particularly conducive to playful language. It is no surprise that there is an explosion of parody accounts on the platform that provide commentary and entertainment, some of which have accumulated tens of thousands of followers (Moreau, 2017). In her analysis of self-presentation on Twitter through trending hashtags, Papacharissi (2012) contends that play is a prevalent performative strategy through which users ‘overcome expressive restrictions’, and notes that shorter messages are more playful (p. 1998).

Broadly speaking, playfulness involves ‘rearranging norms for expression’ and contravening conversational conventions (Papacharissi, 2012, p. 1999). There is a disjunction between intended and literal meanings, and the general effect is humour or wit (Palmer, 1994). The goal of the humorous mode is to render impossible ‘any firm distinction between real and unreal, between actual and non-actual’ (Mulkay, 1988, p. 46). Humour undermines the ‘expectations of serious discourse’, in which meanings are interpreted literally (p. 42). It has been positively associated with both self-presentation and group solidarity, which is enhanced through the common understanding of humorous discourse. As Baym (1995) observes, humour ‘also expresses the complexities and problems of the social world’. For while humour is by definition non-serious, informal humour is often perceived as a means of ‘conveying serious information without appearing to do so’ (Mulkay, 1988, p. 69).

Humour is inherently persuasive because of its connection to emotion as well as to shared assumptions; indeed, such assumptions are the root cause of laughter (Leith, 2011). Many scholars have recognised the effectiveness of humour online. Messages that contain humour are the most appreciated in computer-mediated environments (Herring, 1999). Baym

(1995) demonstrates that humour facilitates ‘the formation of group identity’ through her analysis of a Usenet group. She found that humour distinguishes ‘quality posts, individualized posters, and is a central appeal of the newsgroup’. Because it is performative, it has the power to boost a user’s prestige, status, and control.

Interestingly, most of the humour in the Usenet group was critical. This aligns with Lieberman’s (1977) observation that humour can be positive or negative. Political language is especially conducive to hostile humour because it is adversarial in nature and contains interpretative diversity; opposing groups deconstruct each other’s discourse by representing contrary worldviews as ‘unreliable, illusory, or not to be taken seriously’ (Mulkay, 1988, p. 207). Given the polarised nature of the EU referendum debate, this is even more reason to expect that many referendum-related tweets will contain humour.

Playfulness is difficult to operationalise, as it is grounded upon shared social, cultural, and historical knowledge (Palmer, 1994). The most tangible strategies through which playfulness is achieved are represented by rhetorical figures, which are also called stylistic devices, figures of speech, or figurative language (as opposed to literal language). As noted above, playfulness and humour are defined by incongruity and circumvention of norms (Baym, 1995). Palmer (1994) contends that the transgression of normalcy is the key component of humour. Interpretive multiplicity is also an integral component; participants in humorous discourse inhabit ‘a world in which whatever is said and done necessarily has more than one meaning’ (Mulkay, 1988, p. 4). Incompatible interpretive frames are juxtaposed, resulting in ‘a bifurcated logical process’ (Palmer, 1994, p. 96). Other features of the humorous mode are ambiguity, inconsistency, and contradiction (Mulkay, 1988, p. 26).

Rhetorical figures are ‘artful deviations in language use designed to enhance communication’ (Corbett, 1999; McQuarrie & Mick, 1996; Mothersbaugh, Huhmann, & Franke, 2002, p. 589). They represent different ways to package message content—the style of delivery, as opposed to the core idea (which remains unchanged through the application of different figures). Deviation is from the general expectations that are incurred in communication encounters, captured by the cooperative principle in pragmatics (Grice, 1989).¹⁶ Thus, rhetorical figures create incongruity. This dovetails nicely with the core components of playfulness and humour, and allows for a close association to be drawn. Moreover, the deviation is performed in an *artful* way—wit and wordplay are often involved;

¹⁶ Grice proposed four conversational maxims: *quantity* (giving enough information, but not too much), *quality* (telling the truth), *relation* (being relevant), and *manner* (being clear, succinct, and orderly).

rhetorical devices have become synonymous with *literary devices* given their popularity in poetry and literature.

Hundreds of rhetorical figures have been catalogued since classical times, ranging from more common devices such as hyperbole and metaphor to more esoteric devices such as epanorthosis and anadiplosis. McQuarrie and Mick (1996) have established a framework that categorises the astounding variety of devices, and allows for theoretical predictions to be made about their effectiveness in different advertising contexts. Drawing from Corbett (1999) and Leech (1973), they distinguish between *schemes* and *tropes*: schemes rearrange the surface of the text (through repetition of sounds, words, or phrase structure or the reversal of syntax), whereas tropes affect the semantics of the message by substituting one meaning for another (saying ‘something other than what is meant’) or destabilising meaning entirely (implying that ‘one means more than [what] is said’) (p. 433). Examples of schemes include rhyme, alliteration, and anaphora; examples of tropes include hyperbole, metaphor, pun, and irony. Only tropes will be considered in this study, as they exhibit a greater degree of deviation from the norm.

Why are rhetorical figures effective? Since they deviate from conventional expressions, they naturally arrest attention. Because they are artfully arranged, they can yield a ‘pleasure of the text’ when successfully processed—an aesthetic reward (Barthes, 1990). Berlyne (1971) also argues that incongruity can generate ‘a pleasurable degree of arousal’ (McQuarrie & Mick, 1996, p. 427). Finally, figurative language is more memorable than literal language. Tropes are particularly memorable because they lack closure and invite elaboration (p. 429). Increased cognitive effort might be required for the interpretation of an ironic remark or a metaphor, whereas a normal statement requires no interpretation and thus no additional mental exertion.¹⁷ This extra mental energy not only increases memorability, but also leads to greater message processing (Mothersbaugh, Huhmann, & Franke, 2002). Cognitive activity has been recognised as central to humour as well; studies show that more laughter is elicited from more complex cartoons (Lieberman, 1977). This is further evidence that rhetorical figures are an effective operationalisation of playfulness and humour.

Brief remarks shall now be made on two types of rhetorical devices that are of special relevance to political discourse. *Sarcasm* is a form of verbal irony that expresses hostile humour; it expresses mockery or disapproval and is often targeted at a specific person. Papacharissi (2012) observes that irony on Twitter is often provocative in nature, and includes profanity or

¹⁷ Compare, for example, the relative interpretability of the metaphorical ‘I am intoxicated by words: language is my gin, my vodka, my pinot noir’ and the literal ‘I like words’.

incivility. This might convey honesty or authenticity, and occurs because ‘the premise of play’ on Twitter renders it a safe platform for disclosing private thoughts (p. 1999).

Metaphor is the second device that is of importance to the style of political discourse. Charteris-Black (2005) argues that its persuasive power is derived from the hidden associations that it arouses (p. 2). These associations are emotionally evocative because they trigger shared beliefs and values. Through his analysis of the speeches of six successful 20th-century American and British politicians, Charteris-Black (2005) concludes that metaphor contributes to legitimisation in four ways: establishing ethical integrity, heightening the *pathos* of a speech, illuminating policies, and communicating ideology (p. 198). All in all, metaphor can help explain why political rhetoric is successful, and why leaders are charismatic.

Finally, it is important to note that not all technical instances of rhetorical figures are indeed rhetorical figures. Expressions that do not deviate from the norm cannot be considered as figures. These conventionalised expressions are called *clichés*, and quite a few are metaphors (McQuarrie & Mick, 1996).¹⁸ Because they are not incongruous, they are less persuasive.

The linguistics, communication, psychology, and advertising literature leads to the following hypothesis:

HYPOTHESIS 9 If a tweet is playful, it will be more successful.

2.6 Synthesis of Success Factors

This chapter has introduced a number of factors that might influence the success of a political message on social media. These factors relate to discoverability, content, and style attributes, and are drawn from political science, media and communication studies, social and cognitive psychology, marketing, linguistics, and sociology of culture. No previous study has considered all of them at once or integrated theories from across this broad range of disciplines. Moreover, this is the first study to consider the factors in the context of Twitter, framing the tweet as a new media cultural object.

While each factor might influence success in and of itself, it is undeniable that the interplay of different strategies significantly improves the persuasiveness of a message. As Charteris-Black (2005) observes, discursive modes and rhetorical figures ‘*act in combination with one another*’ rather than in isolation’, and this increases message processing (p. 7; emphasis in original). He demonstrates, for instance, that metaphor is even more powerful when integrated

¹⁸ Some examples are ‘time will tell’ and ‘between a rock and a hard place’.

with antithesis. Division (negative campaigning) is more effective when the message contains a substantive issue, as discussed in Section 2.4.

The next chapter presents the multi-method design that was implemented to answer the research questions. It explains how the concepts introduced throughout this chapter were measured in tweets, justifies the EU referendum as a case study site, describes the interview process, and elucidates how the tweets were collected, processed, and analysed.

3. Methodology

‘We no longer have to choose between data size and data depth.’

—Lev Manovich

Given the interdisciplinary nature of the theoretical framework and the complexities of both resonance and rhetoric, a multi-method approach was adopted to answer the research questions. The EU referendum was selected as the case study site due to its timeliness and socio-political significance, as well as the high volume and diversity of persuasive content produced about it on Twitter. Case studies allow for the understanding of dynamics present within a specific setting, as well as the combination of quantitative and qualitative methods (Eisenhardt, 1989). Because only one setting has been chosen, it can be explored in far more depth (Hammersley, 2004). Moreover, case studies can be used to provide description and to test or generate theory—both of these objectives are relevant.

The design of this thesis encompasses content analysis, multivariate regressions, outlier detection and analysis, semi-structured trace interviews, supervised machine learning algorithms for text classification, and natural language processing. The regressions, informed by both manual content coding and the text classification algorithms, represent overall trends in the dataset—i.e., the main factors that help to explain tweet resonance, as well as the relative sizes of their effects. The qualitative analyses of interview data and exceptional tweets contribute to the explanation of these overall trends by providing more context and nuance: under what circumstances these factors might or might not work. Thus, the methods are ‘integrated during the interpretation phase’ (Robson, 2011, p. 165). boyd and Crawford (2012) emphasise the importance of context in big data studies: ‘Data are not generic. There is value to analysing data abstractions, yet retaining context remains critical... Context is hard to interpret at scale and even harder to maintain when data are reduced to fit into a model’ (p. 671). Over 12 million tweets from the EU referendum campaigning period were collected for this thesis, and although they *were* reduced to fit into a model (with the help of machine learning techniques trained on hand-labelled data), many specific examples are analysed with full contextual detail. In some cases, this context transcends that provided by the tweet and its metadata, as the sender of the tweet was interviewed. As the epigraph of this chapter contends, no compromises have to be made between data size and depth. Indeed, the triangulation of different data types ‘provides stronger substantiation of constructs and hypotheses’ (Eisenhardt, 1989, p. 538).

The first section of this chapter justifies the selection of the EU referendum as the case study site. It subsequently provides more information about the socio-political context of the

referendum, and introduces the main referendum campaigns. Section 3.2 explains how Twitter data was collected, stored, processed, sampled, and manually coded for the presence of rhetorical constructs that might relate to tweet resonance. It also explains how the measures of resonance—retweets and likes—were obtained, and describes the limitations of the data. Section 3.3 focuses on the methodology for the automated coding of the full dataset: natural language processing techniques were conjoined with supervised machine learning algorithms for text classification, four of which were compared. The interview strategy and analysis process are elucidated in Section 3.4, and the chapter concludes with reflections on ethical considerations.

3.1 Case Study Site Selection: The EU Referendum

The United Kingdom European Union membership referendum was chosen as the case study site for three main reasons. First of all, there were two clearly defined sides, both of which were represented by numerous hashtags that could be used in the tweet collection script (tweets were gathered via keyword queries). Because it was (and still is) a very controversial and high stakes issue, many of the tweets were rhetorical—the overarching aim was to persuade people to support Leave or Remain, not simply to disseminate information. Thus, the framework established for this thesis can be effectively applied; the concept of message framing from Atkinson (1984), which is premised upon an us vs. them dichotomy, is especially apropos. Secondly, the referendum was timely: the vote was on 23 June 2016, and the campaigning period began on 15 April. This aligned perfectly with the data collection stage for this thesis, which was the entirety of 2016—tweets from not only the campaigning period, but also the month before and three months after were harvested. Thirdly, the referendum was specific to the UK, so the majority of tweets were from a well-defined (British) audience. As noted in the last chapter, the success of a cultural object is influenced by its recipients and their social attributes. Thus, having a more well-defined audience simplifies the analytical process—only one socio-political context, that of the UK, has to be taken into account; this context will be summarised below.

As aforementioned, the primary advantage of having a single case study is analytical depth—the multi-method design encompasses a wide variety of qualitative and quantitative techniques, and large volumes of data were collected from two sources (Twitter and the interviewees). This would not have been scalable to a second case study site within the four-year timeframe of this research. The main shortcoming of a single case study approach is

generalisability. The findings are applicable to other campaign settings, notably referenda, elections, and social movements. They can also be applied to marketing campaigns, but the framework would have to be slightly adjusted to account for the lack of opposing sides. They are less relevant for communication contexts in which persuasion is not the aim or in which popularity is not important (e.g., private self-expression).

The rest of this section will summarise the socio-political context of the referendum and introduce the main campaigns. There is a long history of Euroscepticism in the UK—demands for an EU membership referendum have been voiced since the original referendum in 1975 concerning the UK's membership of the European Economic Community (BBC News, 2015). This referendum was promoted in the Labour Party election manifesto; over the following decades, as a result of controversies surrounding various European treaties, the Conservatives and Liberal Democrats began to call for referenda as well. On 22 January 2013, Prime Minister David Cameron promised a renegotiation of the UK's relationship with the EU, followed by an EU membership referendum before the end of 2017, if the Conservatives won the next election—which they did (BBC News, 2013). Two of his main renegotiation points, which reflected the British public's concerns, were decreasing benefits for migrants and establishing more sovereignty from the EU 'superstate' (BBC News, 2016c). The reform package that was approved did not appease his Eurosceptic critics on these two fronts. On 20 February 2016, a day after the reforms were approved, the referendum date of 23 June 2016 was confirmed (*ibid*).

To better contextualise Cameron's decision to promise a referendum and the campaigns that emerged, it is important to understand the rise of the UK Independence Party (UKIP), which was launched in 1993 from the remains of the Eurosceptic Anti-Federalist League. Under the direction of Nigel Farage, who became their leader in 2006, UKIP leveraged qualms about rising immigration among the British working class. This policy led the Party to achieve significant victories in the 2014 European Parliament election and the 2015 general election. Their victory in the European Parliament election was heralded as a 'political earthquake'—this was the first British national election since 1906 that neither Labour nor the Conservatives won (Wintour & Watt, 2014). With 24 seats and 27% of the popular vote, UKIP performed nearly twice as well as they had in the 2009 European election, when they came second to the Conservative Party with 13 seats. The Conservatives actually lost seven seats in the 2014 election, which indicates that some of their voters shifted allegiances to UKIP. This allowed UKIP to exert pressure upon the Conservative government for an EU membership referendum.

The rift between UKIP and the Conservative Party led to the rise of two campaigns that battled to represent Leave in the run-up to the referendum period: Vote Leave, supported by Boris Johnson and Michael Gove, and Grassroots Out, championed by Nigel Farage. On 13 April 2016, Vote Leave was designated the official campaign, along with the uncontested Remain campaign Britain Stronger in Europe, which was informally associated with Cameron. The Vote Leave victory significantly shifted the nature of the referendum debate, as its objective was to make the ‘ “positive” case for leaving the EU, with less emphasis on immigration’ (BBC News, 2016d). The leaders of the campaign were wary that anti-immigrant rhetoric would alienate undecided voters, many of whom were in the political centre. Stronger In, on the other hand, focussed on the negative impact of Brexit on the economy.

Vote Leave did, however, learn a valuable lesson from UKIP which contributed to their positive case for Brexit. Matthew Elliott, an interviewee and the CEO of the campaign, disclosed that he and his Campaign Director had conducted polling research to understand UKIP’s remarkable performance in the 2014 European election. They discovered that those who had voted for UKIP were attracted to the idea of *sovereignty*—they felt as if they had lost control over their own country. Matthew noted that the word ‘sovereignty’ itself was too highbrow and might be incorrectly associated with the Queen, but ‘taking control encapsulated what the public were hoping for from Brexit’. Thus, Vote Leave’s main tagline became ‘take control’.¹⁹ The second focus of the Leave campaign’s messaging strategy was another salient theme for the British public: health (more specifically, the NHS). One of their famous pledges became a second tagline, as it was displayed on the side of their Brexit bus—‘Let’s give our NHS the £350 million the EU takes every week’.

Grassroots Out capitalised on the relationship not only between immigration and healthcare, but also between immigration and national security. This is best evidenced by Farage’s ‘breaking point’ poster, which was released on 16 June. The poster depicts a long queue of migrants and refugees, and was unveiled with a subheading that plays with Vote Leave’s main tagline: ‘We must break free of the EU and take back control of our borders’ (Stewart & Mason, 2016). In the aftermath of terrorist attacks in Paris and Brussels and the revelation of the Dusseldorf bomb plot, security was a prominent concern—and it was the angle that Farage emphasised in his defence of the poster, as he claimed that jihadi terrorists might be masquerading as refugees (ibid). Remain campaigners, on the other hand, portrayed security from a far more positive perspective: preventing another world war. This defines the

¹⁹ It was originally ‘take back control’, but the word ‘back’ was removed for succinctness.

very premise of the EU, which was founded to establish and preserve peace in the wake of two world wars (European Union, 2017).

All in all, the contextualisation of the case study site has provided further justification for the five categories of the issue variable introduced in the previous chapter: the economy, immigration, sovereignty, security, and healthcare. Their measurement will be explained in the following section.

3.2 Twitter Data Collection, Processing, and Analysis

Tweets concerning the EU referendum constitute the primary data source for this thesis. Over 12 million in the full dataset pertain to the official campaigning period: 15 April to 23 June 2016. This section reviews how they were collected, stored, cleaned, and analysed; it also highlights the limitations of Twitter data. All code for collecting and processing the data was written in Python because it is ‘a simple, yet powerful, high-level, general-purpose programming language that is well-suited to developing scripts for parsing through large bodies of mostly-uniform data’ (Lewis et al., 2013, p. 43). Many Python libraries are extremely useful for the manipulation of big data: pandas²⁰ and NumPy²¹ were especially invaluable for this research. The scripts that were developed have been uploaded to a GitHub repository for reproducibility purposes.²²

3.2.1 *Tweet Collection*

A keyword—and more specifically hashtag—query was used to obtain referendum-relevant tweets from Twitter’s Streaming Application Programming Interface (API). This API has been developed specifically for individuals with data-intensive needs or who are interested in analytics research. With a modicum of basic programming knowledge, along with a long-lived HTTP connection, the API is quite easy to use. Independent collection has the added advantages of control and timeliness: filters can be specified to tailor the representativeness of the dataset, and all tweets are scraped in real time.

In March 2016, an initial list of 92 hashtags containing referendum-related keywords was established through the examination of Twitter accounts related to Leave and Remain campaigns. This was supplemented by suggestions from RiteTag,²³ a social media tool that

²⁰ <https://pandas.pydata.org>

²¹ <https://www.numpy.org>

²² <https://github.com/Periwyinkle/brexit-tweet-analysis>

²³ <https://ritetag.com>

illuminates the best hashtags to use for a given topic. Registered campaigns were identified from the Electoral Commission webpage,²⁴ and grassroots campaigns were identified from Twitter suggestions and the webpages of partner campaigns. To confirm that the hashtags were contained in tweets that were predominantly concerned with the EU referendum (terms like #EU and #UK were omitted because they were too general), the latest 30 tweets including each hashtag were inspected. The inspection process encompassed not only their text and any multimedia or hyperlinks that they contained, but also the profiles of their users, if additional context was helpful. If all 30 tweets were relevant to the referendum, the hashtag was kept in the list. This process resulted in additional referendum-related hashtags being discovered.

A Python script querying Twitter for tweets including the 92 hashtags was run for approximately 20 hours on 8 March. All 2,826 hashtags from the 25,198 tweets were extracted from the resultant JavaScript Object Notation (JSON) file—this constituted the 92 from the original list along with their collocates. Each collocate hashtag was independently reviewed, and those relevant to the referendum that were not too general were inspected on Twitter using the same process described above. This resulted in 95 new hashtags, which were added to the script and created a total of 187 query keywords. The updated script began running on 10 March. Throughout and after the official campaigning period (15 April-23 June), hashtags were added to the query to capture the evolving discussion (e.g., #CatsAgainstBrexit emerged on 18 June and reached trending status) and new campaign groups (the Electoral Commission website was periodically reviewed). In total, 18 updates were made between 1 May and 30 June. The final update made before the vote was on 22 June, and expanded the hashtag list to 252 (Appendix A).

3.2.2 *Twitter Data Limitations*

As boyd & Crawford (2012) declare, the streaming API does not provide access to all tweets that are published (the firehose). Rather, a limited spritzer of 1% of public tweets represents the upper limit that can be collected in real time. As 500 million tweets are published every day, the daily limit is approximately five million, which is significantly higher than the volume of tweets about Brexit. On 23 June itself, 2.3 million tweets were collected from the script, and this was the highest volume during the referendum period. Thus, it is reasonable to assume that the data gathered are quite comprehensive. Completeness and representativeness cannot be guaranteed, however; multiple studies have shown that the spritzer is biased (Driscoll & Walker, 2014; Morstatter, Pfeffer, & Liu, 2014). Lewis, Zamith, and Hermida (2013) provide

²⁴ <http://search.electoralcommission.org.uk/Search/Registrations>

direct evidence of the significant discrepancies between tweets scraped from the API and the full dataset of tweets published by a specific user; the latter was obtained directly from the user himself. They observed that the sample from the API ‘was missing “chunks” of tweets during some time periods, and it included less metadata or other information about the tweets’ (p. 42). Notably, the sample did not result from random error, but rather systematically missing chunks. This suggests that the Twitter API is not a truly representative sample.

However, after systematically comparing spritzer results from the Streaming API with results from the full firehose (which is provided by Gnip PowerTrack²⁵), Driscoll and Walker (2014) conclude that ‘the Streaming API excels at longitudinal data collection, but is a poor choice for massive, short-term events’ given its 1% limit (p. 1762). In fact, for long-term events with lower tweet volumes per day, there is almost no difference between spritzer and firehose results—having access to the firehose ‘does not necessarily enable more meaningful research design’ (ibid). As Brexit is definitely such an event, this finding lends further confidence to the comprehensiveness of the dataset used in this thesis. Moreover, capturing the tweets in real time (which required knowing about the EU referendum in advance of the campaigning period) was essential for comprehensiveness. The standard Search API, which allows for the capture of historical tweets, has significant rate limits: up to 18,000 tweets over the last seven-day period can be retrieved, whichever limit is reached first, and up to 180 calls are permitted every 15 minutes. Twitter documentation makes clear that ‘the standard search API is focused on relevance *and not completeness*’—tweets and users might be missing from the results (Twitter Developers, 2018b). Complete access is possible with the premium or enterprise versions of the API, but both are extremely costly.

On a higher level, it must be noted that the keyword query approach also introduces constraints on comprehensiveness—but with the benefit of increased precision. Bartlett, Miller, Reffin, Weir, and Wibberley (2014) elucidate the tradeoff between the two: ‘A very precise data collection strategy generally only returns tweets that are on-topic, but will likely miss some. A more comprehensive data collection strategy collects more of the tweets that are on-topic, but will likely include some which are off-topic’ (p. 86). Even though the query used for this thesis was periodically updated with new keywords, the dataset collected by no means encompasses all tweets about the EU referendum. Its most significant limitation is that only tweets with hashtags were retrieved; tweets that contain non-hashtagged relevant keywords were not. Precision was valued over completeness—as noted above, terms like #EU and #UK were not included because they were too general. The dataset is, however, balanced with

²⁵ <http://support.gnip.com/apis/powertrack>

respect to Leave and Remain: 109 hashtags in the list are Leave-oriented, whereas 122 are Remain-oriented (the rest are either ambiguous or neutral).

Even if the dataset was completely comprehensive, it still would not represent the British population. Blank (2017) has shown that Twitter users are more elite than other internet users, who are in turn more elite than non-internet users. More shall be said on this in the Discussion chapter's reflections upon the real-life impact of referendum tweets. It is also important to be aware that accounts and users are not equivalent. As boyd and Crawford (2012) explain, 'Some users have multiple accounts, while some accounts are used by multiple people. Some people never establish an account, and simply access Twitter via the web' (p. 669).

Some accounts do not represent people at all—they are partially or fully automated, and are commonly referred to as *bots*. These accounts constitute a category of exceptional users who have become part of the fabric of Twitter discourse, particularly during political events and controversial topics. Given the public prominence of Brexit and how closely contested the vote was, it is indubitable that a certain proportion of tweets in the dataset were created by bots. As bots are controlled by humans to some extent, it is impossible to identify all of them definitively, but there are some basic behaviours that distinguish automated accounts. The most notable behaviour—what makes them exceptional—is frequency of tweets, many of which are retweets (Ferrara, Varol, Davis, Menczer, & Flammini, 2016). Because bots play such an active role on Twitter in political contexts, they were not removed from the dataset (removing them would decrease its representativeness). Indeed, Chapter 6 provides evidence of tweets that attained their success as a result of bot networks.

3.2.3 *Tweet Storage and Processing*

The entire dataset of referendum-related tweets spans a 6.5-month timeframe, from 10 March to 30 September 2016. The focus of this thesis is on the 9,792,191 tweets published between 15 April and 22 June, the official campaigning period. More specifically, only the 2,598,609 *original* tweets published during this period are examined, as retweets constitute one of the dependent variables. Thus, they were removed from the dataset. Further preprocessing steps are described in this subsection and summarised in Table 3.1.

Given the size of the dataset and its fixed schema, a relational database storage solution was explored to circumvent loading the data into memory. SQLite was the most ideal option because it is supported by Python; moreover, it is designed for device-local storage that is less than 1 TB. However, a Pythonic solution that proved to be more robust than relational

databases was found. The pandas library has an excellent tabular data structure that allows for the efficient manipulation of large volumes of data: the dataframe. Dataframes can be saved as HDF5 files,²⁶ and loaded within a matter of seconds. This solution provided more direct access to the dataset than SQLite. Moreover, as there is only one ‘table’ (i.e., dataframe) containing full information about every tweet, a relational database is not strictly necessary.

A Python script was written to extract relevant data and metadata from the tweet JSON files into pandas dataframes (one JSON file was produced for each day), which were concatenated and saved in HDF5 format. Comprised of all 2,598,609 original tweets, the resultant dataframe was examined and preprocessed in a number of ways. Faulty values (32 tweets with negative user follower, friend, status, or favourite counts) were removed, as were 1,076 tweets from 30 March and 21,326 tweets from 23 June. The former had been included because of a Twitter API glitch; the latter because of the five-hour timezone difference between the location of the server running the script (New York City) and the UK. Temporal control variables derived from tweet metadata were added: date, time, day of week, hour, hour binned, and number of days before the referendum. Two independent variables related to discoverability, URL count and visual media count, were dichotomised (for reasons described in the next chapter). Finally, a separate binary variable was created for each media type—photo, video, and GIF.

A second round of content coding was conducted upon 100 randomly sampled tweets from this processed dataframe. This analysis revealed that those containing only neutral hashtags such as #Brexit or #EURef were less likely to be persuasive, and thus could not be fully labelled. All of the hashtags included in the last-updated Streaming API script before the referendum were reviewed, and partisan hashtags were identified through the manual inspection of 20 random tweets containing each hashtag.²⁷ 184 of the 252 hashtags in the query were considered to be partisan; the six most frequent were #VoteLeave, #Leave, #LeaveEU, #StrongerIn, #Remain, and #VoteRemain. The 1,591,877 tweets that did not contain partisan hashtags were subsequently removed. The content coding process also revealed that tweets that are replies to other tweets are more difficult to interpret, and so the 819,891 replies in the dataset were removed as well.

The final preprocessing stage, the removal of spam and unrelated messages, was inspired by the content coding of the last random sample that was drawn of 3,106 tweets. One

²⁶ HDF5 is a binary data format developed by the HDF group: <https://www.hdfgroup.org>

²⁷ This process revealed some interesting trends: hashtags that are variations of #ProjectFear or that are plural proper noun forms (e.g., #Remainers, #Remaniacs) are almost always part of tweets that exhibit negative campaigning (and, more specifically, name calling). They tend to make broad generalisations about the other side.

hundred and thirty tweets in this sample were identified as falling into either category (see Appendix B for a selection of examples). These 130 tweets were deleted, along with 7,013 other tweets from the 58 users who had sent them (e.g., @TheEdgeKville, an alternative rock radio station that used #TakeControl, and @AlphaPhi_MB, a sorority that used #bettertogether). A subset of these tweets were in the format of a news digest. Their metadata revealed that they were all from the same source: Paper.li,²⁸ an application that aggregates ‘links from the people you follow on ‘Twitter’ and organises them into a virtual newspaper, which can be automatically tweeted (Brooks, 2010). All 3,942 tweets from this source were removed. The resultant final dataset contains 741,411 tweets, and effectively represents original, non-reply, partisan tweets published between 15 April and 22 June. This population dataset, along with the hand-coded random sample of 2,976 tweets (without spam and unrelated content), was used in the regressions presented in the next chapter.

TABLE 3.1: Number of Tweets at Each Preprocessing Stage

PREPROCESSING STAGE	NUMBER OF TWEETS
Referendum period: 15 April-23 June 2016	12,065,669
Removed tweets from 23 June 2016	9,792,191
Removed retweets	2,598,609
Removed faulty values, tweets from 30 March (technical glitch), and tweets from 23 June (timezone difference)	2,576,175
Removed tweets without partisan hashtags and replies	752,496
Removed spam and unrelated tweets to create population dataset	741,411
Drew random sample from population dataset	2,976

3.2.4 *Measurement of Tweet Resonance*

As explained in the previous chapter, success in this thesis is defined as popularity, which is represented on Twitter by two publicly visible forms of engagement: retweets and likes. Both can be directly collected from the Twitter API. A third publicly visible form of engagement, replies, is also part of tweet metadata from the API, but is more difficult to capture. The metadata only indicate whether a tweet is a reply and the ID of the parent tweet if so; to calculate accurate reply counts, a fully comprehensive dataset is required. Moreover, replies measure the social interactions that a tweet receives. This is not the same concept as popularity, although it is related.

It should also be noted that there are many other forms of engagement that are only visible to account holders about their own tweets: impressions, link clicks, hashtag clicks, user

²⁸ <https://paper.li>

profile clicks, follows, etc. (Twitter Help Centre, 2018a). *Impressions* is the most relevant, as it is a direct measure of popularity: how many people have seen the tweet? As such, it is a useful complement to retweet and like count; many digital marketers combine the measures by calculating *engagement rate*—i.e., the total number of retweets and likes divided by the total volume of impressions (Chacon, 2018). Impressions cannot be collected from the Twitter API, but several of the interviewees shared expanded metadata for a selection of their tweets. Thus, references will be made where relevant in the qualitative analysis.

Retweet count and like count were measured in two different ways. First, Twitter's REST API was queried for the IDs of all 2,976 tweets in the coded random sample on 10 July 2017. Given the API's rate limits, this was impossible to replicate for the population dataset. Although it provided the most up-to-date counts for both variables, as the script was run over a year after the referendum period ended, data could not be obtained for 18.45% of the tweets in the sample. These 549 tweets were from users who had removed their accounts, made their accounts private, or been suspended by Twitter between the publication date of their tweet(s) and 10 July 2017.

To circumvent the problems posed by missing data, and to scale up the measurement to the population level, retweet and like counts were manually computed from the tweet JSON files collected from Twitter's Streaming API. As retweets contain the same ID as the original tweet, the 12,671,189 retweets from 15 April to 1 July 2016 (seven days after the last day in the referendum period) were extracted into a pandas dataframe. The dataframe was grouped by the 1,653,121 unique IDs, and maximum retweet and like counts for each ID were calculated. This dataframe was subsequently merged with the dataframe of 741,411 tweets in the full dataset on tweet ID.

The like counts provided by this method are both incomplete and more severe underestimates than those provided by the API. This is because a tweet needs to have been retweeted in order for the like count to be captured; thus, the most up-to-date like count is the count at the time of the last retweet. Subsequent likes are not captured, and tweets that only received likes but no retweets are recorded as having missing data for this variable (this affected over 60% of tweets in both the random sample and the population). The implications of this are discussed in the next chapter.

3.2.5 *Tweet Sampling and Content Analysis*

As the dataset is far too large to be manually labelled in full, a representative random sample was drawn for the content analysis. To avoid having to load the entire dataset into memory, a

reservoir technique was used (Vitter, 1985). The reservoir is small enough to fit into memory, as it is the size of the random sample (N). After it is filled with N tweets, its elements are randomly replaced with decreasing probability: tweets that first entered the reservoir have more opportunities to be replaced, whereas tweets that entered at later stages are less likely to be replaced. When the algorithm has fully executed, every tweet in the dataset has had an equal probability of being chosen for the reservoir. An alternative approach is generating N random numbers, the maximum being defined by the size of the dataset, and treating these numbers as tweet indices. Reservoir sampling is more convenient, however, because the size of the dataset does not have to be determined in advance for an evenly distributed sample to be obtained.

Following the general procedure described by Lewis et al. (2013), coding categories were defined based upon the hypotheses presented in the previous chapter, and a coding scheme was established with ‘explicitly formulated rules’ and examples (p. 36). The full coding scheme is in Appendix C. As most discoverability features are present in tweet metadata (e.g., hyperlinks, media items, @mentions, and hashtags), all but two of the categories relate to content and style: position, organisation, message framing, message format, question, emotion, resolution, playfulness, and issue. A British doctoral student with a background in political science was selected to be the second coder; this background was important due to the numerous cultural and political references made in the tweets (due to their brevity, contextual information for the references is extremely limited). The coding scheme was iteratively revised throughout three rounds of coding, which were required to establish sufficient levels of interrater reliability (IRR) for all of the variables.

IRR scores are important for validity and objectiveness, as they represent a ‘standard measure of research quality’ (Cho, 2008, p. 345). The best scores for nominal data coded by two people were computed: percent agreement and Cohen’s kappa. Percent agreement ranges from 0 to 1 (complete disagreement to complete agreement), and is the most widely used index due to its simplicity. It is defined as ‘the proportion of coding decisions that reached agreement out of all coding decisions made by a pair of coders’ (ibid). Because this measure does not take chance agreement into account (if there are only two categories, for instance, agreement by chance occurs 50% of the time), Cohen’s kappa (κ) was also calculated.²⁹ Cohen’s kappa is regarded as a highly conservative measure of reliability, and never reaches 1 (Clark-Carter,

²⁹ Krippendorff’s alpha (α) also takes chance agreement into account and is generally a more flexible measure than Cohen’s kappa, as it can handle all variable types (Cho, 2008). However, because it is more mathematically demanding and has a tendency to be quite low when applied to ‘binary classification tasks with unbalanced class distributions’, it is not as appropriate as Cohen’s kappa for this thesis (Burscher, Odijk, Vliegthart, de Rijke, & de Vreese, 2014, p. 197). Many of the categories are binary and unbalanced.

1997; Lombard, Snyder-Duch, & Bracken, 2002). It combines percent agreement and expected agreement, which is 'based on the assumption that the distribution of values in the population mirrors the multiplicative marginals of the crosstabulation' (Brader, 2006, p. 246). If percent agreement is 70% and expected agreement is 40%, Cohen's kappa would be 0.50 $(0.70 - 0.40 / 1 - 0.40)$. There are no strict thresholds for either IRR measure, but a kappa of at least 0.41 and percent agreement of at least 65 was considered to be sufficient for this study. Both scores for each category are presented in Table 3.2.

One hundred randomly sampled tweets were coded in the first round for all categories, 50 in the second round for all categories, and 50 in the third round for only the categories that had *not* met minimum thresholds. After each round, the tweets with conflicting codes were reviewed in person; each category was discussed separately. Differences were reconciled through the explanation of how each category had been approached,³⁰ and the coding scheme was updated to reflect agreed upon changes. Before the third round, ten tweets were collectively coded and all category decisions were unanimous. This process reflects the fact that although qualitative judgement is an integral aspect of content analysis, it is 'a research technique for the objective, systematic, and quantitative description of the manifest content of communication' (Berelson, 1952, p. 18).

Moreover, the coding process transcended common practise in that the tweets were viewed in their original context on Twitter, not in a spreadsheet limited to their text. URLs, media items, and user profiles were rigorously scrutinised in order for holistic assessments to be made about tweet content. Hyperlinked articles were read and videos were watched; user descriptions, pictures, and previously published tweets were reviewed. To facilitate this process, a column in the dataframe linking directly to the tweet was added. About one in five tweets was no longer available on the platform at the time of analysis; only their raw text and metadata could be scrutinised. This was almost always sufficient for content coding, however.

³⁰ A 'Notes' column was added to the spreadsheet containing the random samples; justifications of non-straightforward coding decisions were included here.

TABLE 3.2: Interrater Reliability Scores

	CONTROL	DISCOVERABILITY		STYLE	
IRR METRIC	Position	Organisation	Message Format	Question	Playfulness
Percent Agreement	89.8	98.0	83.7	98.0	89.8
Cohen's Kappa (unweighted)	0.809	0.79	0.479	0.921	0.609
Cohen's Kappa z	6.39	5.65	3.93	6.47	4.39
Cohen's Kappa p value	1.70E-10	1.56E-08	8.61E-05	9.83E-11	1.15E-05

	CONTENT					
IRR METRIC	Message Framing	Enthusiasm	Anxiety	Aversion	Issue	Resolution
Percent Agreement	65.3	87.8	81.2	79.6	72.9	83.7
Cohen's Kappa (unweighted)	0.477	0.743	0.419	0.523	0.557	0.532
Cohen's Kappa z	4.84	5.22	3.07	3.73	6.20	3.75
Cohen's Kappa p value	1.30E-06	1.82E-07	0.00214	0.00019	5.77E-10	0.000176

The categories of position are remain, leave, ambiguous, and N/A; the categories of message framing are identification, division, and other; the categories of issue are economy and labour, healthcare, immigration, security, sovereignty, and other. All other variables are dichotomous.

As mentioned above, the coding categories were position, organisation, message framing, message format, question, emotion, resolution, playfulness, and issue. Position, organisation, and question were the most straightforward to code, and this is numerically reflected by their high percent agreement and Cohen's kappa scores. Question is the least subjective category because its presence is indicated by the question mark symbol; all disagreements resulted from careless error. Scrutiny of user profiles was important for both position and organisation; many users explicitly indicate the side they support in their description. Even more impassioned accounts include position-relevant keywords in their usernames, or have profile or header photos that reflect Leave or Remain (e.g., via a UK or EU flag). Because users define their own identities on Twitter, determining whether accounts represent organisations was sometimes slightly complicated, and required external validation via Google searches. In their analysis of Arab Spring tweets, Lewis et al. (2013) found it difficult to classify a large number of users into actor types based upon Twitter profiles.

Message framing and message format were both quite complex. Identification and division (the main categories of message framing) are defined with respect to position, and also encompass a broader range of targets than those that would apply to a normal political context. Any endorsement or criticism of a person, organisation, policy, belief, or value relating to the user's side or the opposite side was labelled as identification or division, respectively. Given this broad definition, examples could not be provided for all subcategories; they had to be recognised through the coding process itself (one prominent subcategory of identification rhetoric was photographs of people participating in campaign rallies). The original version of message format was too complicated, as it was split into four options: contrast, list, list-contrast (i.e., the presence of both), and 'Other'. After the first round of coding, the second coder commented on the lack of an option for 'None'—he used 'Other' to classify tweets that did not contain lists or contrasts. The dichotomisation of this variable into the presence of either message format or neither allowed the IRR scores to be raised above the threshold: in the final round a kappa of 0.479 and percent agreement of 83.7 were achieved.

Emotion was originally comprised of four separate binary variables: enthusiasm, pride, anxiety, and aversion. Pride is a subcategory of enthusiasm, and is defined as appeals to group or national identity. Although it is highly relevant to the EU referendum (many proponents of Leave were extremely nationalistic), it is also highly subjective, and sufficient IRR scores could not be achieved. Thus, it was removed from the coding scheme. This improved the IRR scores for enthusiasm to a Cohen's kappa of 0.743 and percent agreement of 87.8. Recognising anxiety and aversion was less straightforward, as reflected by their low kappa scores even after

the third round of coding (0.419 and 0.523, respectively). Percent agreement was quite high, however (81.2 and 79.6). Interestingly, anxiety was more difficult to systematically detect than aversion.

Resolution was initially divided into three levels: 1, 2, and 3. As 0 was not an option, this conflated no resolution with weak resolution. The category was subsequently recoded as 0, 1, and 2. As the distinction between moderate and high resolution could not be agreed upon, they were collapsed into the same code (the variable was dichotomised). This raised the IRR scores above the threshold: Cohen's kappa was 0.532 and percent agreement was 83.7 after the third round.

Playfulness was another complex variable to code, given its wide-ranging definition: it encompasses any non-literal use of language. Some of the most common rhetorical figures in political language are irony, sarcasm, metaphor, hyperbole, pun, personification, and paradox. Scrutiny of media items was critical for the assessment of playfulness, as it often results from the juxtaposition of text and media. Sufficient IRR scores could not be established until the third round of coding, but ultimately a Cohen's kappa of 0.609 and percent agreement of 89.8 were achieved.

Even though multiple issues can be expressed by a tweet, only the most salient was labelled to reduce the complexity of the coding process. This generally worked, but some tweets were on the borderline between two subcategories, and holistic judgments had to be made that were necessarily somewhat subjective. Immigration featured in many borderline tweets, as it relates to security (in tweets about illegal migrants), the economy (in tweets about migrants taking British jobs), and healthcare (in tweets about migrants using the NHS). After the third round of coding, a Cohen's kappa of 0.557 and 72.9% agreement were achieved.

3.3 Automated Coding of Full Dataset

Grimmer and Stewart (2013) provide an excellent overview of all automated text classification methods, first making a distinction between classification and ideological scaling, and then further distinctions among classification techniques, first between known categories and unknown categories. The former is relevant for this thesis, as content categories of interest have already been identified, and 2,976 tweets have been manually labelled for them. The major division within classifiers for known categories is between dictionary-based approaches and supervised machine learning (SML) approaches.

Given the manually labelled dataset, SML was the better option for most of the variables—the larger the training set, the more powerful it is; dictionary-based methods are often not transferable across datasets or domains, as they are rule-based (and the domain of tweet text is particularly tricky). They were adopted for two variables, however: question and organisation. SML has been extensively applied in recent years in communication and journalism studies, which have had a computational turn. The SML approach adopted for this thesis was inspired by Burscher, Odijk, Vliegenthart, de Rijke, and de Vreese (2014) and van Zoonen and van der Meer (2016). It was also informed by a Demos³¹ report on Islamophobia, xenophobia, racism, and hateful views on Twitter (Miller et al., 2016). Three of the content features they classified with SML in their analysis of Brexit tweets overlap with those measured in this thesis: position (Leave, Remain, Ambiguous), issue (the tweet’s topic), and attacking vs. defending (roughly equivalent to message framing, or identification vs. division).

As Python’s scikit-learn library³² for data science is extremely comprehensive and intuitive, its off-the-shelf algorithms were adapted to classify tweets for the same variables that had been manually labelled (Pedregosa et al., 2011). One customised script was written for each variable, incorporating the procedures described in the rest of this section.

3.3.1 *Text Preprocessing and Vectorisation*

Before any machine learning algorithms could be implemented, tweet text³³ had to be preprocessed and vectorised: URLs, punctuation, and special characters were removed, the text was tokenised by splitting at whitespaces, and all tokens were rendered lowercase. Preprocessing was tailored to the feature; for emotion and playfulness, certain punctuation marks were not ignored, as they are expressive. The tokenised tweets were then converted into vectors that numerically represent the presence and frequency of words (tokens) in each tweet. The best vector to use is *term frequency-inverse document frequency* (tf-idf), as it balances the frequency of words in specific tweets against their overall frequency in the dataset—those that occur too frequently in the dataset are downscaled in weight. To be more precise, *term frequency* is calculated by the number of times a word appears in a tweet, divided by the total number of words in the tweet. *Inverse document frequency* is calculated by the logarithm of the number of

³¹ Demos is a British cross-party think tank that has partnered with the University of Sussex to analyse large-scale social media data: <https://demos.co.uk>

³² <https://scikit-learn.org/stable>

³³ For the position feature, usernames, user screennames, and user descriptions were all added, as these often contain pro-Leave or pro-Remain keywords.

tweets in the dataset divided by the number of tweets that include the word.³⁴ The mathematical formula can be represented as follows:

$$tf-idf(t, d, D) = tf(t, d) \times idf(t, D)^{35}$$

Thus, the tf-idf score highlights the most interesting and unique words: the words that should in principle define the categories of the content features. The ten tokens with the highest tf-idf scores are *campaigning*, *voteleave*, *leadnotleave*, *day*, *remain*, *bremain*, *eddieizzard*, *strongerin*, *voteremain*, and *hurt*. It is interesting to note that six of them relate to position.

The default settings provided by scikit-learn’s CountVectorizer and TfidfVectorizer were used for this process, and experimentations were made with a stemmer derived from Python’s Natural Language Toolkit (NLTK)³⁶ library to reduce the number of features. Stemmers collapse words that contain the same root form into the root form by removing affixes (e.g., *play*, *player*, and *played* all become *play*). All three options were tested: Porter, Snowball, and Lancaster (the most aggressive). A more sophisticated and computationally intensive form of stemming is lemmatisation, which aims to be lexicographically accurate because it classifies tokens according to their part of speech before removing affixes (unlike a stemmer, a lemmatiser would not remove ‘s’ from ‘thus’ because it recognises that ‘thus’ is not a plural noun). As lemmatisation is not easy to incorporate into scikit-learn’s default vectorizers, a custom-built preprocessing pipeline using NLTK was designed and tested on all of the content features as a contrast to scikit-learn. An option to use NLTK’s TweetTokenizer was included in this preprocessor; it is more appropriate for the emotion variables because it recognises emoji as tokens. All in all, however, the significant extra time that the implementation of stemmers and lemmatisers required was not at all justified by the marginal increase in accuracy that they provided over the scikit-learn settings—in some cases they even *decreased* accuracy slightly. This is not a surprise, as tweet text is highly context-specific due to its brevity, and full of abbreviations and slang; part of speech tagging is largely inaccurate.

3.3.2 Algorithm Training and Testing

Once the tweet text had been preprocessed and transformed into tf-idf vectors, machine learning algorithms could be trained and tested on the hand-labelled random sample of 2,976 tweets. Training was performed on 90% of the tweets, and testing on 10% (the decisions of the algorithms were compared with the manual coding decisions to determine accuracy). Three

³⁴ <http://www.tfidf.com>


³⁵ t denotes term, or word, d denotes document, or tweet, and D denotes the entire corpus of documents (Liu, 2015).

³⁶ <https://www.nltk.org>

basic models that are extensively used in the literature on text classification were deployed: Naïve Bayes, Linear SVC (a type of Support Vector Machine)³⁷, and Logistic Regression. Two advanced ensemble classifiers, Random Forest and Stochastic Gradient Boosting (Gradient Boosting Machines), were also tested. As they are less accurate than the best-performing basic classifier, they were only implemented for the issue feature. All of the classifiers were incorporated into a Voting Ensemble (VE) classifier, which balances out the weaknesses of each individual model and uses a majority vote in its decisions.

To better understand how the models were making decisions, the 20 tokens with the strongest coefficients were visualised for the models that produced them. Coefficients measure how correlated the tokens are to each category—*greenerin*, *nhs*, and *immigration*, for example, are the tokens with the highest coefficients for the presence of an issue for the linear SVC classifier, whereas *hope*, *remain eu*, and *catsagainstbrexit* correlate most strongly with no issue (Figure 3.1). It is interesting to note that the unigram *eu* signifies the presence of an issue, whereas the bigram *remain eu* signifies the absence of an issue. This is because tweets that contain *eu* are more likely to contain other keywords related to a substantive theme (such as *economy* or *immigrants*), whereas tweets that contain *remain eu* are more likely to be about general pro-EU campaigning.

FIGURE 3.1: 20 Features with Strongest Coefficients (SVC Issue Classifier)



A better understanding of why the classifiers were making mistakes was established through the manual review of misclassifications. The algorithms for issue, predictably, did not work well on tweets that use metaphorical or allusive language to express themes (e.g., ‘A VOTE TO STAY IS A VOTE FOR INFINITY KEBAB’ and ‘They spread like locusts and

³⁷ This model was compared with a related model, Linear SVM with Stochastic Gradient Descent (SGD) training. Accuracies were very similar, but Linear SVC had slightly better performance for most of the variables.

breed like rats’ were not recognised as relating to immigration). Conversely, the algorithms misclassified tweets that have no theme as having a theme if the language is abstract—i.e., references to reasons to vote Leave or Remain that are made on a high level (e.g., ‘Believe in Britain #EU #brexit #referendum #voteleave’, ‘the stakes are high for every one of us’).


3.3.3 *Model Refinement and Selection*


Scikit-learn’s grid search was used to optimise the parameters of the models. As is best practise, to avoid overfitting 10-fold cross-validation was performed both during and after grid search. Grid search only uses a subset of the training data, so the models with the best parameters that it produced were trained and tested on the entire random sample; 10-fold cross-validation means that the training and testing were performed ten times, on different splits of the data.

Classification reports showing the average accuracies from all ten folds for each classifier were computed. There are three different measures of accuracy: *precision* (the proportion of correct predictions made for a category out of all predictions made for that category), *recall* (the ability of the algorithm to classify all instances of a category), and *F1 score*, which is a balance of precision and recall. These measures were calculated for each individual category of each variable, and overall averages were also calculated. A perfect accuracy is represented by a score of 1.00. Given the fact that most of the variables have imbalanced class frequencies, they are weighted by the number of true instances of each class. Although the overall average is important, what is just as—if not even more—important is the accuracy of the classifier for detecting the presence of the categories of interest (a code of 1), as opposed to the absence of the categories (a code of 0). If the overall average is high but the average for the category of interest is low, it is impractical to implement the classifier.

Barplots of the mean F1 score after 10-fold cross-validation of all of the classifiers were produced for each variable, and are presented in Figure 3.2. This visualisation was used to assess which model (if any) to train on the entire random sample in order to predict the variable in the population from which the random sample was drawn (741,411 tweets). The cutoff for implementation was an F1 score of 0.77. In summary, the classifiers for position, the binary version of emotion, and resolution had the highest accuracies. The classifiers for playfulness and message format were the weakest. Classification reports for the best classifiers for each variable are presented in Table 3.3.

FIGURE 3.2: Boxplots of Classifier F1 Scores for Coded Variables





A binary version of the issue variable was created, as accuracies for the multiclass version were not high enough (the algorithms were especially confused by the ‘None’ and ‘Other’ categories). The Voting Ensemble is the best classifier, as it has the highest median and mean F1 score: 0.77 (+/- 0.02)³⁸. Its mean F1 score for predicting issue is 0.77 as well. This classifier was used to predict the presence of issue in the population.

For position, linear SVC and VE are similar in performance: both have a 0.82 (+/- 0.02) mean F1 score. Linear SVC was chosen because the F1 scores from 10-fold cross-validation are more closely distributed around the median, even though the maximum and minimum scores are higher and lower than those for VE. Moreover, Linear SVC and VE are equally good at predicting Leave and Remain, but Linear SVC is slightly better at predicting Ambiguous (its F1 score is 0.52 vs. 0.50 for the VE, averaged after 10-fold cross-validation). This classifier was used to predict position in the population.

Due to the low accuracies of each individual emotion classifier, a binary version of the variable was created (0 for contains no emotion, 1 for contains enthusiasm, anxiety, and/or aversion). Linear SVC and VE are similar in performance: both have a mean F1 score of 0.72, although VE has less standard deviation (0.01 vs. 0.03). VE is also slightly better at predicting the presence of emotion; the mean F1 score is 0.85 (vs. 0.84). This classifier was used to predict the presence of emotion in tweets in the population dataset.

³⁸ 95% confidence intervals: 95% of the scores fall within this range (+/- 1 standard deviation from the mean).

TABLE 3.3: Accuracy Measures for Best Classifiers

VARIABLE & CLASSIFIER	CATEGORY	PRECISION	RECALL	F1 SCORE	SUPPORT
Issue (Binary) <i>Voting Ensemble</i>	0	0.76	0.76	0.76	1,447
	1	0.77	0.77	0.77	1,529
	Average/Total	0.77	0.77	0.77	2,976
Position <i>Linear SVC</i>	Ambiguous	0.76	0.39	0.52	301
	Leave	0.85	0.90	0.87	1,509
	Remain	0.81	0.86	0.83	1,166
	Average/Total	0.83	0.83	0.82	2,976
Enthusiasm <i>Voting Ensemble</i>	0	0.81	0.89	0.85	2,000
	1	0.72	0.57	0.64	976
	Average/Total	0.78	0.79	0.78	2,976
Anxiety <i>Voting Ensemble</i>	0	0.87	0.92	0.90	2,433
	1	0.52	0.40	0.45	543
	Average/Total	0.81	0.82	0.81	2,976
Aversion <i>Linear SVC</i>	0	0.79	0.78	0.78	1,883
	1	0.62	0.63	0.63	1,093
	Average/Total	0.73	0.72	0.72	2,976
Emotion (Binary) <i>Voting Ensemble</i>	0	0.46	0.25	0.33	713
	1	0.79	0.91	0.85	2,263
	Average/Total	0.71	0.75	0.72	2,976
Message Type <i>Linear SVC</i>	Division	0.74	0.90	0.81	1,619
	Identification	0.65	0.52	0.57	599
	Other	0.77	0.53	0.62	758
	Average/Total	0.73	0.73	0.72	2,976
Resolution <i>Naïve Bayes</i>	0	0.90	1.00	0.95	2,574
	1	0.91	0.31	0.46	402
	Average/Total	0.90	0.90	0.88	2,976
Message Format <i>Logistic Regression</i>	0	0.91	0.90	0.91	2,557
	1	0.44	0.45	0.44	419
	Average/Total	0.84	0.84	0.84	2,976
Playfulness <i>Logistic Regression</i>	0	0.81	0.82	0.82	2,281
	1	0.39	0.38	0.38	695
	Average/Total	0.71	0.72	0.71	2,976

For binary variables, 0 indicates the absence of the variable, and 1 indicates its presence. *Support* is the number of true instances of each variable category.

For message type, Linear SVC has the highest median and mean F1 score: 0.72 (+/- 0.02). The F1 scores for each category reveal that it is not very reliable for identification or

other, but is quite reliable for division. It was used to predict message type in the population, but only the results for division can be interpreted.³⁹

Naïve Bayes is the best at predicting resolution—its average F1 score is 0.88 (+/- 0.01), and its average precision for predicting resolution after 10-fold cross-validation is 0.91. Its average recall for predicting no resolution is a perfect 1.00 (which means that it labels all tweets containing no resolution correctly). Its recall for predicting resolution is quite low, as it only captures a subset of tweets with calls to action. Upon closer inspection, these calls to action include variants of ‘please RT’, ‘read and RT’, voting in the referendum, taking online quizzes, and ‘spreading the word’. Many of these tweets are automated, and multiple instances appear in the dataset. This is why Naïve Bayes can detect them so precisely—the repetition of common linguistic patterns really helps the classifier. It was used to predict the presence of resolution in the population, with the caveat that it only detects a specific subcategory of the feature.

Logistic Regression is the best classifier for predicting message format, with an average F1 score of 0.44 (the average F1 score is 0.39 for Linear SVC and 0.22 for Naïve Bayes). Overall F1 scores are 0.83-0.84 for all classifiers. As accuracies for predicting message format were so low, a Voting Ensemble was not created. Given the complexity of the variable, this is to be expected. It is almost impossible to detect lists, as most users do not include punctuation (commas) in their tweets. Contrast is often from the juxtaposition of text and image, which is not part of the training data for the classifier. Moreover, there is an infinite range of styles for both lists and contrasts (whereas for resolution, there are some call to action patterns, as explained above). Finally, there were only 419 tweets in the training sample that contained message format: this is not enough data to capture any patterns. Thus, no classifier was used to predict the presence of message format in the population.

Logistic Regression is also the best classifier for predicting playfulness, with an average F1 score of 0.38 (the average F1 score is 0.29 for SVC and 0.26 for Naïve Bayes). Overall F1

³⁹ It should be noted that this classifier was trained on tweet text along with manually labelled position, which improves the prediction of message type given that the two variables are related (a position of ‘ambiguous’ means that message type has to be ‘other’). As the position classifier is used to label the entire population of tweets, the message type classifier was run again using position predicted by the Linear SVC classifier, and accuracies were compared with the above figures. The mean F1 score after 10-fold CV is very similar for all classifiers. For Linear SVC, the best classifier, it is 0.72 (+/- 0.02) for the classifier that is trained on predicted position, vs. 0.73 (+/- 0.02) for the classifier trained on manually labelled position. The F1 scores after 10-fold cross-validation for division, identification, and other are almost identical: 0.81 vs. 0.82 (division), 0.57 vs. 0.58 (identification), 0.62 vs. 0.65 (other). Thus, the message type classifier was used with high confidence for the population even though it was trained on manually labelled position. The overall accuracy is better than the accuracy of the position classifier multiplied by the accuracy of the message type classifier (it is not multiplicative). This is because the biggest difference in accuracy is in the ‘other’ category, which is the least important. Particular combinations of position and message type are prominent.

scores are 0.71-0.72 for all classifiers, as shown in the boxplot (Figure 3.2). Once again, a Voting Ensemble was not created due to the low accuracies. Like message format, playfulness is a highly complex variable. It encompasses sarcasm, metaphor, humour, and irony—each of which is in and of itself incredibly nuanced, does not fall into common patterns of expression, and highly context-dependent. There were only 695 tweets in the training sample that contained playfulness: not enough data to capture any patterns. Thus, no classifier was used to predict the presence of playfulness in the population.

All in all, the classifiers that were trained on the entire random sample and implemented on the population level were those for position, emotion (binary), issue (binary), message type, and resolution (when the population was labelled, the random sample used for training was removed). Two caveats must be noted: only the division category for message type can be interpreted, and resolution captures a subset of explicit calls to action. It is no surprise that classifier accuracies were the lowest for message format and playfulness. These style features are extremely complex and nuanced—they are not defined by specific keywords, and are highly context-dependent. Content features such as position and issue, on the other hand, are defined by specific keywords. Tweets containing ‘NHS’ most probably relate to healthcare, whereas references to EU bureaucracy and dictatorship suggest that sovereignty is the topic. Indeed, these keyword associations informed the manual coding process.

3.3.4 Dictionary Approach

A dictionary approach was used to automatically label two variables: question and organisation. Question is the most straightforward variable to automatically code, as it is signified by a punctuation mark (?). Thus, a Python function was written to detect this mark in tweet text. This method is highly accurate but not highly comprehensive, as questions included in media items are not detected.

Any tweet from a list of 174 hand-selected media, referendum campaign, and politician accounts was classified as organisational (the list is part of the coding scheme in Appendix C). This list was assembled from extensive research—all campaigns registered with the Electoral Commission were reviewed and their Twitter handles were added. Prominent politicians were also included: David Cameron, Nigel Farage, Boris Johnson, Michael Gove, Jeremy Corbyn, George Osborne, Daniel Hannan, George Galloway, Douglas Carswell, Suzanne Evans, Nicola Sturgeon, Richard Corbett, Caroline Lucas, Jonathan Bartley.⁴⁰ Campaign leaders

⁴⁰ These politicians were prominent in the context of the referendum because they were actively involved in campaigning, and also were active on Twitter.

(Matthew Elliott, Will Straw) and media outlets were added as well (including broadsheets and tabloids). This list was expanded to encompass 72 accounts that were manually labelled as organisational in the random sample (but had not been included in the list). In addition, tweets that contained ‘MP’ in the user description, user screenname, or username⁴¹ and tweets that contained ‘MEP’ or ‘MSP’ in the user screenname or username were also classified as organisational. These represent Members of Parliament, Members of the European Parliament, and Members of the Scottish Parliament.

3.4 Interview Strategy and Analysis

Semi-structured trace interviews with 22 Twitter users who tweeted during the referendum campaigning period were conducted to complement the tweet data. These interviews were an integral aspect of the design of this thesis. As Lewis et al. (2013) declare, big data are comprised of small data, and ‘micro messages typically come from many different users who are creating content for different purposes and for different intended audiences’ (p. 48; see also Marwick & boyd, 2011). This creates two distinct problems: the brevity of tweets necessitates more effort to determine their meaning, and tweets signify a vast range of user intentions. Such intentions are impossible to detect from manual inspection, even if such inspection is of tweets on the platform. Moreover, as Lewis et al. (2013) further observe, ‘The collective nature of Twitter means that a single tweet from an individual may be relevant at a specific moment for a specific reason, whereas the remainder of the stream may be irrelevant’ (ibid). The temporal context of a tweet is highly personal, and cannot simply be reduced to time and date. Only interviews allow for this context to be revealed.

All of the interviews occurred between 12 January and 28 March 2018, and were recorded with a ZOOM H2N portable device,⁴² which has five built-in microphones and is capable of two-channel surround sound. 15 interviews were in person (in Oxford, London, and Bristol), six were over Skype, and one was over the phone.⁴³ Although only an hour of each interviewee’s time was requested, the average duration was three hours, with a minimum of one hour and a maximum of nine hours. The length of the interviews, coupled with the face to face method of execution, afforded extremely rich insights that transcended the scope of

⁴¹ Only verified users were considered, as there are many parody MP accounts.

⁴² <https://www.zoom-na.com/products/field-video-recording/field-recording/zoom-h2n-handy-recorder>

⁴³ The phone interview was supposed to be in person, but due to an unexpected laptop breakdown just prior to the meeting, alternate arrangements had to be made.

the interview questions. Indeed, relevant questions that were not explicitly asked were answered.

3.4.1 Interviewee Selection

As the objective of the interviews was to determine the strategies that individuals and organisations employ to enhance the success of their tweets, the selection process aimed to encompass both types of accounts. Personal contacts were used to obtain individual interviewees. Organisational interviewees required more effort to secure. The websites and social media profiles of referendum campaigns were scrutinised for contact details; their directors and communications managers were targeted specifically. Email was the predominant contact method, but a research-specific Twitter account, @EURefHashtags, was also established for interview solicitation purposes. All potential interviewees were followed, which led to some accounts following back; Direct Messages could then be sent to these accounts. The response rate was far higher than expected—about half of the emails and direct messages resulted in interviews being arranged. The snowballing method was then used to acquire more interviewees, and was extremely effective. Part of the reason for the high response rate was the fact that many campaigns were still active after the referendum vote, due to the extensive controversy surrounding the exit negotiations.

As a result of these efforts, four managers of organisational accounts, five individuals, and 13 users who managed both organisational and individual accounts were interviewed (Table 3.4). It must be emphasised that they are not at all representative of users who tweeted about the EU referendum during the campaigning period—the majority are journalists, digital marketers, and political campaigners. Thus, no generalisations can be made from the interview data. The purpose was instead to establish to what degree interviewee strategies for tweet success overlap with the theory-derived rhetorical features used in the tweet coding framework. A secondary purpose was to establish how individual strategies differed from organisational ones, if they differed at all.

3.4.2 Interview Structure

A semi-structured approach was adopted for the interviews because it allows for consistency and flexibility to be balanced. As Brinkmann (2013) explains, semi-structured interviews are better than structured interviews at leveraging the ‘knowledge-producing potentials of dialogues’, by virtue of the fact that they allow for follow-up questions to be asked ‘on whatever angles are deemed important by the interviewee’ (p. 21). These angles could not have been

predicted in advance of the interview, and often lead to extremely interesting insights. Moreover, the semi-structured approach also empowers the interviewer, by allowing him or her to become ‘visible as a knowledge-producing participant in the process itself’ (ibid). As such, this approach strikes a balance between interviewing as a research instrument and interviewing as social practice, or in more conceptual terms between phenomenological *reports* (the ‘what’ of communication) and discursive *accounts* (the ‘how’ of talk) (p. 36). The role of the interviewer is twofold: to interpret the meaning articulated by the interviewee and to co-construct it.

Before the interview, all interviewees were asked to fill out a short survey that measured their political and digital activity.⁴⁴ During the interview, they were asked about their tweeting motives and practises, how they define success on Twitter, and how Twitter fits into their overall communication strategy (which encompasses other social media platforms and offline channels). Interviewees were also asked questions about how they *receive* tweets: under what conditions they liked or shared content. The interview templates are presented in Appendix D; slightly different structures were used for individuals and organisations.

A trace interview approach was also adopted for this thesis. Pioneered by Dubois and Ford (2015) in recognition of the complex and multimodal nature of political communication in new media, it is ‘an actor-centric method that employs visualizations of a user’s digital traces during the interview process’ (p. 2067). At least half of each interview was focussed on specific tweets that the interviewees had published during the referendum period, which had been sent to them in advance for review—some were successful, and others were not. The spreadsheet of tweets also encompassed summary statistics about their metadata (number of hashtags, @mentions, URLs, and media items), their average number of retweets and likes, and the interviewee’s follower and friend growth over the course of the referendum. Basic visualisations were provided of temporal patterns: bar plots of hour of day and day of week, as well as histograms and boxplots of days before the referendum; a bar plot was also provided of tweet sources. The integration of computational methods with traditional interview techniques is extremely fruitful, as it improves the quality and utility of the trace data while enabling the generation of novel insights during the interview. Dubois and Ford (2015) highlight four benefits in particular: enhancing recall and illuminating decision-making processes, validating trace data results, addressing data joining issues, and responding to ethical concerns relating to big data (p. 2087).

⁴⁴ https://oii.qualtrics.com/jfe/form/SV_ePpMgeLB0jzve4t

TABLE 3.4: Interviewees and Interview Information

NAME	TITLE & ORGANISATION	INDIVIDUAL OR ORGANISATION	INTERVIEW DATE	INTERVIEW METHOD	INTERVIEW LOCATION	DURATION (HOURS)
Alberto	Europe Editor, BuzzFeed News	Organisation	16/03/2018	Face to face	London	1
Andrew	Head of Campaigns, The Freedom Association	Both	21/02/2018	Face to face	London	3.5
Brendan	General Secretary, Labour Leave	Both	28/03/2018	Phone	N/A	3.5
David	Digital Communications Consultant, British Influence	Both	12/02/2018	Face to face	London	3.5
Denis MacShane	Former Labour Party MP and Minister of State for Europe	Individual	12/03/2018	Face to face	London	1
Elise	Chair, Association of Green Councillors	Both	06/02/2018	Face to face	Oxford	3
Graham	Social Media Officer, The Bruges Group	Both	01/02/2018	Skype	N/A	3
Harriet	Co-founder and Campaign Director, We Are Europe	Both	07/02/2018, 23/02/2018	Face to face, Skype	Bristol, N/A	4
Katherine	N/A	Individual	14/01/2018	Skype	N/A	1
Leni	Social Media Consultant, Green Party	Organisation	15/02/2018	Face to face	London	1
Logan	International Officer, Young European Movement; Director, Sheffield Students Stronger In	Organisation	06/03/2018	Skype	N/A	4
Martha Lane Fox	Board Member, Twitter; Crossbench Peer, House of Lords	Individual	12/03/2018	Face to face	London	1
Matthew Elliott	CEO, Vote Leave and Business for Britain	Both	16/03/2018	Face to face	London	2.5

Michael	Programme Director, Scientists for EU and Heathier In	Organisation	16/03/2018	Face to face	London	3.5
Nick H.	Chairman, London4Europe	Both	12/02/2018	Face to face	London	2.5
Nick T.	N/A	Individual	20/03/2018	Face to face	London	3
Paul	Co-founder and Web Administrator, UKtoStay.EU	Both	24/02/2018	Skype	N/A	9
Robert	N/A	Individual	17/01/2018	Skype	N/A	2.5
Ruthi	Social Media Manager, Oxfordshire Green Party	Both	20/02/2018	Face to face	Oxford	2.5
Sam	Journalist, InFacts	Both	12/01/2018, 26/02/2018	Face to face	Oxford	5
Sean	Co-founder and Director, If Not Now Digital	Both	07/02/2018	Face to face	Bristol	5
Will Straw	Executive Director, Britain Stronger In Europe	Both	15/02/2018	Face to face	London	1

Harriet and Sam were interviewed on two separate dates. The first interview focused on their organisational accounts, the second on their personal accounts. The duration column reflects the cumulative length of these interviews. *Robert* is a pseudonym, as the interviewee requested anonymization.

One challenge emerged consistently during the interview process: probing was often necessary to elicit more substantive explanations of why some tweets were resonant. In their initial explanation, interviewees had a tendency to ascribe success to the tweet being ‘sharable’ or ‘likeable’ (echoing the very metrics used for success), or to use high-level terms such as ‘interesting’ or ‘compelling’ (conversely, ‘dull’ or ‘boring’ were used to describe unsuccessful tweets). This reflects how complex resonance is—it is difficult to operationalise and articulate the features that contribute to it. Follow-up questions had to be asked to establish what exactly makes content ‘shareable’ or ‘interesting’. These probes mostly resulted in more specific explanations, although not always.

3.4.3 *Interview Analysis Process*

All of the audio recordings were manually transcribed. Transcription was not outsourced for two reasons: confidentiality (some of the interviewees divulged private information that they wanted off the record) and accuracy. Notes had been taken during the interviews, and facilitated the transcription process. The transcription feature of a state-of-the-art note-taking app, Otter,⁴⁵ was also tested, but it was counterproductive due to low quality.

Interview transcripts were uploaded into NVivo 12 Pro,⁴⁶ along with screenshots of all tweets that had been discussed. The coding of the transcripts was derived from the features identified in the theoretical framework as well as the structure of the interviews. Each node in NVivo represents a high-level component of the framework or a topic from the interview (e.g., ‘Content Features’ and ‘Twitter Motives, Practises, and Significance’). Features mentioned by interviewees that were not in the theoretical framework were added to appropriate nodes (e.g., ‘Famous Person’ was a sub-node of ‘Content Features’). Tweet screenshots were also coded; this greatly expedited the selection of example tweets in the results chapters.

Coding the interviewee data was not at all straightforward, as the transcripts contained numerous qualifications and inconsistencies. Two individuals representing the same campaign had irreconcilable perspectives about their target audience and Twitter objective; multiple interviewees appeared to have contradicted themselves. As frustrating as these inconsistencies might be, they are also extremely characteristic of interview data. Moreover, there is a logic behind every contradiction. For Watson (2006), ‘inconsistency’ in interview data prompted reflections ‘on the “meaning” inherent in these lapses of coherence’ (p. 367). Language is notoriously slippery, and there are no absolute truths—only relative ones. Interpreting the

⁴⁵ <https://otter.ai>

⁴⁶ <https://www.qsrinternational.com/nvivo/nvivo-products/nvivo-12-pro>

inconsistencies ultimately proved to be a fulfilling process that added more nuance to the interview analysis, which is mainly presented in Chapter 5.

3.5 Ethical Considerations

Given the personal nature of social media platforms, no large-scale examination of Twitter data can be executed without consideration of ethical implications. A tension exists between the Terms of Service documents of social media companies and user expectations of privacy, and Twitter is no exception. When users create their accounts, they have to agree to the Terms of Service, which stipulates that any content they publish on the platform can be published by other parties (including academics):

You agree that this license includes the right for Twitter to provide, promote, and improve the Services and to make Content submitted to or through the Services available to other companies, organizations or individuals for the syndication, broadcast, distribution, promotion or publication of such Content on other media and services... (Twitter Terms of Service, 2018)

Thus, reproducing tweets is legally allowed; indeed, journalists quote tweets all the time. Moreover, all tweets obtained from the Streaming API were public at the time they were sent; tweets from private accounts are not available. Bartlett et al. (2014) contend that ‘those who tweet publicly available messages in general expect a low level of privacy’ (p. 99). In fact, court cases have determined that tweeting is analogous to an act of publishing, and can be ‘prosecuted under laws governing public communications, including libel’ (ibid). Given the high stakes nature of the EU referendum and the fact that only rhetorical tweets are examined, it is also reasonable to assume that the authors of these tweets wanted their messages to be seen by as many people as possible.

All of this being said, users might be uneasy about their tweets being analysed and published without their knowledge, and researchers have an inviolable duty of care to their research subjects (British Sociological Association, 2017b). As the Digital Wildfire project at the University of Oxford⁴⁷ has established, this raises a central question: *In order to publish (anonymised) individual social media posts, is it first ethically necessary to contact the user and solicit their informed consent?* The answer is premised upon the notion of harm: its minimisation is the cornerstone of ethics in research involving human participants. To address this complex concept, the Digital Wildfire project established a risk matrix to identify who is most at harm, according to user type and tweet type (Webb et al., 2017). Low-risk users are public figures or

⁴⁷ <https://gtr.ukri.org/project/3E91917B-2C3A-4995-BC42-47726EEB47CE>

individuals who are not identifiable from their Twitter profiles; low-risk content is mundane (as opposed to provocative or antagonistic). Consent is not required from low-risk users or users who publish low-risk content, but it might be required from high-risk users or users who publish high-risk content. It is important to recognise that consent is necessary because anonymity is often not enough as a privacy-protecting measure; tweets can easily be searched on the platform and their senders identified.

To further complicate matters, tweets and user profiles are both dynamic—tweets might be deleted, and profiles can be made private or public at any time. Indeed, as noted above, about one fifth of tweets in the random sample were no longer available at the time the REST API was queried for their retweet and like counts. These complexities of new media platforms have led the British Sociological Association (BSA) to create a digital research annexe to their Statement of Ethical Practice (2017b). Their general position was adopted in this study: duty of care to research participants was exercised reflexively through ‘situational ethics’, and adapted to the needs and circumstances of each participant; there is no standard template of ethical practise that can be applied to digital research. Although the Digital Wildfire project stipulates that consent is not required from public and commercial organisations, informed consent was actively solicited from both ‘public’ and ‘private’ interviewees through an ongoing open dialogue. This consent applied to the entire research process: from when the interviewees were first contacted to when they were sent drafts of the chapter(s) that featured their quotes. A personalised email was sent to each interviewee to introduce them to the project and ask for their interest in participation: it included the context of the project, why the interviewee’s insights were important, how their contact information had been obtained, and details about the structure, content, and format of the interview (in person meetings were preferred where possible, but video or audio calls were perfectly acceptable alternatives). If the interviewee replied affirmatively, they were emailed a personalised participant information sheet and an informed consent form, and these documents were reviewed at the start of the interview (Appendix E). The informed consent form clearly stipulates that withdrawal from the research can occur at any time without penalty and without an explanation, in accordance with the BSA’s Statement of Ethical Practice (2017a). All of the participants signed and dated this form before the interview began. Although all interviews were recorded so that full and accurate quotes could be obtained, at any point during the interview participants could indicate that they wished to go off the record.

The participant information sheet for interviewees affiliated with an organisation stated that they would be referenced with regard to their title and organisation; the information sheet

for private individuals stated that they would be anonymised unless they consented to being quoted by their first name. All but one interviewee indicated during the interview itself that they were happy to be referenced by their first name. Thus, during the writing process there was an evolution from the degree of privacy stipulated in the information sheet: due to the explicit consent given and the fact that the full identity of almost all of the participants could be determined from their title and organisation, they were also referenced by their first name (full names were used for those who are publicly prominent). Moreover, drafts of the results chapters were sent to all interviewees so that they could review their quotes and indicate if they had been misrepresented or wished to have certain quotes removed; several pointed out technical errors, but none asked for quotes to be removed. Revised results chapters with the modifications were subsequently sent to the interviewees, to ensure that they were represented accurately.

A pseudonym was adopted for the interviewee who asked to be anonymised. As a further precaution, this interviewee's user screenname and username were removed from tweet screenshots, and his profile picture was replaced with the default grey icon. This anonymization process was also applied to screenshots of tweets from non-prominent users that are included as examples in Chapters 4 and 6. As none of the tweets contain high-risk content and none of the users are identifiable from their Twitter profiles, informed consent was not deemed to be strictly necessary.

Finally, ethical considerations also apply to data storage and security. All of the tweets and interview recordings are held on a password-protected private server. When the dataset is made public after the publication of this thesis to enable reproducibility, only tweet IDs will be released.

4. Factors that Influence Tweet Resonance

‘For me context is the key—from that comes the understanding of everything.’

—Kenneth Noland

This chapter focuses on the main research question posed by this thesis from a statistical perspective: what makes a message resonate with its audience on social media? To be more specific, *to what degree do discoverability, content, and style variables influence the success of a tweet?* All of the hypotheses presented in Chapter 2 will be tested.

The measure of tweet resonance adopted by this thesis is *popularity*: how often the tweet is retweeted (shared) or liked. The more popular a tweet is, the more visible it is and the more influence it can exert; this influence is especially significant in a campaign context. It is important to note that both retweets and likes are *publicly visible* metrics—they are visible not only to the messenger, but also to anyone on or off the platform (provided the tweet is not private). This has both technical and conceptual implications. The technical implication is that retweets and likes can be collected from the Twitter API, whereas private metrics such as impressions or link clicks cannot be, as noted in the previous chapter. The conceptual implication is that tweets that have already received retweets and likes have a greater likelihood of receiving more, because of the phenomenon of *social influence* (Salganik, Dodds, & Watts, 2006). The Twitter algorithm encourages this phenomenon by displaying which tweets your friends have shared or liked. This reinforces the distinction that must be made between two levels of popularity: receiving at least one retweet or like, and receiving many more. Given that the majority of tweets are not retweeted or liked, *what makes a message on Twitter get liked or shared at all?* Consequently, of the minority of tweets that *are* retweeted or liked, what are the characteristics of those that are exceptionally popular? It is clear that discoverability matters more for the first level of success, as a tweet must be seen in order to attract engagement. Because tweets that have attracted engagement are inherently more discoverable, what might differentiate between the two levels is content and style: what is being tweeted and how it is expressed could play a larger role for greater success.

This chapter will focus on the first level of popularity through a series of logistic regression models derived from a hand-coded random sample of 2,976 tweets, as well as larger scale regressions based upon the population dataset of 741,411 tweets. Supervised machine learning algorithms for text classification trained on the random sample were used to label content features at the population level (some features had to be dichotomised or excluded because accuracies were not high enough). The results from the regressions represent general characteristics associated with successful tweets, averaged across all user and tweet types.

Chapter 5 complements these results with data from interviews with individuals and organisations who actively tweeted during the referendum period: it assesses the strategies they employed to increase the popularity of their messages, and provides more context for the features included in the models. Drawing upon the regressions as well as the interview data, Chapter 6 focuses on the second level of popularity through a qualitative examination of exceptional tweets and users.

The first section of this chapter provides an overview of the tweet data: descriptive statistics of the dependent, independent, and control variables included in the models, at both the random sample and population levels. Section 4.2 presents five models for each dependent variable based upon the random sample ($N=2,976$): control variables only, control and discoverability features, control and content features, control and style features, and all variables. Section 4.3 scales up the retweet regression models to all tweets in the population dataset ($N=741,411$), after text classification algorithms were used to automate content coding. The results from the random sample and population level are compared, and overarching remarks made, in the final section.

4.1 Summary Statistics

As introduced in the previous chapter, two datasets are used in the analyses presented in this chapter: a random reservoir sample of 2,976 tweets, and the population of 741,411 tweets from which it was drawn. This population represents a subset of the data collected from the Streaming API: it is restricted to tweets published during the referendum period, excluding referendum day (15 April-22 June 2016), it contains no retweets, replies, or quotes, and it only includes tweets that contain hashtags supporting Remain or Leave. These filters increase the probability that the tweets employ rhetorical strategies and are designed to persuade, and decrease the probability that they are uninterpretable due to lack of context.

The 2,976 tweets in the random sample were hand coded for the presence of the discoverability, content, and style features introduced in Chapter 2: organisational authority, message framing and structure, emotions (enthusiasm, anxiety, aversion), relevant issues, resolution, questions, and playfulness. All of these features are treated as independent variables in the regression models. Position (whether the tweet is pro-Remain, pro-Leave, or ambiguous) was also hand coded, and is treated as a control variable.

The manually coded sample was subsequently used to train and test text classification algorithms to label the features in the population. After ten-fold cross validation, average

accuracies were calculated to assess a) the best classifier for each feature and b) whether that classifier could be implemented. As explained in Chapter 3, the cutoff for implementation was an F1 score of at least 0.77 for predicting the presence of the feature. Message format and playfulness were not classified at the population level because accuracies were too low; issue was dichotomised, and the three emotion variables were combined into one binary variable. Only the ‘division’ category of message type was included in the population-level models, as accuracies for ‘identification’ and ‘other’ were too low.

4.1.1 *Dependent Variables*

Chapter 2 contextualises and justifies the two outcome variables in the models: whether the tweet was shared or liked. The previous chapter describes how they were measured directly from the Twitter REST API (for the sample only), and indirectly from the tweet JSON files scraped from Twitter’s Streaming API (for both the sample and the population). Table 5.1 indicates that the retweet counts produced by the latter method are underestimates of the more up-to-date counts provided by the former method: the mean and maximum values are slightly lower for the retweet counts computed from the dataset for the random sample, as some tweets continued to be retweeted after 1 July 2016. Given the missing data issue, however, and the fact that the difference is not too extreme (the API mean is only 0.77 retweets greater than the mean computed from the dataset), the retweet count computed from the dataset is used for all statistical analyses at both random sample and population levels.

TABLE 4.1: Descriptive Statistics for Continuous Dependent Variables

	RANDOM SAMPLE				POPULATION	
	RETWEET COUNT	LIKE COUNT	RETWEET COUNT (API)	LIKE COUNT (API)	RETWEET COUNT	LIKE COUNT
Valid N	2,976	1,170	2,427	2,427	741,411	278,143
Mean	6.79	20.66	7.56	10.98	4.22	9.01
SD	162.17	461.56	185.45	377.87	48.03	86.64
Median	0	1	0	0	0	1
Minimum	0	0	0	0	0	0
Maximum	8,796	15,774	9,098	18,607	14,258	18,107

As manually computed like count and retweet count from the API are not used in the models, they are presented in grey.

As noted in Chapter 3, the like counts provided by this method are more problematic, as they are incomplete (no like count is provided for tweets that were not shared). Because the valid *N* is much smaller, *average* like counts are *overestimates* compared to those provided by the

API—tweets that received no retweets are not included, many of which also received no likes (they thus would have significantly lowered the average). Table 4.1 reveals that the average like count from the API is 9.68 lower than the average like count computed from the dataset at the sample level. Moreover, the median like count from the API is 0, whereas the median like count computed from the dataset is 1, and the difference between the maximum like counts is 2,833. Due to the severity of this discrepancy and the missing data issue, the like count from the API is used for the statistical analyses of the random sample, and like count models are not produced at the population level.

The descriptive statistics illuminate that both retweet count and like count are extremely right-skewed: a very small proportion of tweets received a large number of shares and likes; at least half of the tweets received no shares or likes at all, as the median for both variables is 0 (like counts computed from the dataset are not considered here for the reasons given above). 0 is also the mode for both variables. The fact that their mean and maximum values are quite high compared to their medians (6.79 and 8,796 for retweet count and 10.98 and 18,607 for like count at the sample level; 4.22 and 14,258 for retweet count at the population level) indicates that the minority of successful tweets were very successful. This reflects the heavy-tailed distribution that is commonly seen in internet data (Java, Finin, Song, & Tseng, 2007). It should also be noted that the skew is greater in the population dataset than it is in the random sample—the population has a lower mean and a higher maximum for retweet count, which suggests that it contains more extreme outliers as well as more tweets that were not shared. These outliers will be examined in Chapter 6, as they represent a different level of resonance.

Moreover, the high standard deviations (162.17 for retweet count and 377.87 for like count at the sample level) and large ranges (0 to 8,796 for retweet count and 0 to 18,607 for like count) reveal that the distributions for both variables are not very concentrated around their means—there is a vast ‘spectrum’ of success. The standard deviation for retweet count at the population level is significantly lower (48.03), indicating less dispersion, but the range is greater due to more extreme values (0 to 14,258). It is interesting to note that the distribution of like counts is more right-skewed and more spread out than the distribution of retweet counts; the mean is also noticeably higher. This indicates that liking a tweet is more common than retweeting a tweet. Users are more selective about what they share because retweets are published in their main feed and the Home feeds of their followers, whereas likes are far less salient—they are stored in the Likes tab on profiles and not shared with followers.

Because of the extreme right skew of both distributions and the fact that the median of both variables is 0, they were dichotomised for logistic regression: *which features increase the likelihood of a tweet receiving any shares or likes?* 60.69% of tweets in the random sample and 62.48% of tweets in the population did not receive a retweet, and 62% of tweets in the random sample did not receive a like (using the retweet count computed from the dataset and the like count from the API). Thus, it is clear that even obtaining one like or share should be considered as success in and of itself. As this like or share becomes part of the tweet object, visible to all users who see the tweet in their feed, it becomes far easier for a tweet to receive another like or share after the initial threshold is met.

4.1.2 Independent Variables

The predictor variables in this analysis fall into three thematic categories, as introduced in Chapter 2: discoverability (tweet and user attributes), content, and style. These categories are all components of rhetoric. Their crosstabulations and stacked bar charts with the dependent variables, which provide a general sense of whether they might be related to tweet success, are presented in Appendix F.

As explained in Chapter 2, many affordances of Twitter increase discoverability: hashtags, @mentions, visual media (images, videos, GIFs), and hyperlinks. All four of these features represent Twitter strategies that will be detailed and contextualised in the next chapter. Hypotheses 1a, 1b, 1c, and 1d test the impact of discoverability features on tweet resonance.

Interestingly, apart from hashtags, Twitter-based discoverability features are not very pervasive (all tweets have at least one hashtag because they were harvested from Twitter's API via a hashtag search). The range for hashtag count is 12 for the sample and 20 for the population. Although the overall distribution of this variable is right-skewed, sizable proportions of tweets contain one, two, or three hashtags (Table 4.2). Given this distribution, hashtag count is treated as a continuous variable.

TABLE 4.2: Frequencies for Hashtag Count

	RANDOM SAMPLE		POPULATION	
	FREQUENCY	PERCENT	FREQUENCY	PERCENT
1	929	31.22	228,107	30.77
2	874	29.37	222,643	30.03
3	533	17.91	127,588	17.21
4+	640	21.51	163,073	21.99
VALID TOTAL	2,976	100.00	741,411	100.00

FIGURE 4.1: Tweet with Numerous Hashtags



Retrieved 31 August 2018 from Twitter. Screenshot by author.

Although including hashtags increases tweet discoverability, there might be a boomerang effect after a certain number, as the presence of too many hashtags is indicative of spam (and reduces the space for substantive content due to the 140-character limit). Indeed, the tweet with 21 hashtags was sent by a pro-Trump user whose account has since been deleted, and lists abbreviations of US states: ‘#NV #CO #IA #OH #NH #VA #NC #FL #OR #AR #NM #MN #GA #GA #PA #WI #MI #NY #CA WAKE UP AMERICA! #Brexit #VoteLeave <https://t.co/tfHXgCWTwS>’. Another user sent three tweets with at least 17 hashtags, all of which consist of sentences of hashtagged words; one example is in Figure 4.1. Because the hashtags are blue and preceded by the hash symbol, they detract visually from the message substance and create a spam-like effect.

TABLE 4.3: Frequencies for Mention Count

	RANDOM SAMPLE		POPULATION	
	FREQUENCY	PERCENT	FREQUENCY	PERCENT
0	2,358	79.23	580,350	78.28
1	466	15.66	121,938	16.45
2	108	3.63	27,442	3.70
3+	44	1.48	11,681	1.58
VALID TOTAL	2,976	100.00	741,411	100.00

Almost four out of five tweets in the sample and population do not contain any @mentions (Table 4.3). As the proportions of tweets containing one or two mentions are not negligible, mention count is also treated as a continuous variable. The maximum number of mentions in any tweet in the sample is six; in the population it is 10. Many of the mentioned users are prominent individuals or organisations: e.g., @StrongerIn, @David_Cameron, @SkyNews, @BBCNews, @UKIP, @Nigel_Farage, @BorisJohnson, @vote_leave, and @eddieizzard. Figure 4.2 presents one of the tweets with the most mentions from the

population dataset, which demonstrates this technique: popular news outlets, petition sites, and the leader of the Labour Party are included.

FIGURE 4.2: Tweet with Numerous Mentions



Retrieved 31 August 2018 from Twitter. Screenshot by author.

FIGURE 4.3: Tweet with Multiple Media Items



Retrieved 31 August 2018 from Twitter. Screenshot by author.

Over three-quarters of tweets in the sample (77.39%) and population (77.08%) contain no media items. This variable was dichotomised because only 1.01% of tweets in the sample and 1.31% of tweets in the population contain at least two media items (four is the maximum). From a design perspective, including more media items does not significantly increase the visibility of a tweet, because the visual space occupied by one media item is similar to the space occupied by multiple media items. Subsequent images and videos are represented as smaller icons on the right side of the first image, as demonstrated by Figure 4.3.

TABLE 4.4: Frequencies for Visual Media Types

	RANDOM SAMPLE		POPULATION	
	FREQUENCY	PERCENT	FREQUENCY	PERCENT
Image	649	21.81	163,322	22.03
Video	18	0.60	3,813	0.51
GIF	6	0.20	2,794	0.38

As Table 4.4 reveals, images are by far the most popular visual media type, followed by videos and animated GIFs. These differences make sense, because much more effort is required to produce and upload a video (it is a larger file type). Images and videos are easy to capture on a mobile device, whereas GIFs have to be taken from a preexisting library or created via an online tool, and have pop culture connotations (which restricts their usage and appeal). Because there are qualitative distinctions among the visual media types, they are included in the population-level models as three separate variables to assess the influence of each on tweet resonance. Frequencies are too low at the sample level, so the dichotomous version of visual media count is used.

TABLE 4.5: Frequencies for Hyperlink Count

	RANDOM SAMPLE		POPULATION	
	FREQUENCY	PERCENT	FREQUENCY	PERCENT
0	2,016	67.74	488,574	65.90
1	948	31.85	249,921	33.71
2	11	0.37	2,847	0.38
3	1	0.03	58	0.007
4	N/A	N/A	11	0.001
VALID TOTAL	2,976	100.00	741,411	100.00

Approximately two-thirds of tweets in the sample and population contain no hyperlinks (Table 4.5). This variable was also dichotomised because only 0.40% of sample tweets and 0.39% of population tweets contain more than one hyperlink. As with media items, a larger number of URLs does not increase tweet visibility, because a preview can only be

generated for one URL. The majority of URLs are to news articles or YouTube videos. Remain campaigners mostly linked to the *Guardian*; Leave campaigners favoured the *Daily Mail*, the *Daily Mirror*, and the *Spectator*. One of the most popular YouTube videos was *Brexit the Movie*; many Leave campaigners linked to the full documentary or to select excerpts.

The other feature that increases discoverability relates to the *user* as opposed to Twitter affordances: whether or not the user represents an organisation. In the context of the EU referendum, media and political organisations are the most prominent (journalists, news outlets, politicians, registered campaigns) and exert more authority than ordinary individuals. Due to their public prominence, their tweets are inherently more attention-arresting. Hypothesis 2 tests the impact of organisational authority on tweet resonance.

Just under five percent of sample tweets and approximately three percent of population tweets are from organisations; this variable is dichotomous. The population frequency is an underestimate that is skewed towards politicians (accuracy was prioritised over completeness). Even at the sample level, journalists and news organisations are far less prevalent than politicians: a wide range of MPs from both sides are represented. The most prevalent organisations are the official campaigns, Vote Leave and Stronger In. Both had press accounts for communication with the media as well as general accounts for communication with the public. The press accounts sent more tweets than the general accounts, and for Stronger In the difference was more extreme: there are 3,025 @StrongerInPress tweets and 504 @StrongerIn tweets in the population dataset; by way of contrast, there are 1,201 @Vote_LeaveMedia tweets and 1,129 @vote_leave tweets. The greater activity of the press accounts is no coincidence, and will be explained in Chapter 5.

The second thematic category for predictor variables is *rhetorical content*: what is actually expressed in the tweet? Message type, or the framing of the message, is the first component of tweet substance.⁴⁸ Does the tweet contain positive campaigning (identification rhetoric) or negative campaigning (division rhetoric)? The second component of tweet content is message tone, or the presence of high-arousal and politically-relevant emotions: enthusiasm, aversion, and/or anxiety. To complement emotion, referendum-related issues is the third component. Does the tweet reference one of the salient themes in the referendum—economy and labour, immigration, sovereignty, security, or healthcare—or another relevant issue? Finally, the fourth predictor variable related to rhetorical content is resolution. Does the tweet explicitly ask for a specific action to be taken (e.g., voting in a poll, attending an event, retweeting to followers)?

⁴⁸ Although Atkinson (1984) used the term message type, *message framing* will be used throughout this thesis because it is more transparent.

All of the content features were automatically labelled at the population level because the classification algorithms achieved high enough accuracies. However, as mentioned above, issue had to be dichotomised, the three emotion variables were combined into one, and only the category of ‘division’ in message framing is interpretable. This makes sense computationally, because these are the features for which keyword analyses are applicable—they are closely connected to the presence or absence of specific terms (message framing is slightly more complicated, as it is based on keywords relative to position). Tweets containing ‘NHS’ are likely to be about immigration; tweets containing ‘disgusting’ are likely to exhibit aversion. Style, on the other hand, is not closely connected to keywords. Hypotheses 3a, 3b, 4a, 4b, 4c, 5, and 6 assess the impact of content features on tweet resonance.

Overall, rhetorical content features are more present than discoverability features. As Table 4.6 reveals, over half of the tweets in the sample and population contain division;⁴⁹ identification is present in one-fifth of the sample (note that if a tweet contains both identification and division, it was coded as division). A quarter of sample tweets contain neither. This indicates that a large majority of users are politically oriented, and that critique of the opposition is far more popular than expressing the common values of one side.

TABLE 4.6: Frequencies for Message Framing

	RANDOM SAMPLE		POPULATION	
	FREQUENCY	PERCENT	FREQUENCY	PERCENT
Identification	599	20.13	120,223	16.22
Division	1,619	54.40	521,446	70.33
Other	758	25.47	99,742	13.45
VALID TOTAL	2,976	100.00	741,411	100.00


The values for ‘identification’ and ‘other’ for the population dataset cannot be interpreted, as accuracies for predicting these categories were too low.

Most tweets that express divisive rhetoric criticise the EU (and are thus from Leave supporters), but a prominent subcategory exhibit name-calling, as they castigate prominent politicians on both sides: namely David Cameron, Boris Johnson, and Nigel Farage. Tweets that express identification rhetoric, on the other hand, are far less personal—most praise the EU or express national pride. A distinct subcategory report on campaign events (e.g., outdoor rallies, leafleting) and include images of campaigners in action. Exemplary tweets for both strategies are presented in Figures 4.4 and 4.5; Figure 4.3 also serves as a good example of positive campaigning. Because there are three mutually exclusive categories, message framing

⁴⁹ The discrepancy between the proportions of division tweets in the sample and population strongly suggests that the classifier overestimated the prevalence of this feature, but its accuracy was 0.82.

is treated as a discrete variable in the analysis (dummy variables for the categories were created for the regressions).

FIGURE 4.4: Example of Division



Retrieved 31 August 2018 from Twitter. Screenshot by author.


FIGURE 4.5: Example of Identification



Retrieved 31 August 2018 from Twitter. Screenshot by author.

One out of every three tweets (32.80%) in the sample expresses enthusiasm, almost one out of every five (18.25%) expresses anxiety, and over a third of tweets (36.73%) express aversion; frequency bar plots are in Figure 4.6. It is interesting to note that enthusiasm and aversion are about equally prevalent, and that anxiety is not very prevalent despite accusations of fearmongering from both sides (Leave campaigners labelled the Remain campaign as Project Fear). Moreover, despite the large volume of divisive rhetoric, many tweets in the sample contain positive sentiment.

FIGURE 4.6: Bar Plots of Emotion Frequencies



Because a tweet can contain more than one emotion, each emotion is treated as an independent dichotomous variable at the sample level. While almost two-thirds of the tweets (64.78%) exhibit only one of the three emotions, 6.38% contain anxiety and aversion, 3.23% contain enthusiasm and aversion, 1.18% contain enthusiasm and anxiety, and 0.47% contain all three emotions. Only one quarter of the tweets (23.96%) express either a weaker sentiment (e.g., sadness) or no sentiment at all. At the population level, the three emotion variables were combined into one binary variable: 86.47% of tweets contain at least one of the three.⁵⁰ It is clear that the EU referendum was a highly affective event. Examples of enthusiasm, anxiety, and aversion are provided in Figures 4.7, 4.8, and 4.9, respectively.

FIGURE 4.7: Example of Enthusiasm and No Theme



Retrieved 31 August 2018 from Twitter. Screenshot by author.

⁵⁰ Like the proportion of population-level tweets exhibiting division, this is probably an overestimate given its discrepancy with the sample proportion (76.04%). The accuracy of the best classifier for predicting emotion was 0.85.

FIGURE 4.8: Example of Anxiety and Healthcare



Retrieved 31 August 2018 from Twitter. Screenshot by author.

FIGURE 4.9: Example of Aversion and Immigration



Retrieved 31 August 2018 from Twitter. Screenshot by author.

Almost half of the tweets in the sample contain no substantive issue. Economy and labour is the most prevalent issue, followed by 'other', immigration, sovereignty, security, and healthcare (Table 4.7). The most common themes in the 'other' category are the environment, EU corruption, and education/research; more shall be said on them in the next chapter, as some interviewees represented niche interest groups. Figure 4.8 contains healthcare, Figure 4.9 contains immigration, and Figure 4.7 is an example of no theme. Although a tweet can contain multiple issues, to simplify the coding process the most salient issue was chosen. Thus, as the issue categories are mutually exclusive, it is treated as a discrete variable. At the population level, it is dichotomised due to low classifier accuracies for the multiclass version of the variable; 'other' is grouped with the five recognised issues. Just over half of tweets express a theme, which reflects the proportion in the sample.

TABLE 4.7: Frequencies for Tweet Issue in Random Sample

	FREQUENCY	PERCENT
Economy and labour	526	17.67
Healthcare	61	2.05
Immigration	252	8.47
Security	130	4.37
Sovereignty	244	8.20
Other	316	10.62
None	1,447	48.62
VALID TOTAL	2,976	100.00

About one in seven tweets in the sample and 4.20% of tweets in the population express a call to action.⁵¹ The most popular action is to vote for either Leave or Remain. Other popular imperatives are to read an article or blog post, watch a video (often on YouTube), and retweet content—all actions pertaining to the internet. Apart from casting a vote on 23 June, imperatives relating to offline action include registering to vote and participating in a campaign event, but these are far less pervasive. Resolution is treated as a dichotomous variable, and an example of a tweet containing this feature is given in Figure 4.10.

FIGURE 4.10: Example of Resolution



Retrieved 31 August 2018 from Twitter. Screenshot by author.

The third and final thematic category for predictor variables is *rhetorical style*, the complement to rhetorical content: how is the message expressed? Message format, like message framing, is derived from Atkinson's (1984) framework (see Chapter 2). It represents the presence of either or both of the two stylistic devices that Atkinson found to be most correlated with success: contrast and lists. The second component of style is the presence of

⁵¹ As explained in Chapter 3, resolution at the population level represents a subset of the feature—tweets that exhibit common and unambiguously action-oriented phrases such as 'please RT', 'read and RT', and 'spread the word'. Accuracy was again prioritised over completeness: while the classifier's precision for predicting resolution was 0.91, its recall was only 0.31. This explains the noticeable discrepancy between the sample and population proportions.

questions of any kind, and the third component is playfulness—non-literal use of language. Only question was coded at the population level, as it is signified by a punctuation mark (a rule-based approach was thus appropriate). As noted above, message format and playfulness transcend keyword analysis—there are no consistent terms that signify either feature, and classifier accuracies were extremely low. Indeed, both features are arguably more effective if they are more idiosyncratic (if, e.g., a joke has already been heard, it is less humorous). Hypotheses 7a, 7b, 8, and 9 test the relationship between stylistic features and tweet success.

TABLE 4.8: Frequencies for Style Variables


	RANDOM SAMPLE		POPULATION	
	FREQUENCY	PERCENT	FREQUENCY	PERCENT
Message format	419	14.08	N/A	N/A
Question	445	14.95	92,661	12.50
Playfulness	695	23.35	N/A	N/A

FIGURE 4.11: Example of Message Format (List)



Retrieved 31 August 2018 from Twitter. Screenshot by author.

FIGURE 4.12: Example of Message Format (Contrast)



Retrieved 31 August 2018 from Twitter. Screenshot by author.

All three style features are treated as dichotomous variables, and their frequencies are presented in Table 4.8. One out of every seven tweets in the sample contains a list, a contrast, or both. Popular lists are of reasons to Remain or Leave, or of supporters of either campaign (Figure 4.11). Many instances of contrast explicitly or implicitly compare the campaigns (Figure 4.12). About the same proportion of tweets in the sample and slightly fewer tweets in the population contain a question, many of which are rhetorical.⁵² A quarter of tweets in the sample play with language, often with humorous effect. Sarcasm is the predominant form of playfulness; literary devices such as metaphor are not as pervasive. A rough corroboration of the prevalence of sarcasm is the fact that 27.73% of the tweets that exhibit divisive rhetoric also exhibit playfulness, whereas only 10.02% of tweets that exhibit identification rhetoric are also playful. Figure 4.13 presents an example of playfulness coupled with negative framing.

FIGURE 4.13: Example of Playfulness and Division



Retrieved 31 August 2018 from Twitter. Screenshot by author.

4.1.3 Control Variables

The control variables in this analysis fall into four categories: *user influence*-based (follower count), *time*-based (time of day, day of the week, days before referendum), *user experience*-based (status count), and *practically interesting* (position).

Follower count is the most obviously influential control variable on tweet resonance—as tweets are published in the feeds of a user’s followers, the more followers a user has, the greater the visibility of his or her tweets. Like the dependent variables, it is extremely right-

⁵² The population-level frequency represents an underestimate, as questions contained in media items or URL previews are not included (more sophisticated machine learning algorithms would be necessary for this). They were included in the sample.

skewed; the median (508 in the sample and 471 in the population) is quite far below the mean (4,241.88 in the sample and 4,600.39 in the population). The sample standard deviation of 31,861.62 and population standard deviation of 74,662.67 indicate that the distributions are very widely dispersed around the means—far more so at the population level, due to the greater variety of data points included (Table 4.9). This variety is reflected in the ranges: the maximum follower count is 1,480,967 in the sample and 22,752,666 in the population; the minimum in both datasets is 0. Follower count is kept continuous, but it is $\log_{10}$ transformed because it has a non-linear relationship with the dependent variables. An increase of one follower has an extremely insignificant effect on whether a tweet from the user is liked or shared; a more substantive increase is necessary to improve the likelihood of engagement. Indeed, logit coefficients of the non-transformed variable and the square root version of the variable are close to 0 in all models, which translates to odds ratios that are close to 1. The odds ratios of the $\log_{10}$ transformed variable are noticeably higher (Tables 4.11 and 4.12).

TABLE 4.9: Descriptive Statistics for Continuous Control Variables

	RANDOM SAMPLE			POPULATION		
	FOLLOWER COUNT	ACCOUNT AGE	DAYS BEFORE 23 JUNE ⁵³	FOLLOWER COUNT	ACCOUNT AGE	DAYS BEFORE 23 JUNE
Valid N	2,976	2,976	2,976	741,411	741,411	741,411
Mean	4,241.88	3.87	17.16	4,600.39	3.90	17.68
SD	31,861.62	2.58	17.05	74,662.67	2.59	17.17
Median	508	4.22	12	471	4.25	13
Minimum	0	0	1	0	0	1
Maximum	1,480,967	9.54	69	22,752,666	9.95	69

Another potential control variable related to user influence is verified status. If an account is verified (indicated by a blue badge with a white check mark next to the username and evidenced in Figure 4.3), it is of ‘public interest’ and the user has some degree of social prominence (Twitter Help Centre, 2017c). Celebrities, politicians, companies, organisations, journalists, and leaders in various fields are the most popular categories of verified users. The purpose of verification is to confirm that the user is authentic (i.e., not a fake account), and thereby confers a certain degree of authority and trust. This is important given the large number of fake accounts on Twitter; more prominent individuals are more likely to be impersonated. Twitter selects which accounts can be afforded verified status, and automatically grants this status to some users. Anyone can apply for verification, but they have to explain

⁵³ The number of days before the referendum that the tweet was published.

their ‘impact in the field’ (if the account represents an individual) or their ‘mission’ (if the account represents an organisation) (ibid). Thus, verified status is a rough proxy for public prominence. Only 2.25% of users in the random sample and 2.00% of users in the population are verified. This variable is not included in the final regression models because it has a strong and statistically significant relationship with whether the tweet is from an organisation, as revealed by the correlation matrix of independent and control variables at the sample level (Appendix G). The Pearson’s correlation coefficient is 0.585. Given the types of users who are verified, this is to be expected.

Friend count, the number of accounts that a user follows, is a third potential measure of user influence but is not included in the models. The first reason for its exclusion is also statistical: there is an extremely strong and significant correlation between follower count and friend count; the Pearson coefficient for the $\log_{10}$ transformed version of both variables is 0.804 at the sample level (Appendix G). This collinearity is unsurprising because following users is one strategy of acquiring more followers, as users are notified about who follows them (and might decide to follow back). The second reason for the exclusion of friend count is technical: a user’s ‘friends’ do not see his or her tweets unless they are following the user.

Time-based variables must be controlled for because activity on Twitter is not constant: the hour of day, day of the week, and date of publication of a tweet might increase or decrease its visibility. Unless the user is as prominent as Donald Trump, there is little chance of his or her tweets being shared or liked if they are published at 3am. Many social media analytics companies have commented on the optimal time and day to post content: the general conclusion is that afternoons and weekdays are more ideal (Ellering, 2017). Tweets published closer to the referendum date (23 June) might have had a larger audience, as users were more active on Twitter in the weeks and days leading up to the vote.

Hour of day was binned into four categories and is treated as a discrete variable: small hours (1am-6am), morning (6am-12pm), afternoon (12pm-6pm), and night (6pm-1am).⁵⁴ Keeping the 24 original categories is impractical because they are too fine-grained; there are no substantially meaningful differences between hours in specific bins. These bins reflect distinct behaviour patterns—the average adult tends to sleep during the small hours, commute to and begin work in the morning, work in the afternoon, and return home in the evening. A population-level crosstabulation of hour and whether the tweet was shared reveals that tweets sent between 1 and 6am have, unsurprisingly, the lowest retweet frequencies; this trend is also

⁵⁴ Before hours were binned, they were rounded to the nearest hour: e.g., 5.05pm is 5pm, but 5.35pm is 6pm. A tweet sent at 5.35pm would thus be classified as ‘night’.

reflected in the crosstabulation of the binned variable with whether the tweet was shared (Appendix F). Differences between the other time of day categories are subtle. A population-level bar plot of average retweet counts for every hour of the day (Figure 4.14) reveals that tweets sent between 2.30 and 3.30pm have the highest average (5.52), followed by tweets sent between 7.30 and 8.30am, 7.30 and 8.30pm, 6.30 and 7.30am, and 6.30 and 7.30pm (the mean retweet counts are 4.79, 4.67, 4.62, and 4.61, respectively). The modal category of the binned variable is night (40.93%); morning and afternoon are similar in frequency (25.50% and 29.70%), and the least frequent category by a wide margin is small hours (3.87%) (Figure 4.15).

FIGURE 4.14: Bar Plot of Mean Retweet Count for Hour of Day






FIGURE 4.15: Bar Plot of Binned Hour Frequencies



Because there are qualitative differences between different days of the week (which will be illuminated in the next chapter), this variable was kept categorical with seven classes. Figure 4.16 shows their frequencies at the population level: Tuesday and Wednesday are noticeably the most popular (the reason for this will be disclosed in the discussion of the next variable); Friday, Saturday, and Sunday are the least. This distribution is reflected in the random sample.

FIGURE 4.16: Bar Plot of Day Frequencies



The number of days before the referendum that the tweet was published is a continuous variable, and is quite skewed to the right. Although the dataset represents a 68-day campaigning period (15 April–22 June 2016), the average tweet in the sample was published 17.16 days before the referendum, and the median number of days is 12 (Table 4.9). The population-level distribution is strikingly similar, with a median of 13 and a mean of 17.68. This is unsurprising, as users were more incentivised to tweet about the referendum as the voting date drew closer—which helps to explain why Tuesday and Wednesday are the most frequent days of the week (23 June was a Thursday). Dispersion is reasonably large relative to the 68-day range: the sample and population standard deviations are both close to 17. Tweets from 23 June were removed because most were not trying to persuade (and instead just reported the results).

User experience is the third control variable category. It is important to take into account because the more time a user has spent on Twitter, and the more tweets the user has published, the better he or she is at composing them. Account age and status count (the number of tweets a user has sent) are both proxies for Twitter experience provided by tweet metadata. As they overlap conceptually, only account age is used in the models because it is less correlated with followers count: the Pearson coefficients are 0.365 and 0.668 in the sample, and both are highly significant (Appendix G). It must be noted, however, that status count is

arguably a better proxy because it indicates *active engagement* with the platform, as opposed to passive engagement: i.e., just reading tweets (such a practise is called *lurking*). As of December 2017, 1.3 billion Twitter accounts had been created, but only 550 million users had sent a tweet; 44% of new accounts are abandoned (Smith, 2017). However, some users have extremely high status counts because they use automated tools to produce tweets at substantial volumes. The most prolific account in the sample produced 2,253,139 tweets; the maximum in the population is 6,206,247. It is impossible for these numbers to have been manually achieved.

Account age is expressed in years. Unlike the other continuous control variables, it is not extremely right-skewed: the median (4.22 in the sample and 4.25 in the population) is higher than the mean (3.87 in the sample and 3.90 in the population). The minimum in both datasets is zero; the sample maximum is 9.54 and the population maximum is 9.95 (this is close to the absolute maximum, as Twitter was launched in March 2006, ten years prior to the referendum period). Given these ranges, the standard deviations are quite moderate: 2.58 in the sample and 2.59 in the population (Table 4.9). Account age is kept as a continuous variable.

The last control variable is the position expressed by the tweet: does it support leaving the EU, remaining in the EU, or is it ambiguous?⁵⁵ Given the highly political and extremely contested nature of the referendum, it is practically interesting to examine which campaign was more effective, even though position is not theoretically relevant to this thesis. It is also interesting to compare success on Twitter with the real life outcome of the referendum: Leave might have won the vote, but did they also ‘win’ on social media? Broad comments will be made on this matter in Chapter 7. For now, it will be noted that the population and sample-level crosstabulations of position with whether the tweet was shared reveal that Leave and Remain tweets are retweeted at similar frequencies, and that both categories are more often shared than ambiguous tweets (Appendix F). At the sample level, Remain tweets are more likely to be *liked* than Leave tweets, and both categories are more often liked than ambiguous tweets. What is interesting is that in the sample, Remain tweets also have double the average retweet count and quadruple the average like count of Leave tweets. At the population level, however, Leave tweets have a higher mean retweet count: the difference, as shown by Figure 4.17, is just over one tweet. This is due to the fact that the most shared and liked tweet in the sample was exceptionally successful and endorsed Remain—it received 8,315 more retweets and 15,496 more likes than the next most successful tweet, which was Leave-oriented, and will be discussed in Chapter 6. Tweets that are ambiguous have a substantially lower average of shares and likes at the sample level (Figure 4.18). Leave tweets, however, were more pervasive:

⁵⁵ This category includes those that are neutral, undecided, or dislike both sides equally.

half of the tweets in the sample and population support Brexit, which is about ten percentage points higher than those in favour of Remain (Table 4.10). Position is treated as a discrete variable with three categories in the analysis.

FIGURE 4.17: Bar Plot of Mean Retweet Count for Position (Population)




FIGURE 4.18: Bar Plot of Mean Retweet and Like Counts for Position (Sample)




TABLE 4.10: Frequencies for Position

	RANDOM SAMPLE		POPULATION	
	FREQUENCY	PERCENT	FREQUENCY	PERCENT
Leave	1,509	50.71	396,870	53.53
Remain	1,166	39.18	303,072	40.88
Ambiguous	301	10.11	41,469	5.59
VALID TOTAL	2,976	100.00	741,411	100.00

4.2 Sample-Level Models

This section presents ten sample-level logistic regression models that help explain tweet resonance, as measured by retweet and like counts; five models pertain to each dependent

variable (Table 4.11 and Table 4.12). The first model in each table serves as a baseline, as it includes only the control variables described in Section 4.2.3. The subsequent models are each focussed upon a category of features introduced in Section 4.2.2: discoverability, rhetorical content, and rhetorical style. The final model in each table includes all of the variables.

Apart from the extremely skewed distributions of both dependent variables, another reason for choosing logistic regression is that the independent variables contain a large number of zeros, as described above (i.e., many explanatory features are *not* present in most tweets). Moreover, receiving even one like or share is a form of success in and of itself, given that the majority of tweets receive no likes or shares.

4.2.1 *Regression Diagnostics*

As described above, collinearity was diagnosed through matrices showing Pearson pairwise correlation coefficients for all independent and control variables; their associated p values were also computed. The models were refined accordingly so that highly correlated variables were not included. Two of the substantively interesting independent variables are collinear: emotion and message framing. There are statistically significant relationships between all combinations of their subcategories that are at least moderate in magnitude; the strongest correlations are between identification and enthusiasm (0.611) and division and aversion (0.594) (Appendix G). This is not surprising, as positive framing often involves positive sentiment, whereas negative framing often involves negative sentiment (especially anger). To address this collinearity and to reflect the population-level regression, emotion was dichotomised in the models that include message framing (the correlation coefficients of the binary emotion variable dropped to 0.168 with identification and 0.224 with division). Separate full models were created with the three discrete emotions and without message framing, but as none of the emotions are even weakly significant for retweets, they are not used as the final models and are instead presented in Appendix G.⁵⁶ It should be noted that in the like model, emotion and anxiety are both weakly significant.

Tests for multicollinearity and the presence of outliers were also run to confirm the appropriateness of the models. To informally detect the presence of collinear variables, variance inflation factors (VIFs) were calculated. VIFs close to 1 indicate no collinearity; Hair, Black, Babin, and Anderson (2010) suggest a threshold of 10. The position categories Leave and Remain (included as dummy variables in the regression) have noticeably higher VIFs than

⁵⁶ Models were also created including emotion as a binary variable and excluding message framing, but even this version is not significant.

the other variables: they score 4.85 and 4.55 in the sample and 6.92 and 6.25 in the population, respectively. Wednesday has the next highest VIF of 2.910 in the sample and 2.77 in the population.

Although VIFs below 10 are generally perceived to be non-problematic, as VIF does not have a strict diagnostic boundary, the eigenvalues and condition numbers of the correlation matrices were calculated. Eigenvalues close to zero indicate multicollinearity, and their eigenvectors consist of coefficients that indicate which variables might be involved (high coefficients are signs of multicollinearity). The lowest eigenvalue in the correlation matrix for the sample is 0.097, and for the population it is 0.078. Their eigenvectors reinforce the results from the VIF test: Leave and Remain are the variables with the highest coefficients (0.507 and -0.569 in the sample, -0.557 and -0.674 in the population).

The condition numbers of the correlation matrices are computed from eigenvalues. They approach infinity if the variables have perfect linear dependence (i.e., a variable approximates the linear combination of other variables), and are reported in the regression tables. The largest condition number is associated with the full models because they contain the most variables: it is 24.91 for the sample and 30.96 for the population. As the diagnostic boundary for condition number is 30, there is a potential multicollinearity problem at the population level—and the eigenvectors of the lowest eigenvalue reveal that Leave and Remain are the main cause (Hair et al., 2010). However, as Leave and Remain are substantively interesting control variables, neither is removed from the model. Moreover, the reason for their collinearity is clear: they are mutually exclusive categories of the same variable. A tweet that is pro-Leave is by default not pro-Remain, and vice versa. The high condition numbers result from a technicality, and as such can be disregarded.

To detect outliers, influence plots were created for the full models. Studentised residuals, which indicate outliers on the dependent variable, are visualised against leverage values (the diagonal of the hat matrix), which indicate tweets that are substantially different with regard to at least one independent variable. The size of a point is weighted by its Cook's distance score, a combination of residuals and leverage. Cook's D scores represent the best measure of influence on overall model fit, and 1 is a common threshold (Hair et al., 2010). The largest Cook's D score in the sample is only 0.034 (twice the value of the second highest); in the population it is 0.001. The sample tweet with this score received 8,796 retweets and 18,607 likes, and will be analysed in more detail in Chapter 6.

Because outliers on the dependent variable are inconsequential for logistic regression, only tweets with the highest leverage were removed (the seven with scores above 0.030 in the

sample, as 0.037 is the maximum value and 0.013 is the 75% percentile). Regressions with all of the variables were computed and compared with the original results (see Appendix G). The removal of the outliers has little impact on the models: significance levels do not change by much, and coefficients only change slightly, if at all. Only in the like model are there noticeable p value changes: small hours and healthcare decrease by one significance level, whereas emotion and question increase by one significance level. These changes are not consequential—none of the significant independent variables become insignificant, and no directions of effect are reversed. Thus, the outliers on the independent variables are retained in both models.

4.2.2 *Control Features*

The control-only models in the leftmost columns of Tables 4.11 and 4.12 serve as baselines against which the influence of context, content, and style can be assessed. All of the control variables are statistically significant predictors for whether a tweet is liked, and almost all are significant predictors for whether a tweet is shared. Unsurprisingly, the user influence variable follower count is positively related to both measures of engagement: for each tenfold increase in follower count, the odds of receiving a retweet increase by 73.8%, and the odds of receiving a like increase by 60.1%, holding all other control variables constant.⁵⁷ The user experience variable account age has, interestingly, a weak negative relationship with both dependent variables: for each yearly increase in account age, the odds of receiving a retweet decrease by 10.7%, and the odds of receiving a like decrease by 4.1%. This is a marginal difference, but is counterintuitive nevertheless. In the context of the EU referendum, however, it must be noted that many accounts were created specifically for campaigning purposes—thus, younger accounts in this dataset are more likely to be more active.

The reference category for the hour of day variable is night, as it is the mode. For retweet count, afternoon is not statistically significant when compared to this category, but tweets sent during the morning and small hours are significantly less likely to get retweets. Tweets sent in the morning, afternoon, or small hours are significantly less likely to get likes. It is no surprise that the largest decrease in odds is for small hours in both models: almost no one is online at this time.

⁵⁷ All percent changes are calculated from the odds ratios shown in the tables.

TABLE 4.11: Odds Ratios from Five Models for Was Retweeted (Sample)

	CONTROL	DISCOVERABILITY	CONTENT	STYLE	FULL MODEL
CONTROL					
Leave	0.714 **	0.769 *	0.618 **	0.713 **	0.637 **
Remain	0.984	0.878	0.798	0.988	0.725
Followers count ($\log_{10}$)	1.738 ***	1.627 ***	1.759 ***	1.744 ***	1.647 ***
Account age	0.893 ***	0.908 ***	0.895 ***	0.894 ***	0.911 ***
Small hours	0.413 ***	0.513 **	0.451 **	0.410 ***	0.547 *
Morning	0.764 **	0.746 **	0.772 **	0.769 **	0.757 **
Afternoon	0.927	0.941	0.932	0.929	0.949
Monday	0.263 ***	0.320 ***	0.269 ***	0.270 ***	0.323 ***
Tuesday	0.233 ***	0.271 ***	0.239 ***	0.240 ***	0.278 ***
Wednesday	0.247 ***	0.297 ***	0.253 ***	0.254 ***	0.306 ***
Thursday	0.254 ***	0.292 ***	0.260 ***	0.261 ***	0.296 ***
Friday	0.297 ***	0.367 ***	0.306 ***	0.304 ***	0.372 ***
Sunday	0.247 ***	0.301 ***	0.255 ***	0.255 ***	0.308 ***
Days before referendum	1.008 **	1.011 ***	1.008 **	1.008 **	1.011 ***
DISCOVERABILITY					
Hashtag count		0.908 ***			0.912 ***
Mention count		1.061			1.043
URL count (binary)		0.655 ***			0.634 ***
Visual media count (binary)		2.095 ***			2.166 ***
Organisation		7.847 ***			7.414 ***
CONTENT					
Identification			1.553 **		1.290
Division			1.225		1.247
Emotion			1.030		1.054
Economy and labour			0.866		1.020
Healthcare			0.518 *		0.629
Immigration			0.870		0.999
Other issue			0.804		0.995
Security			0.931		1.097
Sovereignty			0.694 *		0.746
Resolution			0.766 *		0.842
STYLE					
Message format				0.983	0.885
Question				0.918	0.924
Playfulness				0.882	0.844
Number of Observations	2,976	2,976	2,976	2,976	2,976
McFadden's R^2	0.06	0.11	0.07	0.06	0.12
AIC	3,766.67	3,581.77	3,763.53	3,770.01	3,590.02
Largest Condition Number	18.76	20.14	22.95	18.92	24.91

* $p < .05$, ** $p < .01$, *** $p < .001$

Odds ratios are presented in the table, but are reported as percent changes in the text for ease of interpretation. The omitted reference categories are ambiguous (position), night (hour of day), other (message framing), Saturday (day of week), and none (issue). The discoverability, content, and style models all include control variables.

TABLE 4.12: Odds Ratios from Five Models for Was Liked (Sample)

	CONTROL	DISCOVERABILITY	CONTENT	STYLE	FULL MODEL	
CONTROL	Leave	0.758 *	0.853	0.721 *	0.736 *	0.771
	Remain	1.298 *	1.231	1.149	1.244	1.057
	Followers count ($\log_{10}$)	1.601 ***	1.564 ***	1.630 ***	1.608 ***	1.579 ***
	Account age	0.959 *	0.973	0.958 *	0.959 *	0.976
	Small hours	0.419 ***	0.569 *	0.464 **	0.406 ***	0.590 *
	Morning	0.742 **	0.778 *	0.764 **	0.739 **	0.787 *
	Afternoon	0.659 ***	0.710 ***	0.667 ***	0.659 ***	0.711 ***
	Monday	0.247 ***	0.300 ***	0.252 ***	0.251 ***	0.299 ***
	Tuesday	0.289 ***	0.332 ***	0.292 ***	0.297 ***	0.338 ***
	Wednesday	0.291 ***	0.351 ***	0.292 ***	0.296 ***	0.353 ***
	Thursday	0.332 ***	0.381 ***	0.336 ***	0.338 ***	0.383 ***
	Friday	0.274 ***	0.326 ***	0.278 ***	0.280 ***	0.328 ***
	Sunday	0.263 ***	0.313 ***	0.270 ***	0.267 ***	0.321 ***
	Days before referendum	0.993 **	0.996	0.993 **	0.993 **	0.996
DISCOVERABILITY	Hashtag count		0.904 ***			0.905 ***
	Mention count		0.904			0.888
	URL count (binary)		0.456 ***			0.468 ***
	Visual media count (binary)		1.780 ***			1.797 ***
	Organisation		5.734 ***			5.674 ***
CONTENT	Identification			1.228		1.051
	Division			1.067		1.041
	Emotion			1.191		1.223
	Economy and labour			0.713 **		0.901
	Healthcare			0.409 **		0.514 *
	Immigration			0.754		0.892
	Other issue			0.682 **		0.900
	Security			0.790		0.952
	Sovereignty			0.609 **		0.649 *
	Resolution			0.800		0.923
STYLE	Message format				1.116	1.015
	Question				0.732 **	0.743 *
	Playfulness				1.109	1.024
Number of Observations		2,427	2,427	2,427	2,427	2,427
McFadden's R ²		0.05	0.11	0.06	0.05	0.11
AIC		3,775.05	3,568.15	3,763.58	3,771.27	3,572.39
Largest Condition Number		18.76	20.14	22.95	18.92	24.91

* $p < .05$, ** $p < .01$, *** $p < .001$

Odds ratios are presented in the table, but are reported as percent changes in the text for ease of interpretation. The omitted reference categories are ambiguous (position), night (hour of day), other (message framing), Saturday (day of week), and none (issue). The discoverability, content, and style models all include control variables.

Saturday is the reference category for day of the week, as crosstabulations of both dependent variables with this variable indicate that tweets sent on Saturday have the highest proportion of engagement to no engagement (Appendix F). The regression results corroborate this: tweets sent on other days of the week are significantly less likely to be shared or liked; the decrease in odds ranges from 70.3% to 76.7% for retweets and 66.8% to 75.3% for likes. This illuminates that users are probably most active on Twitter on Saturdays, although this is counterintuitive in light of the interviewee strategies presented in the next chapter. The average retweet count for each Saturday in the dataset was computed to detect whether an exceptional event occurred on one particular Saturday, but no such evidence was found—this day of the week appears to be systematically better than the others.

The number of days before the referendum that the tweet is published is the final time-based control variable, and it is highly statistically significant. Its relation with both measures of engagement is quite strong: for each one-day increase, the odds of being retweeted increase by 0.8%, whereas the odds of being liked decrease by 0.7%. As the entire campaigning period was 68 days, tweets sent on 22 June had a 54.4% greater odds of being shared and a 47.6% lower odds of being liked than tweets sent on 15 April. Because the referendum was more timely as 23 June drew closer, more people were motivated to tweet. Sharing content was more important than liking content, as spreading the message as far and wide as possible was critical.

The reference category for position is ambiguous. Pro-Leave tweets have a lower likelihood of receiving a share or a like than tweets that do not clearly align with a position, whereas pro-Remain tweets have a higher likelihood of receiving a like (Remain is not significant in the retweet model). It must be noted that although Leave is highly significant in the retweet model, the position variables are only weakly significant in the like model. As will be discussed later in this section, in other like models the position variables are no longer significant. Nonetheless, this suggests that Remain-oriented tweets were more likely than Leave-oriented tweets to receive engagement.

4.2.3 *Discoverability Features*

The second columns in Tables 4.11 and 4.12 reveal the impact of discoverability features on retweet and like count, while controlling for user influence, user experience, time, and position. The addition of these features—hashtag count, @mention count, binary versions of URL and visual media count, and organisation—does not change the effect of the control variables in the retweet model by that much. The only discoverability feature that is not highly significant in both models is mention count. All control variables are at the same level of statistical

significance (the only noticeable difference is that Leave drops from the 0.01 level to the 0.05 level). The like model with discoverability features is more divergent from the baseline model: Leave, Remain, account age, and days before referendum are no longer statistically significant; small hours and morning drop from the 0.01 to the 0.05 level. Odds ratios in both models are only slightly different, if at all.

The goodness of fit measures indicate that the models including discoverability features are improvements over the baseline models. The pseudo- R^2 value almost doubles for the retweet model and more than doubles for the like model after discoverability features are included. The decrease of AIC from 3,766.67 to 3,581.77 in the retweet model and from 3,775.05 to 3,568.15 in the like model signifies that the added complexity introduced by discoverability variables is justified by the explanatory power that they provide. Thus, discoverability features in conjunction with the control variables better explain why some tweets receive engagement than the control variables alone.

4.2.4 *Rhetorical Content*

Once a tweet has been seen, its content can actually be read. To what extent does the substance of a tweet impact its success? The third columns in Tables 4.11 and 4.12 reveal the impact of rhetorical content features on retweet and like count. The inclusion of these features—message framing, emotion, issue, and resolution—has a negligible effect on the baseline models. Apart from the fact that Remain in the like model is no longer significant, there are no significance level changes.

The presence of identification rhetoric increases the odds of being retweeted by 55.3%, compared to tweets that do not contain identification or division rhetoric. Division rhetoric, on the other hand, is not significant; nor is the presence of strong emotion (compared with no emotion). Of the resonant issues, only healthcare and sovereignty are significant relative to the reference category (no issues): including them in a tweet actually *decreases* retweet odds by 48.2% and 30.6%, respectively. Resolution is also significant and negatively correlated—it lowers the odds of receiving a retweet by 23.4%. The only statistically significant content variables in the like model are issue categories: economy and labour, healthcare, sovereignty, and other.⁵⁸ Here, too, they are negatively related with engagement—references to these issues decrease the odds of a tweet being liked by 28.7%, 59.1%, 39.1%, and 31.8%, respectively.

These results indicate that positive campaigning is more effective than negative campaigning, which supports hypothesis 3a; they also provide evidence against hypotheses 5

⁵⁸ Immigration and resolution are borderline significant, with p values of 0.0589 and 0.0619, respectively.

and 6, as it appears that including salient issues or calls to action is rhetorically ineffective. All of these conclusions are tentative, however, because the goodness of fit measures indicate that adding content features does *not* improve either baseline model. There are only minor increases in McFadden's R^2 , and the improvement of fit is also marginal (the AIC for the retweet model only decreases by three, and the AIC for the like model decreases by 11). It is clear from this that content features overcomplicate the models, and are less ideal predictors than discoverability features.

4.2.5 *Rhetorical Style*

If *what* is expressed in a tweet does not matter for success, what about the manner in which the message is expressed? This is captured by the three variables that constitute rhetorical style: message format, questions, and playfulness. The fourth columns in Tables 4.11 and 4.12 represent their impact on the baseline models. This impact is even more negligible than that of content features: the only noticeable change is the fact that Remain in the like model is no longer significant. Moreover, no style variable is significant in the retweet model. In the like model, only question is significant: it decreases like odds by 26.8%. This is evidence against hypothesis 8; nothing can be said about hypotheses 7a, 7b, or 9.

The goodness of fit measures confirm that the models with rhetorical style features added to the control variables are poorer fits than the models that include rhetorical content features. The pseudo- R^2 value in both models remains the same; AIC decreases by four in the like model and actually *increases* by three in the retweet model. Thus, style features not only add too much complexity, but also reduce the fit.

4.2.6 *Full Models*

The above subsections have demonstrated that while discoverability features are quite important, content and style features alone do not hold much explanatory power for tweet success, and in fact overburden the baseline models. In the context of all features, however, they might be more effective. The full models are represented by the fifth columns in Tables 4.11 and 4.12.

A comparison of the full retweet model with the baseline model reveals that in the context of all predictors, the explanatory power of the control variables only slightly decreases. Apart from Remain, which is no longer statistically significant, and small hours, which drops from the 0.01 to the 0.05 level, all other control variables remain at the same significance levels. The odds ratios for account age, days before referendum, and days of the week do not change

by much. In the context of all predictors, the odds of pro-Leave tweets being shared are 36.3% lower than the odds of ambiguous tweets being shared; this is a stronger negative relationship than that in the baseline model. Two variables have noticeably weaker relationships in the full model: follower count and small hours. For every tenfold increase in follower count, the odds of receiving a retweet increase by 64.7%; tweets sent between 1 and 6am have a 45.3% lower odds of being retweeted than tweets sent during the night.

What is more interesting is the fact that of the independent variables, only those that relate to discoverability remain statistically significant: hashtag count, whether the tweet has a URL or a media item, and whether the user is an organisation. The relationships illuminated by the discoverability and control model are reflected in the full model; odds ratios are extremely similar. With every additional hashtag that is included, the odds of being retweeted decrease by 8.8%. Including a hyperlink decreases retweet odds by 36.6%, whereas including visual media increases retweet odds by 116.6%. Organisation has by far the strongest positive relation with whether a tweet is retweeted: tweets from organisations have a 641.4% greater odds of being shared. None of the content or style features is significant—not even identification, healthcare, sovereignty, or resolution, which had been significant in the model with content and control variables.

The pseudo- R^2 measure for the full retweet model is twice that of the baseline, but this is only marginally better than the model with discoverability and control variables. Moreover, it must be noted that the pseudo- R^2 measure is extremely low: only 12% of the variance in whether a tweet is retweeted is explained by the full model; it is evident that important explanatory variables are missing. What's more, the added complexity introduced by the content and style variables decreases the fit of the full model: its AIC is 177 units lower than that of the baseline model, but 9 units *higher* than the AIC of the discoverability model. Thus, the discoverability model is superior to the full model for whether a tweet is shared.

The full model for whether a tweet is liked has a greater divergence from the baseline model—the control variables lose explanatory power in the context of all predictors. Leave, Remain, account age, and days before referendum are no longer significant; small hours and morning decrease from the 0.01 to the 0.05 level. The odds ratios remain about the same for followers count, days of the week, morning, and afternoon. The only noticeable change is for small hours: tweets sent between 1 and 6am have a 41.0% lower odds of being retweeted than tweets sent during the night (this is a weaker negative relationship than that in the baseline model).

The four discoverability features that remain significant for the retweet model remain significant for the like model as well: hashtag count, whether a tweet contains a hyperlink or media item, and organisational status. As with the full retweet model, the full like model reinforces the relationships revealed by the discoverability and control model. With each additional hashtag, the odds of being liked decrease by 9.5%. Including a URL decreases like odds by 53.2%, whereas including visual items increases like odds by 79.7%. Tweets from organisations have a 467.4% greater odds of being liked—this is by far the strongest variable in the model.

Two of the variables that had been significant in the content model are no longer significant in the full model: economy and labour and other issue. Healthcare, sovereignty, and question are still significant, although all three drop from the 0.01 to the 0.05 level. They have a moderate negative relationship with whether a tweet is liked, and decrease like odds by 48.6%, 35.1%, and 25.7%, respectively.

The full model for whether a tweet is liked has comparable predictability to the model with discoverability features—they have the same McFadden's R^2 , which is over twice that of the baseline but still extremely low; only 11% of variance in whether a tweet is liked is explained. Once again, the content and style variables introduce too much complexity, and decrease the fit of the full model: its AIC is 203 units lower than that of the baseline, but four units higher than that of the discoverability model. The discoverability model is thus also superior to the full model for whether a tweet is liked.

4.3 Population-Level Models

As tremendous effort was exerted to collect millions of tweets from the Twitter API pertaining to the referendum, larger scale regressions were conducted on the entire population of 741,411 original tweets from which the random sample was drawn. As explained in Subsection 4.1.1, only retweet models can be computed due to missing data issues for like count. The results of these regressions (Table 4.13) are compared with and contrasted against the results of the sample-level regressions (Table 4.11). To what degree are trends in the manually-coded sample reinforced?

Due to low classifier accuracies, many of the content and style variables had to be modified or excluded: issue was dichotomised, enthusiasm, aversion, and anxiety were combined into one binary emotion variable, only the division category of message framing was included, and message format and playfulness were not classified. Resolution is retained, with

the caveat that it represents only a subset of common and unambiguous calls to action. The extreme complexity of tweet text accounts for the poor classifier accuracies. Regardless, the sample-level regressions have revealed that most content and style features are inconsequential, so the more simplified population-level models should represent a better fit to the underlying data.

The results from four population-level logistic regression models for whether a tweet was shared are presented in Table 4.13. Because there is only one style variable, it was collapsed with content into one model (represented by the third column). Given the extremely large number of observations, it is unsurprising that every variable is statistically significant—only the strengths of relationships can be compared. The large number of data points also significantly decreases the standard errors of the estimates. Thus, the population-level models have higher predictive accuracy, with the caveat that they do not fully represent the content and style variables.

The odds ratios of the control variables in the baseline and full population models reflect those in the corresponding sample models, albeit to a slightly weaker degree: they are closer to 1 in the population models for most variables (which indicates lower percentage changes). The odds ratios for account age and days before the referendum are about the same: in the context of all predictors, an increase in one year for the former decreases the odds of being retweeted by 9.2%, and an increase in one day for the latter increases the odds of being retweeted by 0.70%. Leave and follower count have slightly stronger relationships with success at the population level than at the sample level: in the full model, the odds of pro-Leave tweets being retweeted are 42.5% lower than the odds of ambiguous tweets being retweeted; for each tenfold increase in follower count, the odds of receiving a retweet increase by 85.0%. The odds of pro-Remain tweets being shared are 34.7% lower than the odds of ambiguous tweets being shared, but this variable was not significant at the sample level.

The impact of discoverability variables is also extremely similar at the population and sample levels. Organisation remains the strongest variable by a wide margin: in the context of all predictors, tweets from organisations have a 624.1% higher odds of being retweeted. Every hashtag that is added decreases retweet odds by 11.9%, and tweets with hyperlinks have a 30.1% lower odds of being retweeted. Given the larger number of observations, visual media was split into its three constituent categories: photo, video, and GIF. The population-level regression reveals the differences in effectiveness of these types. Video is by far the most resonant visual item: including a video increases retweet odds by 162.6%, in the context of all

predictors. Including an image increases retweet odds by 64.8%, and including an animated GIF only increases odds by 25.9%.

TABLE 4.13: Odds Ratios from Five Models for Was Retweeted (Population)

	CONTROL	DISCOVERABILITY	CONTENT & STYLE	FULL MODEL
CONTROL	Leave	0.457	0.480	0.604
	Remain	0.604	0.559	0.739
	Followers count ($\log_{10}$)	1.858	1.834	1.896
	Account age	0.897	0.908	0.896
	Small hours	0.490	0.632	0.514
	Morning	0.770	0.821	0.791
	Afternoon	0.757	0.799	0.776
	Monday	0.370	0.449	0.399
	Tuesday	0.314	0.377	0.339
	Wednesday	0.318	0.383	0.342
	Thursday	0.348	0.423	0.375
	Friday	0.386	0.470	0.415
	Sunday	0.393	0.480	0.425
	Days before referendum	1.005	1.007	1.006
	Hashtag count		0.877	0.881
DISCOVERABILITY	Mention count		1.026	1.027
	URL count (binary)		0.696	0.699
	Photo		1.661	1.648
	Video		2.625	2.626
	GIF		1.280	1.259
	Organisation		7.308	7.241
	Division		0.875	0.936
CONTENT & STYLE	Emotion		0.799	0.813
	Issue		0.883	1.004
	Resolution		0.669	0.757
	Question		0.805	0.835
	Number of Observations	741,411	741,411	741,411
	McFadden's R^2	0.06	0.09	0.06
	AIC	925,283.5	889,922.77	921,565.54
	Largest Condition Number	19.79	21.89	28.68

Because everything in the model is highly statistically significant ($p < .0001$), significance stars are not included.

The omitted reference categories are ambiguous (position), night (hour of day), other (message framing), and Saturday (day of week). The discoverability, content, and style models all include control variables.

The odds ratios of the content and style variables are all closer to 1, which indicates that even if they are highly significant, the sizes of their effects on retweet odds are not as high as those of discoverability features. The binary version of issue has the smallest effect size: including a substantive theme only increases retweet odds by 0.40% in the full model. The

other four content and style features decrease the odds of being shared: division by 6.4% (compared with other), emotion by 18.7%, resolution by 24.3%, and question by 16.5%. These effect sizes are moderate, and indicate that the impact of content on tweet resonance is not negligible. Division had been the closest to being significant in the sample-level full model, with a p value of 0.0877.

The goodness of fit measures indicate that the best model is discoverability and control (Table 4.13, second column). Its pseudo- R^2 value is 1.5 times that of the baseline model, whereas the content and style model has the same pseudo- R^2 as the baseline. The full model has the same pseudo- R^2 as the discoverability model, but at a cost of added complexity: its AIC is 37,000 units less than that of the baseline, but only 1,639 units less than that of the discoverability model (given the larger number of observations, a difference of 1,639 at the population level is marginal). Just as significantly, the pseudo- R^2 measure is even lower than that in either sample-level full model: only 9% of the variance in whether a tweet is retweeted is explained by the independent variables. This underscores the fact that important predictor variables are missing; further remarks shall be made upon this matter in the next section.

4.4 Discussion

The statistical results presented in this chapter provide evidence that on social media, content and style *do* matter, but only to a limited degree. What matters more is arresting attention; the regressions at both sample and population levels clearly indicate that discoverability features explain far more than rhetorical features, and have larger effects on tweet success.

All things considered, these results provide strong evidence in support of hypotheses 1c and 2:

HYPOTHESIS 1C If a tweet contains visual media, it will be more successful.

HYPOTHESIS 2 If a tweet is from an organisation, it will be more successful.

Tweets that include visual media and tweets from organisations are overwhelmingly more successful; this is illustrated at both sample and population levels. The population-level regressions reveal that there are also qualitative differences among different types of visual media. Videos are by far the most effective, followed by images and animated GIFs.

The sample and population-level results provide evidence *against* hypotheses 1a and 1d:

HYPOTHESIS 1A If hashtags are included in a tweet, it will be more successful.

HYPOTHESIS 1D If a tweet contains hyperlinks, it will be more successful.

Tweets that include multiple hashtags are less successful than tweets that contain at most one; reasons for this are illuminated in the next chapter, which focuses on interview data with individuals and campaigners who tweeted actively during the referendum. Although hyperlinks can be visually arresting due to image previews, they are qualitatively different from media items because they do not represent native media content—the previews generated are automatically extracted from the URLs. It must also be noted that sometimes previews are not generated at all.

The like models provide evidence *against* hypotheses 5 and 8; the population-level retweet models provide weak evidence for hypothesis 5 and also do not support hypothesis 8:

HYPOTHESIS 5 If a tweet contains relevant issues, it will be more successful.

HYPOTHESIS 8 If a tweet contains questions, it will be more successful.

The presence of certain relevant issues (namely healthcare and sovereignty) or questions *lowers* the odds of tweet resonance. On the whole, however, including a substantive issue in a tweet slightly increases retweet odds. Perhaps specific issues are too controversial, and perhaps questions (like @mentions) encourage replies as opposed to likes or shares.

Less can be concluded about content and style variables, almost all of which remained insignificant at the sample level. Partial evidence is provided *against* hypotheses 3b, 4a, 4b, 4c, and 6 by the population-level retweet models:

HYPOTHESIS 3B If a tweet contains division rhetoric, it will be more successful.

HYPOTHESIS 4A If a tweet contains enthusiasm, it will be more successful.

HYPOTHESIS 4B If a tweet contains anxiety, it will be more successful.

HYPOTHESIS 4C If a tweet contains aversion, it will be more successful.

HYPOTHESIS 6 If a tweet contains resolution, it will be more successful.

Including divisive rhetoric, strong emotions, or resolution actually decreases the odds of a tweet being shared. While their effect sizes are not at all negligible, it is difficult to make firm conclusions about these features, as they are not significant in the sample-level regressions. Given the large *N* of the population dataset, every variable is extremely statistically significant. The qualitative analysis in the next chapter will help illuminate why these content features might backfire.

Although hypothesis 1b is partially supported by the population-level regressions, no conclusions can be drawn from the sample-level models. No comments can be made about hypotheses 3a, 7a, 7b, and 9 from either sample or population-level analyses:

<i>HYPOTHESIS 1B</i>	If user mentions are included in a tweet, it will be more successful.
<i>HYPOTHESIS 3A</i>	If a tweet contains identification rhetoric, it will be more successful.
<i>HYPOTHESIS 7A</i>	If a tweet contains a list, it will be more successful.
<i>HYPOTHESIS 7B</i>	If a tweet contains a contrast, it will be more successful.
<i>HYPOTHESIS 9</i>	If a tweet is playful, it will be more successful.

Interestingly, @mentions are not statistically significant in any of the sample-level models, even though they technically bring tweets to the attention of specified users. It could be that they trigger different forms of engagement, such as replies. The inconclusiveness around the effects of identification, message format variables, and playfulness might stem from the fact that there are diverse manifestations of these strategies—some of which work, others do not.

Overall, it is clear that discoverability features are more important at a *global* level. As the next chapter will reveal, there are certain *local* contexts in which rhetorical content and style features have an impact. The discoverability features, however, only capture a small proportion of the variation in tweet success. This is reflected by the fact that the pseudo- R^2 values are very low: the highest in the sample-level retweet and like models and the population-level retweet models are 0.12, 0.11, and 0.09, respectively. It is clear that the models are missing significant features that help to explain tweet engagement. This is in and of itself an interesting finding: the low explanatory power of the models is strong evidence that the theoretical framework of rhetoric that was tested is incomplete. This framework was derived from studies of traditional media; to better understand resonance on social media, new theory is unquestionably required. The results from the qualitative analysis presented in the next chapter illuminate several features that should be added to the model. These features, however, raise a measurement problem—they are quite high-level and require additional context that cannot be obtained through the close inspection of tweets alone. Indeed, measurement is a challenge for the content and style features that were included: Twitter data is noisy, and its analysis is limited by the brevity of tweets, as well as the wide range of motivations behind the inclusion of @mentions, hashtags, links, and visuals.

All in all, the results from the quantitative analysis in this chapter suggest that for ordinary users to be retweeted or liked, the optimal strategy is to include a video or image in the message. Organisational users, however, have an extreme competitive advantage—their tweets are far more likely to attract engagement regardless of visual content. This variable has the highest odds ratio in all of the models by an enormous margin, and unlike the other content and style features, it cannot be modified by the user. On a broader level, the regressions indicate that new media reinforces the power structures in old media and offline; politicians, journalists, and leaders of organisations are as dominant on Twitter as they are in newspaper headlines and in the ‘real world’, despite the vast literature on the democratisation potential of social media (Loader & Mercea, 2012). The models for tweet resonance are far simpler than expected: people engage with influential figures and with videos and images. As noted above, however, the models are missing significant features that could not be measured.

Despite these findings, all is not lost for non-organisations. A unique minority of users who do not have organisational support manage to become extremely influential on social media. They are the exceptional cases, the outliers. New theory is required to understand why they are so successful, and the following two chapters assess empirical evidence from tweets and interviews that helps to build this theory.

5. Strategies for Tweet Resonance

‘The stuff that we do, day in and day out, even if it’s not tangible—I think sometimes we get too obsessed with the numbers, which is why these qualitative interviews are so helpful [to] get context and understanding.’

—Sean, Director of If Not Now Digital

‘Some people have a talent and flair at social media, and some people just don’t. It’s got an analytical side to it, but it’s also got a very EQ side to it—which is understanding the mood of the audience there, both that are on your side and that are against you. Understanding what arguments fly and will go down well. You’re doing that on a very rapid basis. You’re constantly message testing but in real life.’

—Mike, Programme Director of Scientists for EU and Heathier In

Sean’s meta-commentary on the interview process perfectly encapsulates the motivation for this chapter: to obtain additional context for fuller understanding of what makes tweets resonate. This intangible understanding is arguably more difficult to grasp than tangible numbers, and it is just as important. Mike’s comments further underscore the unique blend of skills necessary for social media messaging: there is an analytical or quantitative aspect to it, as well as an ‘EQ’ or qualitative aspect. It is fair to say that social media managers are mixed-methods researchers as well as practitioners: they are ‘constantly message testing but in real life’. Both Sean and Mike represent individuals who consciously reflected upon their own tweets and those of others in order to develop more resonant messages. Their strategies are not static but rather dynamic, evolving in response to what works and what doesn’t, and to fit the circumstances.

The features included in the regression models all represent strategies for tweet resonance, and as such constitute the hypotheses of this research. The models reveal that apart from visual media and whether the user represents an organisation, most of the features are not inherently useful or have weak effects. The models fail to consider, however, the fact that *there is a context associated with every feature*: temporal, emotional, thematic, as well as a broader socio-political context that undergirds the entire campaign. They provide a useful summary of overall trends across the population of 741,411 original partisan tweets from the referendum period, but are limited to this bigger picture.

To complement the quantitative results from the previous chapter, this chapter addresses both research questions by transcending the analysis of individual tweets to include the motivations and practises of the users who sent them. It explores the features included in the models and a much richer context, afforded by trace interviews that incorporated specific examples of both successful and unsuccessful tweets from the interviewees. When, why, and how are specific strategies for tweet resonance used? Even if a strategy is not successful at the quantitative level reflected by the models, it might be successful under certain circumstances. Rhetorical content does not matter very much at the global level, but in which local contexts

does it matter? There is immense variation in what constitutes each content feature; certain subcategories of the feature might be more effective than other subcategories. Including a call to action or a specific resonant theme, for example, might not be effective if it's not the right *type* of call to action or the right *timing* for the theme. Images are important, but what kinds of images are particularly successful? Message framing is another illustrative example: the question 'is positive or negative campaigning more effective?' presents a false dichotomy. The question should rather be 'under what circumstances and from whom is positive or negative campaigning more effective?' This form of contextual questioning applies to all of the content features, and is not answerable from a numerical analysis alone. Recognising that organisations are on average significantly more successful than individuals, this chapter also addresses how individual strategies differ from organisational ones.

Moreover, the features included in the regression models are limited, as they are focussed on what is directly measurable from tweet metadata and the inspection of user profiles. There is much more to content and context than what these features can capture: some aspects are not countable or even meant to be counted. It is no wonder, then, that the R^2 values in the models are all so low: much is not explained by the numbers alone; the 'EQ side' that Mike highlights is needed to better understand what makes tweets work.

Thus, in assessing the strategies employed by individual and organisational Twitter users to increase the resonance of their messages, this chapter will contextually embed the features included in the regression models, as well as introduce new features that cannot be counted. It complements the statistical and computational analyses presented in the last chapter with data from interviews with 22 users who tweeted actively during the referendum campaigning period. These interviewees and their conceptions of success will be introduced in the next section, and the following three sections (5.2-5.4) explore the context, content, and style features mentioned in the interviews, many of which overlap with the features included in the models. Each feature is part of or constitutes a strategy. It should be noted that many interviewees discussed an aggregation of features, as opposed to features in isolation. These combinations represent contextual depth, and case study tweets will be presented that best illustrate not only specific features, but also particular combinations: both effective and ineffective examples.⁵⁹ Conclusory remarks will be made in Section 5.5.

⁵⁹ *Note:* the screenshots of case study tweets were taken between 29 September and 4 October 2018, so retweet and like counts may be slightly different than they were at the time the Twitter API was queried for them (just before the interviews in January, February, and March 2018). These marginal differences do not affect the analysis in this chapter. Also, some usernames have changed: Healthier In has become NHS Against Brexit; Stronger In has become People's Vote HQ.

5.1 Interviewees

The 22 interviewees all produced tweets that appeared in the dataset, from private accounts and/or from organisational accounts (a summary is provided in Table 3.4). As such, although they are not a representative random sample of users who tweeted about the EU referendum, they represent an important dimension of discoverability that is included in the models: whether the user is an individual or an organisation. More shall be said on this dimension in Section 5.2. At this point, two caveats must be noted regarding the interviewees. There is a large overlap between individual and organisational accounts, as many interviewees had both. Some had different strategies for these accounts, others did not—indeed, using the accounts in tandem is a strategy in and of itself. Finally, not all representatives of organisations wrote the tweets that they discussed: Will Straw, Matthew Elliott, Leni, Logan, and David commented on tweets published by other people. Their commentary will still be included in this chapter with that caveat, as they are deeply familiar with the organisations and their communication strategies. Will and Matthew’s comments are particularly important, as they were the directors of the official campaigns.

5.1.1 Interviewee Backgrounds and Motives

Almost all of the interviewees have a background in journalism, politics, and/or online marketing. This is no surprise given that these subgroups are particularly active on Twitter—indeed, there is such a critical mass that it is a professional necessity for politicians, digital marketers, and journalists to be on the platform. According to Sam, a journalist at InFacts, ‘The politicians are there, the journalists are there... It’s how we talk to each other—it’s LinkedIn for political nerds’.

Content producers have a self-marketing incentive to be on Twitter. Denis MacShane, Nick T., Sam, Robert, and Katie all use the platform to promote their articles from their personal accounts. Robert observed that he can reach more people with his content through social media than on the publication platform itself: ‘If you write for the *Huffington Post* or for openDemocracy—how many people actually go to that site and read it?’ As freelance journalists, he, Sam, and Katie also use Twitter to obtain more commissioned work. Sam emphasised how important it is for being discovered: ‘I would quite like people to pay me to write articles at some point, and [tweeting] is quite a good way of being on their radar’. Robert has been approached by international publications on Twitter to both write articles as well as be interviewed for them. Indeed, Twitter can be more effective than traditional application techniques for getting a job in journalism (Madrigal, 2013).

The fast, real-time nature of Twitter makes it a perfect platform for sharing event updates or breaking news. Alberto, the Europe Editor of BuzzFeed, and Ruthi, the Social Media Manager of the Oxfordshire Green Party (represented on social media as Oxfordshire Greens), began tweeting because of events that they attended: South by Southwest and the 2009 United Nations Framework Convention on Climate Change Conference of the Parties. Elise, who co-managed the Oxfordshire Greens account, started tweeting because of the Arab Spring, one of the most significant political events that was reported—and shaped—by Twitter. She used the platform to become an expert on the situation: ‘I’d get a news story and I knew about it ten hours earlier because of Twitter. There were young, independent journalists—almost bloggers really—based in places like Tahrir Square who were tweeting’. Robert was actually one of these young, independent journalists in 2013. He lived next to Tahrir Square, where the military coup took place:

One thing I find annoying is that on TV you get a certain idea of what it’s like. At some point the BBC had a video with a headline ‘Cairo is burning’, and they showed quite a few houses on fire. Cairo is a huge city, so if a few houses burn that is bad, but that doesn’t [mean] that Cairo is burning... In that sense, I had a certain urge to share my views.

Robert’s compulsion to use Twitter was not just to fill in the gaps in mainstream media coverage, but also to *counter its narrative* by presenting a more accurate version of what was happening.

The pre-interview survey results reflect the fact that regardless of their professional background, all of the interviewees are quite politically active (Table 5.1). Nineteen were involved in public referendum-related activities during the campaigning period, the most common of which were attending campaign group meetings, distributing leaflets, and attending marches. The 16 interviewees who were eligible to vote in the referendum did. Nineteen voted in the last national election of their home country (the other three were ineligible). The interviewees also exhibit a high degree of party allegiance: eighteen of them either prefer to consistently vote for the same party or are members of a political party. This intense political and referendum-related activity illuminates the interviewees’ motivations to strategically use Twitter during the campaign. They had a social or political agenda; they were not tweeting carelessly or as neutral suppliers of information.

Apart from being politically active, all 22 interviewees are digitally active as well—during the referendum period, they not only read online content pertaining to Brexit from both mainstream and new media sources, but also published original comments or replies on Internet platforms (Table 5.1). Twitter was not the only platform they used: Facebook, news websites,

and blogs were also quite popular. This activity reflects a certain degree of experience with online content as well as political motivation.

TABLE 5.1: Summary of Pre-Interview Survey Responses

POLITICAL ACTIVITY	YES	NO
Prefer voting for a political party	18 (82%)	4 (18%)
Member of a political party	15 (68%)	7 (32%)
Participated in public referendum-related activities	19 (86%)	3 (14%)
<i>Attended a campaign group meeting</i>	16 (73%)	6 (27%)
<i>Distributed print material (e.g., leaflets)</i>	14 (64%)	8 (36%)
<i>Attended a march, demonstration, or rally</i>	10 (45%)	12 (55%)
Voted in the EU referendum	16 (73%)	6 (27%)
Voted in the last national election	19 (86%)	3 (14%)
Voted in the last local election	18 (82%)	4 (18%)
DIGITAL ACTIVITY		
News and information sources for the referendum	N/A	N/A
<i>Websites of print publications</i>	21 (95%)	1 (5%)
<i>Websites of online-only publications</i>	22 (100%)	0 (0%)
<i>Websites of campaign groups</i>	17 (77%)	5 (23%)
<i>Websites of political parties</i>	13 (59%)	9 (41%)
<i>Blogs</i>	17 (77%)	5 (23%)
<i>Twitter</i>	20 (91%)	2 (9%)
<i>Facebook</i>	16 (73%)	6 (27%)
<i>YouTube</i>	13 (59%)	9 (41%)
Made a political comment about the referendum online	22 (100%)	0 (0%)
<i>Twitter</i>	21 (95%)	1 (5%)
<i>Facebook</i>	19 (86%)	3 (14%)
<i>News websites</i>	16 (73%)	6 (27%)
<i>Blogs</i>	12 (55%)	10 (45%)

5.1.2 Conceptions of Success and Intended Audiences

Before strategies for tweet success can be discussed, definitions of success from the interviewees' perspectives need to be established. All interviewees were asked the same question prior to their tweets being introduced: 'What does success on Twitter mean to you?' (Appendix D). Although there was general consensus that retweets and likes were important in a campaigning context (the former more so than the latter because it spreads the message), there was confusion between success as tweet engagement and success as persuading someone to vote for Leave or Remain. Upon hearing the introduction to this research project, some interviewees immediately assumed that success was being defined as persuading someone to vote. This relates to Schudson's (1989) concept of *influence* discussed in Chapter 2. Before the

interview even began, Will, the Executive Director of Stronger In, preemptively declared: ‘What’s really important for you to know is that we didn’t use Twitter for persuasion’.⁶⁰ Although this definition of success is not the focus of this thesis, potential ‘real life’ implications of Twitter will be reflected upon in Chapter 7.

A distinction, then, must be made between persuading people to vote in the referendum and persuading people to share or like a tweet. As Will noted, in campaign parlance the word *persuasion* evokes persuading ‘the body of the electorate who have not made up their mind’. There is also a distinction between persuading someone to vote (i.e., take real life action) and persuading someone to change their view. Logan, the International Officer of the Young European Movement (YEM) and the Director of Sheffield Students Stronger In, emphatically upheld this distinction: ‘Persuasion for me is always: someone has one view and goes to another; you’ve put them on your side now. Maybe they won’t vote your way, but they believe your side is right... The whole “get out the vote” operation is very different, and is a whole huge campaigning issue in and of itself’. Irrespective of whether this second distinction is made, operationalising the persuasion of the electorate is incredibly difficult, if not impossible.

There is, however, a numerical correlation between the two definitions of persuasion, which aligns with the theoretical correlation between *influence* and *diffusion* (Chapter 2). According to Logan, if many people share a tweet, this demonstrates that the tweet might be ‘the kind of thing which convinces people to change their mind’. Conversely, ‘if one person retweets something they’ve probably not been convinced by it’. Logan was quick to append an addendum to this: because so many people retweet content they already agree with, ‘there are certain tweets maybe which are getting a lot of retweets but they’re convincing no one’. Persuading someone to change their mind is thus a higher-order measure of success, that might in some circumstances be indicated by tweets with high levels of engagement.

When assessing definitions of success, the *intended audience* is critical. The interviewees can be divided into roughly three groups with regard to Twitter audiences and motives: some targeted the influencers, others focussed on individuals who already supported their side, and others focussed on the undecided middle. The more established organisations aimed their messages at Westminster politicians and journalists: both official campaigns used their media accounts to do this, and used war metaphors to describe the process. Matthew, the CEO of Vote Leave, commented that Twitter was used to ‘fight battles with people’. In Will’s words,

⁶⁰ It must be noted that multiple interviewees used Facebook for persuading people to vote, because it reaches a bigger audience. Will forcefully emphasised that Twitter is *not* designed for this purpose: ‘If you meet somebody from a political point of view who tells you that they are using Twitter for persuasion, they either don’t understand Twitter as a tool or they’re misleading you’. In light of the diversity of Twitter motives and practises detailed in this chapter, the strength of his conviction is extremely interesting.

within the SW1 network, ‘you want to win the battle of resonance and message’. British Influence, a pro-Remain think tank with ties to Stronger In, targeted politicians as well. David, their Digital Communications Consultant, noted that the objective was to ‘give them ammunition’ for their speaking engagements. It is interesting that he also used a war metaphor. This metaphor emphasises the highly rhetorical nature of the referendum—there are two sides competing for attention and engagement, so rhetorical strategies are extremely helpful for increasing the chances of tweet success (they are analogous to weapons). Success for an audience of influencers is less defined by large numbers of retweets and likes, as these audiences are quite narrow—what is most important is that the tweet is seen.

We Are Europe represents a campaign that focussed on activating latent support. As Sean, the Director of If Not Now Digital,⁶¹ explained: ‘This campaign was to get people that already wanted to vote [Remain] to vote. It was really a “get out the vote” operation. It wasn’t a persuasion operation’. The campaign, as a matter of fact, was inspired by the concern that young people in particular were getting ‘bored to death’—‘the main risk was inertia’. The word *activating* is deliberate: the idea was that those who responded positively to We Are Europe’s tweets would, in turn, serve as activists to convince others to vote. As Harriet, the Campaign Director, affirmed, ‘We were hoping that they were going to lead the people that didn’t know either way into the debate’. They focused not just on the final vote, but also on getting registered to vote and applying for a postal vote. Speaking on behalf of InFacts, Sam also discussed the ripple effect of a tweet, using the same metaphor as David: ‘If nothing else, you can give people ammunition for the argument they’ll have at the pub. For every person that reads your article, they’re going to spread it to a few other people’. InFacts actually had a strategy (and hence audience) shift during the referendum campaigning period. They began by focusing on the influencers, but given the perceived ineffectiveness of the main Remain campaign, decided to expand their messaging to the wider public. For this broader audience, retweets and likes are more important; the active purpose, in Sean’s words, was to ‘take over Twitter’ by ‘flood[ing] it with tweets’.

One local campaign that was interviewed had mixed objectives because of its more ad hoc organisation: Oxfordshire Greens. Ruthi, their Social Media Manager, claimed that they were about *reminding* as opposed to *convincing*: ‘Either reminding people that we’re there and you should follow us, or reminding people “This is important and you should go and vote”’. Interestingly, Ruthi’s colleague Elise claimed that they were *not* preaching to the converted: ‘My hope was that people who were toying with Leave—because that’s what we were about; we

⁶¹ We Are Europe hired this agency to manage their digital strategy.

were about the people who were undecided actually more than the firm voters—would think *oh* and hopefully click on the article’. This aligns more with the classic campaign definition of persuasion. Paul, the Web Administrator of UKtoStay.EU, argued that ‘in a referendum it’s the undecided middle that really matters’—speaking with committed Leavers would have been a waste of time and resources. In the case of the EU referendum, the undecided middle was larger than usual. For this audience, retweets and likes are also quite important. Leni, the Social Media Consultant for the Green Party, noted that retweets allow campaigns to break out of their echo chambers by reaching people that are not following them. As likes are more closely connected to endorsement, they are especially significant from those who are undecided.

Nonetheless, even in the aggregate retweets are less important than persuading people to change their views for some campaigners. Logan captured the tension between persuasion and engagement with a pointed question: ‘Would you rather change one person’s mind or get 1,000 retweets?’ In his opinion, the former is clearly preferable. Brendan, the General Secretary of Labour Leave, agreed with this perspective in his justification of one retweet being successful: ‘You don’t know who that person is and you don’t know what audience that person has. If one person reads one message online that influences them to vote to leave, that was positive’. In his comparison of political and sales campaigns, Sean provided an alternate perspective:

We do sales campaigns, and one like or share on that would be a success because you’re trying to get a sale. But when you’re doing this, you need it to get beyond a certain amount before it starts to feel like it’s really had an impact. You can’t really guarantee conversions or anything; it’s all about sentiment and getting into the consciousness.

Persuasion is more important in a sales context because a sale can be more closely connected with a tweet; the connection between convincing someone to vote and tweeting is far more tenuous. Indeed, no examples of mind-changing tweets emerged during any of the interviews. Tweet popularity is, however, important in and of itself for some campaigners.

5.2 Organisational Authority

Now that the interviewees have been introduced, their strategies aimed at achieving tweet success can be examined. It is important to recognise that organisations and individuals have different strategies. As the sample-level regressions reveal, organisations have a significant advantage over individuals: the odds of their tweets being shared are 684.7% greater and the odds of their tweets being liked are 473.4% greater. User identity is, in fact, the feature with the largest effect on tweet success in the models. This aligns with the theory cited in Chapter 2—organisations have on average more financial and human resources than private individuals, as

well as more salience; all of this increases their discoverability and authority. Sam, a journalist at InFacts, spoke about tweeting from both an organisational account and a personal account. He highlighted the difference clearly: ‘If I write the article for InFacts, I know it will get a certain number of views. If I wrote the same article on my own personal blog, I know I wouldn’t get any engagement at all’. This same logic applies to tweets that share the article. This section will illuminate the reasons for the significant advantage of organisations, and provide qualitative evidence in support of hypothesis 2: *If a tweet is from an organisation, it will be more successful.*

5.2.1 *Time on Twitter*

Regardless of prominence, the key difference between organisations and individuals is that organisations generally invest far more time on Twitter. This is an advantage in and of itself, as it allows for more message testing and thus understanding of which strategies are most effective for retweets and likes. In fact, many individuals do not have tweet strategies; organisations are far more likely to care about tweet engagement. Nick T. and Robert are good examples of individuals without conscious strategies: they just want to communicate important information. As Nick T. put it, ‘I genuinely never think “Everybody will like this” or “Everybody will read this”. I just think “Do I think this is a good idea? Is this something that needs to be said?” ’ Like Denis, he hardly spends any time and effort on tweeting. Brendan, Paul, and Ruthi manage both organisational and private accounts, and focus far more on the former. As Brendan explained, ‘The emphasis was by far on Labour Leave. All the big messages went out on Labour Leave’.

Martha Lane Fox, a Twitter board member, commented directly on the time investment difference for individual and campaign accounts:

I just haven’t got the brain space or the time to spend ages compiling tweets. If I was using it for campaigning more effectively, I probably would—if it was my job. But this is much more just me thinking about things that I like or that made me laugh or that I’m interested in or are important.

Martha’s comment underlines the difference between using Twitter as a social platform, which is more characteristic of individual accounts, and using Twitter as a broadcasting platform, which is more characteristic of organisational accounts. It also highlights the professionalization of tweeting—many organisations hire a full-time social media officer or an external agency, given the importance of digital communications. Being paid clearly increases the incentives for spending more time on Twitter. Sean, the director of the digital agency that produced We Are Europe’s tweets, provided an estimate of their time investment based upon logs of hours worked:

A general sense of it would be probably 25-40 hours a week. It definitely takes up the most time out of all the things. That does not include reporting and planning which inevitably is for Twitter; this is just spending time thinking about directly impacting Twitter: content creation, time on the channel, influence building...

It is unfeasible for most individuals to spend this many hours on Twitter, unless they have a financial or professional incentive to do so.

Other organisations, such as Labour Leave, were logged into the platform at all times. As Brendan explained, ‘I would say 24 hours a day: it was being either monitored or used. I thought it was so important to be in a position where we knew exactly what was happening and when and who was saying what and when’. Brendan’s quote illuminates that one of the reasons it is so essential to be ‘always on’ is research—monitoring the conversation is as important as producing content. Graham, the Social Media Officer of the Bruges Group, Leni, the Social Media Consultant for the Green Party, and Paul also noted that they were constantly logged into the platform, so that they could keep up to date and respond when necessary, as quickly as possible. This connects to *timing*, which will be discussed in the next subsection. For Leni, monitoring is more time-consuming than scheduling content: she spends 90% of her time on Twitter reviewing ‘the reactions we get to our tweets and the general chatter about the Green Party’. Monitoring reactions is also a strategy for improving tweet resonance. In his capacity as a journalist at InFacts, Sam assesses reactions to understand what to tweet next: ‘You want to see how your stuff is being engaged with: is it being shared, is it being liked? If it’s not being shared, should I change the draft to the other post that I’ve scheduled for later?’ Paul noted that UKtoStay.EU had a two-part strategy: to understand what worked, and then to share it ‘over and over and over’. This reinforces the connection between having a digital communications strategy and increased tweet engagement.

Matthew drew an interesting distinction between the main account for Vote Leave, which was designed to reach a general audience, and the media account, which was targeted specifically at journalists and politicians. Not much time was invested in the main account, which tweeted less, but the media account relied heavily upon Twitter:

Twitter was massively important to what the media team were doing, both in terms of following what was going on in the debate more generally [and] getting our message out there. It was in a sense their main social media platform. And also their main conversational platform with journalists... There was no limit to how many tweets it would do in a day.

As Matthew explained, having a separate media account is an important strategy because tweeting too frequently might repel followers. Moreover, the accounts had different purposes: to communicate information to the public and to ‘fight battles’ in Westminster. Stronger In also had separate accounts for general communication and PR-type content.

5.2.2 *Orchestrated Success*

Organisations use digital marketing tools to assist with their relentless tweeting and constant monitoring, both of which are especially important in a campaign setting. Of all the interviewees, We Are Europe leveraged scheduling tools to the fullest extent—and were thereby able to orchestrate the success of their tweets to a certain degree. Sean stressed the importance of organisation given the political nature of the campaign, and provided an overview of their main activities:

There's a lot of process under the hood: 'These are the streams and things that are happening; we're doing interactions on these channels; we're doing content creation; we're doing content posting; we're doing advertising'. In a political campaign, which is known for having the highest engagement rates I've ever seen on any channels, it's what we would normally do multiplied by three.

To facilitate all of this organisation, We Are Europe utilised four different digital marketing tools—each for a specific activity. TweetDeck⁶² was used for *interactions*: someone would track the campaign hashtag (#infor) every 30 minutes and engage with accounts that used it. They qualitatively assessed the *sentiment* of the tweets containing the hashtag; a separate application, Audiense,⁶³ was used to quantitatively measure this sentiment. TweetDeck was also employed alongside the Twitter Web Client to publish *live tweets*. *Scheduled tweets* were divided into two categories: evergreen content, which could be recycled, and time-sensitive content. Edgar⁶⁴ was used for the former, Buffer⁶⁵ for the latter. Both of these automated content posters are premised upon the importance of *timing*, which will be discussed in Section 5.4. Audiense was also used for auto-replies to new followers. Its core function, however, was *influencer building and outreach*, which entails developing relationships with proven and potential supporters of the campaign. These users are ranked according to their number of followers, and leveraged to increase the dissemination of important tweets.

Although these tools were focussed on separate tasks, all of the tasks fit together into one overarching strategy. They should not be enacted at the same time; the order is important, and reflects both chronology and value. As Sean explained:

You have to have X amount of evergreen content before you start moving on to more scheduled and Buffer content. You'd have to have that before you start moving onto influencer outreach. They're like layers of a strategy that you build... It's like a pyramid.

⁶² <https://tweetdeck.twitter.com>


⁶³ <https://audiense.com>

⁶⁴ <https://meetedgar.com>

⁶⁵ <https://buffer.com>

The best example of the effectiveness of We Are Europe’s highly orchestrated approach is their ‘Kiss of Death’ tweet (Figure 5.1). This was their most successful tweet as measured by reach and engagement by a wide margin—it had 782 retweets, 1,573 likes, and 580,322 impressions, which is almost eight times that of the next most successful tweet. To put this into perspective, at the time the tweet was sent the campaign had 1,550 followers. Taking all of their original posts from the referendum period into consideration, their average tweet received 744 impressions, three retweets, and two likes.⁶⁶

FIGURE 5.1: We Are Europe’s ‘Kiss of Death’ Tweet



Retrieved 29 September 2018 from Twitter. Screenshot by author.

This tweet combines many features that will be discussed in this chapter, but to fully understand its success, its background should be taken into account. Like the campaign, the tweet content’s origins were spontaneous—Harriet had the idea after seeing the *Fraternal Kiss* mural in Berlin. She mentioned this to her co-founder Bethan over coffee, who immediately liked the idea; they very quickly commissioned a street artist via a friend. The spontaneous creation of the mural is a contrast against the extensive organisation for its publicity through Twitter. Due to the mural’s prominent location (a main road in Bristol), social media and

⁶⁶ Sean kindly supplied expanded tweet metrics for all We Are Europe tweets sent during the campaigning period, from which these numbers are drawn. The averages are more accurate than those calculated from my dataset of their tweets, which is not complete, but practically the same nevertheless (my rounded averages were also three retweets and two likes).

mainstream media had already picked up on it—it was successful before the campaign shared it themselves.⁶⁷ This represents the complete inverse of the usual situation. The timing of the tweet was thus critical; Sean deliberately delayed it until the next morning: ‘It’s the first thing that people see on their feeds. It’ll blow up’. Moreover, after the post began to spread, the campaign’s PR manager could call journalists to obtain further media coverage. This strategy worked, as the mural was featured on all national media outlets in Britain and even went as far as CNN in the United States. Another reason for the delay was that they wanted to leverage the highly anticipated popularity of the tweet to direct people to a register to vote link, which had not yet been set up. The tweet was extremely successful on this measure as well: there were 1,261 URL clicks.

Because the tweet had such high impact potential, Sean spent 30 to 60 minutes thinking only about how to post it: they ‘needed to make sure that the framing was right’. The framing applies to both the text and the photograph, which was chosen from an album. According to Sean, the text is ‘standard good Twitter practise: short and snappy, call to action’. The feature that he immediately emphasised to explain the tweet’s success, however, was something more conceptual: *connecting to something in the real world*. His digital team supported the promotion of the mural, but they could not have originated it. The tweet would not have worked if it had been a graphic: ‘If you put something in a real life context and show that it actually happened somewhere in the country, someone’s a lot more likely to engage with it. This is actually something that is real, instead of just a figure on a chart’. This is an example of added context for the discoverability feature of visual media: some types of images are more resonant than others, and they are more resonant in specific circumstances. More shall be said on visual media in Section 5.3.

5.2.3 *Strategies to Improve Prominence*

To conclude this section, a qualifying factor highlighted by the interviewees must be noted: not all organisations and individuals are created alike. Mike and Paul claim that account size—measured by number of followers—is what confers an advantage for tweet success, as opposed to whether the user is an organisation. This tangible feature is treated as a control variable in the models, and is associated with a less tangible feature: *prominence*. Is the organisation established or grassroots? Individuals can also be subdivided along this dimension: a popular blogger such as Nick T. is more similar perhaps to an established organisation than he is to a

⁶⁷ In fact, the PR manager of the campaign sent an ex-*Daily Mail* paparazzi photographer to cover the mural’s construction.

lesser-known individual. Prominence is more problematic for the quantitative analysis because it is relatively subjective; there is a spectrum, and it is very difficult to draw a distinguishing line. During the referendum period, both UKtoStay.EU and We Are Europe were grassroots accounts run by volunteers—and are thus qualitatively different from Vote Leave and Stronger In. We Are Europe, however, managed to immediately secure financial and celebrity support, and hired a digital agency to manage their social media accounts. It thus had more prominence than UKtoStay.EU, which had no financial support and deliberately did not involve personalities—in Paul’s words, ‘we were trying to have the facts stand on their own and speak for themselves’.

Taking this qualifying factor into consideration, several strategies that can be used to enhance account prominence will be presented in this subsection. These strategies are derived from interviews with grassroots organisations, although they can also be adopted by individuals to a certain extent. In marketing terms, enhancing account prominence is equivalent to *brand building*, and it is especially important for unknown accounts. For these accounts, there is a *trade-off between self-promotion and short-term engagement*: sharing content from more prominent sources might lead to more retweets, but sharing original content helps with brand awareness and long-term growth. This is exactly what UKtoStay.EU chose to do: ‘We prioritise our own content even if it gets less shares than something much more viral’. Their nine fact-checking resources article is an excellent example of this: it reflects the premise of the campaign, but is far less exciting than, for example, exposing a Leave lie using one of these resources (Figure 5.2).

FIGURE 5.2: UKtoStay.EU’s Fact-checking Tweet



Retrieved 29 September 2018 from Twitter. Screenshot by author.

Paul made a connection between visibility and brand awareness, which is difficult to measure without the use of surveys:

On Twitter that translates into have lots of people seeing your logo and name. Maybe they saw your tweet a few days ago but they look you up the next day, and then maybe they become a follower, maybe they engage with a different tweet. Because of the volume and the speed that the conversation moves at, it's not a bad thing to occasionally have tweets that aren't getting actively engaged with. It moves so fast that nine times out of ten it doesn't go anywhere, but if that tenth

one takes off, it more than makes up for the nine that didn't go very far. Or if it really takes off and goes viral, even the 99 tweets before it.

A brand cannot be built overnight, and Paul's reflections highlight the temporal dimension of tweet engagement: strategies that are effective in the short-term for attracting retweets and likes might be counterproductive in the long-term, and vice versa.

Mike emphasised another concept that relates to brand building and authority: *credibility*. A powerful way to establish immediate credibility is by influencing the mainstream media, thereby defining its narrative. This was a strategy that Mike and his Scientists for EU co-founder leveraged early in the campaign:

A real power of social media is that you tell media outlets 'We would like to write an article and then of course we'll be promoting it'. So you write an article for them, you say 'Tell us when it's online', it comes online, and then you promote it, and you can even spend money promoting that article. They see from Guardian central or Independent central: 'Oh that article seems to be doing really well; people like articles about science and Brexit'. That way you create that positive feedback loop of getting articles into media and then making sure that those media articles do well, and then you show to your social media audience 'Look we're not just chatting with you; we're linking in with these politicians who are saying these things for us'.

This positive feedback loop is a double layer of credibility, and thus creates a win-win situation:

As far as the media are concerned, your credibility is that you've got this big social media following that taps into your community. As far as your community are concerned, your credibility as a campaign outlet is that you can drag and drop things onto the BBC or into the Guardian, Independent, or Mail on Sunday.

The first media article that Mike and Rob wrote was a letter to the *Times*, which had famous signatories including Astronomer Royal Lord Rees and Nobel Prize winner Paul Nurse. This article was actually requested by the *Times*, as they noticed the campaign immediately after their Facebook and Twitter accounts were launched (over one weekend they acquired 2,500 Facebook followers, which was shortly followed by 1,500 Twitter followers). As of 10 October 2018, it is still their pinned tweet, and has over 7,000 retweets and 4,400 likes (Figure 5.3). Pinning the tweet was a deliberate strategy, as it not only reflects the campaign's double layer of credibility in mainstream and social media, but also its third layer of credibility within the scientific community, which is represented by the letter itself.

Mike also observed that having publicly popular tweets is important for *credibility*, especially if the account is also popular—so much so that he deletes Scientists for EU and Healthier In tweets with fewer than ten retweets (at the time of the interview in March 2018, these accounts had 91,000 and 30,000 followers, respectively). Thus, once an account reaches a certain level of prominence, a strategy for maintaining this prominence involves removing

unpopular tweets. This reflects the importance of retweets for not only tweet, but also account success.

FIGURE 5.3: Scientists for EU's *Times* Letter Tweet



Retrieved 29 September 2018 from Twitter. Screenshot by author.

5.3 Discoverability

The four tweet attributes that improve discoverability are discussed and contextualised in this section: visual media, hyperlinks, hashtags, and @mentions. The regression models have shown that only visual media are positively related with tweet success—hyperlinks and including more than one hashtag are actually negatively related, and @mentions are not significant at all. This suggests that the benefit of added discoverability that these features confer is outweighed by other factors, which are illuminated by the interview data. Hypotheses 1a, 1b, 1c, and 1d are thus further tested and qualified in the following four subsections.

5.3.1 Visual Media

Multiple interviewees believed that including visual media is an effective strategy for tweet success. Sean made a direct connection between visual media and arresting attention, which conforms with the theory presented in Chapter 3: 'There's a tendency to scroll past things on timelines that don't have a video or image... because there's nothing that draws the attention'. David argued that irrespective of their content, visual media represent a 'disruption in your news feed' and are thus 'more likely to be seen than a bit of text'. Many interviewees maintained

that visuals were inherently engaging, and therefore shareworthy; several with digital marketing backgrounds were convinced that they were favoured by the Twitter algorithm.

Images were the most mentioned visual media category, which makes sense as they were the most commonly included media type (at both the interviewee tweet level and at the population dataset level). It is no coincidence that the two successful tweets discussed in the previous section contain a visual item. Andrew, the Head of Campaigns of the Freedom Association, is convinced that whenever he includes a visual in a tweet, it always receives at least ten engagements. This is actually an underestimate: of his tweets in the population dataset, those that contained images were retweeted a minimum of 27 times and an average of 61.50 times (a marked contrast to 0 and 16.91, respectively, for those without images).

Andrew's most successful personal tweet of a 'street in a very, very poor part of Beverley' on the morning of referendum day has a quintessential example of an effective image (Figure 5.4).⁶⁸ He asserted that the image in and of itself contributed to its success: 'Without a doubt people wouldn't have really read that; they'd have said *blimey!*' The attractiveness of the image is clearly evidenced from the way he told the story of the tweet. It was so strong that he could not *not* tweet a photograph of what he saw:

I was around with a representative from Vote Leave. We were just doing a tour. We'd already looked at one polling station and we were driving past, and I said 'Let's have a walk down here'. And I saw all these signs. And I just couldn't believe it. I seriously couldn't believe how many signs there were. It was nearly every garden. It was getting close. So I took a picture. It was incredible. I had to take a picture.

The tweet was so resonant that its engagement was immediate and sustained for an entire year: 'It just took off. I was still getting likes for that middle of last year [2017]'. Moreover, it also caught the attention of a traditional media outlet: the local newspaper contacted Andrew to ask for permission to use it.

From both a technical and a timing standpoint, images were important during the referendum campaigning period. According to Elise, in 2016 'more and more people were getting smartphones, so images were easier to deal with'. David expressed a similar observation, and contended that Brexit was 'one of the first political events where it was almost if you're not using an image, you won't get seen in the stream'.

⁶⁸ Andrew actually brought up this tweet, as it was not part of the dataset collected for this research (given that it doesn't contain any referendum-related hashtags). He had no difficulty at all in retrieving it, as he remembered it so clearly.

FIGURE 5.4: Andrew's Street in Beverley Tweet



Retrieved 1 October 2018 from Twitter. Screenshot by author.

Returning to the idea of persuasion vs. engagement, Logan contended that although Twitter cannot convince people to change their mind, it can change their *perspective*. As such, it is an important marketing and PR tool. Sheffield Students Stronger In used Twitter to show that the campaign was both well-supported and extremely active, and real-life photographs were integral to this objective (which could lead to, e.g., a media interview). For Logan, the physical and digital strategies were interconnected: 'If we do this event, we can take some pictures and then put them on social media'. This reflects Sean's assertion that the 'Kiss of Death' tweet was effective because it was connected to something in the real world.

A pair of tweets from YEM provide empirical support for the effectiveness of images. Both tweets have a similar structure: they open with a question, contain enthusiasm, are positively framed, link to a Facebook video, were sent near the end of May, and are concerned with niche topics. The first is about LGBTQ rights, includes a graphic, and is the second most successful YEM tweet in the dataset (Figure 5.5). The second is about gender violence, does not include a graphic, and received no engagement (Figure 5.6). Logan declared that 'the image makes a huge difference'—it was in his opinion the determining factor.

FIGURE 5.5: YEM's LGBTQ Rights Tweet



Retrieved 1 October 2018 from Twitter. Screenshot by author.

FIGURE 5.6: YEM's Gender Violence Tweet



Retrieved 1 October 2018 from Twitter. Screenshot by author.

Four Labour Leave tweets promoting the same Indiegogo campaign illuminate the effectiveness of certain types of images over others, as they contain the same text with different images and were all sent on the same day. Brendan maintained that having a person speaking and a nice quote increase the effectiveness of the image, and this technique can be seen in Figures 5.7 and 5.8. Figure 5.9 features stylised text; Figure 5.10 features an image of London. The tweet with Tony Benn (Figure 5.7) was quite successful compared to the average Labour Leave tweet, which received 86 retweets and 64 likes, but the other three were not. What accounts for the decreased engagement with the Kate Hoey graphic (Figure 5.8) is its complexity: 'It's her with a very large quote on a very colourful background and perhaps not necessarily the most straightforward image'. The advantage of simplicity applies equally to text, as will be seen in Section 5.6.

Thus, while for the average tweet having an image helps, images in and of themselves are not inherently appealing. Apart from having a picture of a person and a quote, features that increase an image's appeal are *visual contrast* and using the recommended dimensions—otherwise, according to Paul, the image might not appear in the feed. Visual contrast is created by using different colours: UKtoStay.EU prefers white on red, whereas Andrew highlights important text in bold red capitals. This contrast allows the text to ‘[fly] off the page’.

FIGURE 5.7: Labour Leave Indiegogo Campaign Tweet (Tony Benn)



Retrieved 1 October 2018 from Twitter. Screenshot by author.

FIGURE 5.8: Labour Leave Indiegogo Campaign Tweet (Kate Hoey)



Retrieved 1 October 2018 from Twitter. Screenshot by author.

FIGURE 5.9: Labour Leave Indiegogo Campaign Tweet (Stylised Text)



Retrieved 1 October 2018 from Twitter. Screenshot by author.

FIGURE 5.10: Labour Leave Indiegogo Campaign Tweet (London)



Retrieved 1 October 2018 from Twitter. Screenshot by author.

An effective strategy for making images even more attention-arresting that is not captured by the models is *tagging users*; up to ten users can be tagged in each image. Sean argued that this is better than @mentions ‘because you don’t use up character limit’. Logan tagged Natalie Bennett and Madeleine Kay in many images, even if they had nothing to do with the image, because they are popular and active users who are likely to retweet content they support. Thus, this technique ‘leaves it less up to chance’. For Sean, choosing to tag users is already an indication of tweet quality—he only does so if he thinks the message will be successful. If the message is extremely important, influencer outreach via image tags can guarantee extreme success. The ‘Kiss of Death’ tweet is a prime example of this (Figure 5.1). As Sean explained:

If a big thing comes out like the mural, you would tag ten of your biggest influencers in that mural. And you’re like ‘Right, we want these ten people’, tag them, and then write a tweet and tag another ten people underneath it and get them all to retweet this tweet. Then it should start trending, and then it should take off.

As can be seen at the bottom of the tweet, Caroline Lucas, Jonathon Porritt, and *Bristol24/7* (a local online newspaper) are three influential users tagged in the image. Because Caroline and Jonathan are on the board of the campaign, tagging them was especially effective.

Videos are arguably more attention-arresting than images, as they contain moving images and will autoplay if they are natively embedded. They are also more favoured by Twitter’s algorithm, according to Paul. This is supported by the population-level regressions, which include visual media types as separate variables: including a video increases retweet odds by 162.6%, whereas including an image or GIF increases retweet odds by 64.8% and 25.9%, respectively. Why are videos computationally favoured? Apart from the fact that they arrest attention, they also account for the biggest share of advertising revenue. Traditionally, such advertising is placed on television programmes, but in recent years social media platforms have attempted to ‘build up enough of the infrastructure that they can start almost acting like conversational television channels’, in Paul’s words. Twitter’s livestreaming of political and sports events in partnership with broadcasters is one manifestation of this business model shift. As Paul contended, ‘There’s a real push for them all to become video-friendly, and to start getting as much original content as they can in video content from a larger, long-term growth perspective’. This financial and future-proofing incentive leads to a self-perpetuating effect—because videos are inherently more likely to receive engagement, more videos are created.

The algorithmic favouring of visual media creates a tension between content creation (originality) and content sharing. Videos are extremely engaging, but they take a lot of time and effort to produce. Downloading and reuploading someone else’s video, however, breaches social media etiquette. As Paul explained:

It's the eternal dilemma for someone like me working on Twitter branding: how much content can we lift? It's a bit different taking a link and reposting it with your own wording attached to actually grabbing a video and reposting... If you were purely looking at it from a 'How can I get the most persuasive, visible tweet?', that's how you would do it.

This is another example of the trade-off between short-term and long-term success: including native videos is an extremely effective strategy for increasing tweet resonance, but at the cost of one's reputation in the Twittersphere.

As with images, videos are not inherently appealing in and of themselves, and some are more effective than others. *Brevity* is an essential feature of resonant videos. As Matthew put it, 'They need to be soundbites; you can't put a five-minute video. It's got to be... the maximum point, the most interesting point'. Including subtitles is also important, as many people do not have the volume turned on.

Only one interviewee made a comment about GIFs, and it was negative. Nick T. referenced a prominent journalist who automatically 'blocks anyone who sends him a GIF', and he himself ignores them. In his opinion, 'They're really overused and often as a lazy way of not making an argument'. The fact that GIFs are overused indicates that many are not original—this lack of novelty counteracts the advantage of a visual item. Thus, unlike an image, a GIF might not be worth 140 characters. It is probably no coincidence that they were the least frequent visual category in both my interviewee tweets and the full dataset, and have a far weaker positive relation with success than images or videos in the population-level regressions, as noted above.

5.3.2 *Hyperlinks*

Hyperlinks are closely connected with images in two ways: functionally and technically. Functionally, they allow for information to be provided beyond the character constraint of a tweet. Information dissemination represents one of the main uses of Twitter, and it was especially important during the referendum period. Graham, the Bruges Group's Social Media Officer, believed that tweeting links was an effective strategy for this very reason: 'To get more information and more nuanced arguments out there you put in hyperlinks'. Sean affirmed that including links is 'quite a standard digital marketing strategy'.

Links can also make a user more attractive, even if they are not to self-authored material. They define Harriet's personal account: 'My Twitter style is to link people to interesting stuff that I've read. Some people follow me because they know that I do that'. Paul also shares links to interesting articles on his personal account, to 'boost [his] credibility of being able to talk about these issues'. Martha enjoys sharing external links because it creates *community*: 'You're

not just ranting on and on and on and on, but you've paid attention to other people and other ideas'.

In Elise's opinion, a 50/50 balance should be struck between visuals and links, as they attract different people—images attract people who are busy, whereas links attract people who have the time to discover more. An effective strategy is actually to include both, as both audiences can thus be reached; this strategy was deliberately employed by Oxfordshire Greens, Scientists for EU, Healthier In, and UKtoStay.EU. The image and the URL work in tandem, as the image helps convince people to click on the link.

There is, however, a technical combination of the two—and this combination represents the second way in which hyperlinks are connected with images. As mentioned in Chapter 2, hyperlinks sometimes contain a visual item, which is displayed as a link preview in the tweet. Link previews can come with either a thumbnail-sized picture or a large one. Figure 5.6 is an example of a tweet with no link preview, and Figure 4.8 is an example of a simple link preview with a thumbnail image icon. One of Harriet's personal tweets is an example of a preview with a large image (Figure 5.11).

FIGURE 5.11: Harriet's We Are Europe on BBC News Tweet



Retrieved 1 October 2018 from Twitter. Screenshot by author.

For some interviewees, the link preview was far more important than the link itself. In describing the hierarchy of content favoured by the Twitter algorithm, Paul added that links

with previews are between images and links (all of which are favoured over pure text). UKtoStay.EU modified their website code to enable links to their content to always ‘fully display with a big Twitter-sized image’. InFacts put Twitter Cards at the same level as video. Logan stressed the importance of not only having a link, but also of having the right preview display: ‘If you put up a link and it doesn’t [have a preview], your tweet is basically dead. Sometimes the preview that it gives you is different to how it actually looks a little bit. It’s incredibly annoying. I have sometimes deliberately erased the tweet and redone it, just to avoid that’. Apart from providing visual appeal and taking up more space on the Twitter timeline, link previews are also important for *credibility*, as Mike asserted. Firstly, previews are so common that if they are not shown, ‘it feels a bit weird’. Secondly, the text and image in the preview can reinforce the argument made in the tweet:

If you say ‘Brexit’s going to be a disaster for science’ and then there’s a picture showing science stuff with the title ‘Brexit will be a disaster for science—new information out today shows that...’ It first has the visual draw in, and then it has the multiple confirmations that this hyperlink is good through to a quality article associated with the title that agrees with what’s just been said.

Katie views the preview as ‘a second attempt to persuade people’, and deliberately keeps her tweet factual. Thus, the content of the link preview can constitute an extension of tweet text. This provides further justification for the manual coding strategy described in Chapter 3, which took the full context of the tweet into account—including hyperlinked content. It also helps to explain the regression results, which are counterintuitive in light of the benefits of hyperlinks described above. Many unsuccessful tweets with links might not have appropriate previews, or might lack them altogether. This is impossible to measure from tweet metadata, and could not be included in the models.

In marketing and political contexts, hyperlinks are also important because they are tied to *action-based campaigning*. Sean admitted that his agency includes URLs wherever possible: ‘Usually there’s an action of what you want people to complete’. Logan emphasised that links help people leave Twitter and turn toward ‘doing real things in the real world’ (albeit via other platforms): ‘You want to be directing them towards Facebook pages which they can like, and to Facebook events that they can go to’. Interestingly, he associated links with persuasion, a higher order measure of success: ‘In terms of convincing people, they’re hugely useful’. In terms of engagement, however, they are not useful unless they come with a preview—ironically enough, by taking people off of Twitter, they *decrease* the likelihood of likes and shares. As Alberto explained, ‘In order to see why it’s amazing, you have to click through. And if you want to retweet, you’ve got to go back’. This also helps to explain the negative relationship between hyperlinks and tweet success in the regressions.

Some links are, naturally, more appealing than others. One example of an appealing link is in Harriet's personal tweet above (Figure 5.11). It actually received 27 link clicks, which is over triple the number of retweets and likes.⁶⁹ What accounts for this? One feature was already mentioned: it has an expanded Twitter Card with a 'huge compelling image'. Secondly, it represents a *newsworthy* story—the image is of the 'Kiss of Death' mural, which is the focus of the article. As Alberto asserted, 'I don't think it adds engagement unless the link leads to a huge story'. Newsworthiness is a dimension of *timing*, which will be discussed in Section 5.4. Thirdly, the link is from an established source: the BBC. Brendan observed that institutional links are more successful than links to personal blogs: 'Personal blogs most people don't take seriously because it could be anyone... People prefer that command and authority that comes with recognised brands'. Brendan's comments reflect one of the conclusions from the previous chapter, which is that new media reinforces structures in traditional media. There are, however, exceptions to this—and Nick T. represents one of them.

To conclude this subsection, several technical points will be made. Some interviewees commented on *link type* and *link transparency*. Because of the 140-character limit, link shorteners such as bit.ly were very popular. As Sean noted, bit.ly is recognisable, whereas something like buff.ly (a link shortener from Buffer) is not, so people will not click. Leni argued, however, that bit.ly is ineffective because it is not transparent. A normal link, on the other hand, includes words that indicate what it is about. Sean and Leni's comments suggest that many unsuccessful tweets probably contain shortened URLs.

When describing hashtags and @mentions, Elise implied that they serve the same informational function as links, as they are meant to be clicked. In her opinion, the more of these features that are included, the better: 'So people have a multiple choice of where they go'. Hashtags and @mentions will be discussed in the next two subsections.

5.3.3 *Hashtags*

The quantitative results indicate that including more than one hashtag in a tweet is ineffective (the more that are included, the lower the odds of tweet success). This is contradicted by the strategies of some interviewees. As Paul asserted, 'it was best practise—on Twitter you should always try and put hashtags in, because it's just what you do if you're trying to get something trending'. During the referendum campaign, UKtoStay.EU included #UKtoStay as well as #EU or #Brexit in every tweet. They tried to 'capitalise on some of the other trending Brexit-related hashtags' if 'there were enough characters left'. Graham also believes that hashtags are useful,

⁶⁹ Harriet kindly supplied screenshots of expanded metrics for her personal tweets discussed in the interview.

because they are used to search on Twitter for certain topics; Sean included #EURef, #Brexit, #Bremain, and #EU in We Are Europe tweets because they were ‘being searched a lot’.

Mike claimed that hashtags are not useful because people search for them (in his opinion they don’t), but rather because of their *aesthetics*: ‘They help colour the words differently and make them pop out... “These hashtags are the issues that I care about: #EU, #Brexit, #science”’. Thus, they can be used to visually emphasise keywords. Having too much blue text, however, is not good. Logan pointed to the YEM tweet about gender violence as an example of this: ‘It’s very hard to read, especially with the link being hidden in the middle of a giant cloud of hashtags’ (Figure 5.6). Another factor that needs to be considered is *hashtag number*: several interviewees said that including too many backfires. As Leni declared, ‘I find it looks a bit spammy and I switch off when I see people using a lot of hashtags’. She claimed that two is ideal. This relates to Logan’s point about the ‘giant cloud of hashtags’: too many hashtags increases the ratio of blue text to black text. Sam provided both aesthetic and conceptual justifications for the ineffectiveness of hashtags: ‘They’re clunky, and it looks like you’re looking for attention’. This is an interesting point in light of his earlier statement about using Twitter to attract more commissions for articles. Although attention-seeking is a popular Twitter motive, it is best not to be too overt about it. Leni, Logan, and Sam’s assertions all support the regression models.

To better understand the quantitative results, it is important to recognise that just as media items and hyperlinks are not created alike, hashtags are not created alike. Sometimes specific hashtags are not helpful: Harriet commented that she should have used #women as opposed to #infor and #voting in one of her tweets, because it would have linked the tweet to wider conversations. Many interviewees mentioned that trending hashtags, which are by definition time-sensitive and short-term, are more effective than long-term hashtags—this conflates hashtags with *timeliness*, which will be discussed in Section 5.4. Logan provided several hypothetical examples:

Say Boris Johnson is making a speech and it trends: #BoJoSpeech. It will trend just for a few hours. Or there’s an election going on, and you’ll get the hashtag #GE17. Or there’s a natural disaster happening and people will be like #tsunami.

As Nick T. put it, these hashtags ‘are very of the moment’. Many people search for topical hashtags at the same time, and it is good to appear ‘in that substream’.

Multiple interviewees pointed out that using hashtags during live events was indispensable, irrespective of whether the hashtag was trending. Hashtags play a pivotal role in organising real-time Twitter streams of conversation; the most prominent event category during the referendum campaigning period was televised debates. Sean referred to the practise of using

hashtags in the context of a TV debate as *hashtag seeding*, which allowed We Are Europe to ‘take over the channel’.

Tweets sent during live events constitute a special category that should be treated differently for two reasons. Firstly, there are far more of them, and hence the chances of any tweet being engaged with are lower. This helps to explain the regression results; quite a few of the unsuccessful interviewee tweets contained event hashtags. The flip side is, as Sam contended, that content has a greater impact: ‘Live tweets are more competitive based on your Twitter content... because everyone is tweeting about it at the same time’. Thus, events constitute a more democratic playing field—prominent organisations and individuals do not have as much of an advantage over non-influential accounts.⁷⁰ They represent a chance to break out of the echo chamber, and the Green Party took advantage of this. As Leni explained, ‘There would’ve been a lot of people watching this who wouldn’t normally have been following us. They might have been following the #BBCDebate hashtag’.

Consequently, tweeting actively during the referendum debates using the stipulated hashtags was a strategy in and of itself. UKtoStay.EU and Labour Leave tried to take advantage of live tweeting as much as possible to boost their engagement levels and spread their messages wider. InFacts journalists covered not only the referendum debates, but also the Sunday morning political shows. Brendan and Paul both emphasised that the *timing* of live tweets is critical—the comments have to occur as close to when things are happening as possible. Even being a few minutes late significantly decreases a tweet’s chances of being noticed, as so much content is being produced. Thus, it appears as if live tweets have to be spontaneous.

Spontaneity, however, can be planned. Leni made broad comments about *planned spontaneity* when discussing the Green Party’s general reactive strategy: ‘Even when we are reacting to something, in some scenarios we would actually have some sort of draft content or a graphic that we have ready to go, because we are kind of anticipating it’. UKtoStay.EU planned specifically for referendum debates, as they had built up a ‘library of graphics and information’. As Paul explained, they could then ‘throw out’ relevant items from this library at event hashtags and say: ‘Hey look, here’s what the facts say vs. here’s what the people debating are saying’.

Outside of live events, during which they are always important, Paul made a qualifying observation about hashtag usage related to account size: hashtags are necessary for smaller accounts, but not for larger accounts. UKtoStay.EU’s follower count ballooned from 4,393 the

⁷⁰ One qualification that must be mentioned is that hashtag streams can be viewed in different ways: they can be ranked by popularity or time posted. As Will noted, ‘Twitter’s algorithms will do it if you want it latest, but it’ll also do top, then you’d hope that your ones will get driven to the top’. Top, however, is the default tab—and influential organisations have a natural advantage in this view of the stream.

evening before the referendum to around 13,000 in February 2018. Because the account now has its own momentum, it is less fixated on using hashtags to improve the chances of tweets being discovered. Ironically, this very prominence increases the effectiveness of hashtag use—more popular accounts are more visible in the hashtag stream.

5.3.4 *Mentions*

Nothing conclusive could be said about @mentions from the quantitative analysis, as it is the only discoverability variable in all of the models that is not statistically significant. The interview data indicates why this might be the case: a primary function of @mentions is *supporting others* (e.g., by crediting users when sharing their content), which reflects Twitter’s conversational nature. What is important to note is that this usage of @mentions is viewed more as Twitter etiquette as opposed to a deliberate technique, even though it can attract engagement as a side effect (usually from the person being supported). Martha stressed that it should *not* be viewed as a strategy, but rather as ‘being respectful to people’s ideas and work’—which is really vital and ‘part of what gives Twitter its community and spirit’. It also builds credibility for both parties.

UKtoStay.EU, We Are Europe, and Brendan (from his personal account) sent many thank you tweets to supporters not because they wanted retweets and likes, but rather to keep them engaged in the campaign. Brendan emphasised that this practise was especially important because social media can be ‘faceless’ and ‘voiceless’, but added that there is a limit: ‘Obviously you can’t thank everyone because you’ve got hundreds of people writing to you, but you can pick out one or two’. We Are Europe, however, literally *did* thank everyone—via an automatic responder (Audiense) that tweeted their gratitude to each new follower along with a link to reasons to remain; by referendum day they had almost 3,000. Paul noted that acknowledging, endorsing, and thanking users is critical for community building and long-term reputation. Sean also emphasised the long-term benefit of tweeting directly at users. It is the first step in a hierarchy leading to real-world action: ‘We talk about marketing funnel: reach, awareness, engagement, action. You’re trying to move them towards actually registering to vote, or voting’. This marketing framework clearly situates four metrics of success in the same context—the ultimate metric being persuasion defined in campaigning terms.

Irrespective of their purpose, Mike asserted that @mentions can be helpful if they are *relevant* and *recognisable*. He provided a hypothetical example: ‘“Dear @BBCNews, your coverage today was appalling because you didn’t recognise this”, and people think “yeah BBC News”, retweet’. This example also illustrates the importance of capturing a sentiment that people agree

with. If the message is ‘followed by a junky list of names’, however, @mentions do not work—people do not share messages if they include unfamiliar accounts, as it can be perceived as spamming.

An additional dimension that must be taken into consideration for @mentions is the *prominence of the user* being mentioned. According to Elise, mentioning more famous people is more effective. Her colleague Ruthi’s most resonant personal tweet provides evidence for this, as it supports Molly Scott Cato, a Green Party MEP (Figure 5.12). Cato was the first person to retweet the tweet; Natalie Bennett, the leader of the Green Party from September 2012 to September 2016, also shared the tweet. Ruthi attributed the tweet’s success to this: ‘The fact that she retweeted meant that more people saw it, but then any further retweeting is more about either supporting her or about the particular question’. Indeed, this support was the motive behind tagging Cato in the first place, as opposed to a desire for engagement: ‘It’s more “that’s a good idea that I tagged her” rather than thinking “I’m going to tag her so she retweets”’. It should be noted that knowing the prominent user increases the effectiveness of this technique, which might also explain why Ruthi’s tweets were successful (she is a Green Party councillor).

FIGURE 5.12: Ruthi’s Molly Scott Cato Tweet



Retrieved 1 October 2018 from Twitter. Screenshot by author.

A desire for engagement is the motive for a different technique: using @mentions to *ask for support*, as they bring the tweet to the attention of tagged users (via a Twitter notification). The Bruges Group employed this strategy with users who already supported them: ‘We used to like @The_bat_gotham—we know that he’s got quite a few followers who like our stuff, and he usually retweets it for us’. Conversely, UKtoStay.EU focussed on attracting the attention of experts that they could not normally reach: ‘It’s using the power of Twitter and the fact that people can help by retweeting our @mention to reach people we couldn’t in a normal circumstance have any chance of talking to’. We Are Europe had a highly organised influencer strategy, as indicated in Section 5.2. After compiling and ranking a detailed list of influencers,

they used Audiense to tweet directly at them at volume. As Sean explained: ‘We would go tweet, enter, tweet, enter and we would send people tweets at probably five times the speed we could do it on the actual channel’. They would tweet 50 influencers to share a particular link, and the strategy always works, provided the list is high-quality. From Sean’s perspective, this is ‘a great way to get a tweet started’.⁷¹

5.3.5 *Other Discoverability Strategies*

The advantage to conducting interviews is that they allow strategies to emerge that are not represented by tweet metadata. This subsection addresses a selection of techniques aimed at increasing engagement through elevating the discoverability of a tweet, all of which are enabled by Twitter’s affordances. None could have been detected by observing tweets in and of themselves.

A technique that guarantees engagement is *having multiple accounts*. As noted above, almost all of the interviewees who tweeted from both personal and organisational accounts used them in tandem. The more people involved in the campaign, the greater the likelihood of retweets. Nick H., the Chairman of London4Europe, mentioned his campaign account in his personal tweets, as it was the quickest method: ‘I would tag in @London4Europe as a prompt for them to retweet. I couldn’t send an email each time saying “I’ve sent this tweet”’. As Andrew managed the Better Off Out account, he directly retweeted himself: ‘I just wanted to have maximum reach’.

Several discoverability strategies relate to Twitter-specific features. The first feature, retweeting, is quite well-known—but not in the context of *retweeting one’s own tweets*. This technique can be employed multiple times on the same tweet, provided the retweet is first undone. It is effective not only because it returns the tweet to the top of the timeline, but also because existing engagements are not lost (they are if the tweet is reposted). Mike uses this technique ‘almost religiously’ on tweets that he thinks are successful: ‘I give it two hours and then retweet it. If it’s successful the next day, undo retweet, retweet it. Particularly first thing in the morning when people are going into work’. More shall be said on timing in the next section. Paul actually discovered this strategy when trying to grab a user’s attention with a reply, which ‘gets lost in their feed’. Retweeting, on the other hand, allows the tweet to appear in his public

⁷¹ It should be noted that influencer outreach does not have to be via public @mentions. In fact, there are more effective ways of doing this privately. On Twitter, Direct Messages can be sent to individuals or groups; email is another private alternative (Stronger In used the latter to send model tweets to influencers for important messages). Both of these methods are effective, Mike contended, because they are ‘more personal and more likely to be seen’. An even more personal method is to establish real-world connections, and this was also part of We Are Europe’s outreach strategy: they ‘had a lot of people working offline to get those relationships warm’, as Sean noted.

feed, and his ‘supporters can retweet it further to amplify it’. This is a good example of creative manipulation of platform design.

Another Twitter feature is pinning tweets. This technique was already mentioned in the context of Scientists for EU’s letter to the *Times* tweet (Figure 5.3). Pinning the tweet afforded the campaign a triple layer of credibility. As Mike noted, this set a flavour that provided people with a *raison d’être* to follow the account. Robert pins tweets about his latest articles to promote the tweets themselves: ‘It’s more likely to get more retweets, because it’s there longer, so whenever someone accesses my profile they’re going to see it’. Thus, this strategy can be used for multiple metrics of success.

Finally, not all tweets are organic. *Promoting organically successful tweets*, in fact, is a strategy employed by We Are Europe and Scientists for EU. Sean’s rationale is that because the tweet is already popular, ‘people will be more likely to retweet it again because they’ve seen that it’s popular’—their initial scepticism about the quality of promoted content is overcome. Although Sean called this a viable strategy, promoting content is a far more common technique on Facebook.

5.4 Timing

Timing is an aspect of tweet context that was operationalised as three different control variables in the regression models: hour of day, day of the week, and number of days before the referendum. It is undeniable that these definitions of timing have an impact on tweet success, as audience sizes on Twitter shift systematically with regard to all three of them.

Only the third of these variables represents a broader conceptualisation of timing that transcends tweet metadata. The interviewees expressed this concept in two ways: as *timeliness* or *topicality*, which strongly relates to *newsworthiness*, and as *novelty*, which strongly relates to *originality*. To identify whether a tweet contains either or both of these features, an awareness of current events as well as social, cultural, and political knowledge are necessary. A tweet is timely if it reflects contemporary headlines in the media, or more conceptually provides information at the moment in which it matters; a tweet is novel if it is the first to express something, or provides information that is outside of general knowledge. These features overlap to a certain degree—tweets that are the first to share breaking news, for instance, are examples of both. The number of days before the referendum captures a narrow dimension of timeliness, as the referendum in and of itself was a topical event; tweets sent closer to voting day were thus more timely.

Timeliness is at the core of Twitter and its philosophy. The company's About page directly reflects this: 'Twitter is what's happening in the world and what people are talking about *right now*' (Twitter, 2018; emphasis added). Harriet, the Campaign Director of We Are Europe, eloquently encapsulated the platform's *modus operandi* when commenting on tweet success:

Twitter makes it unbelievably timely, and it's all about the moment that the person sees that tweet. So for me it's got to be saying something that's touching a nerve, or encapsulating a thought, or expressing an opinion, or awakening someone to something that they are feeling or wondering or thinking about in that moment.

Papacharissi and Oliveira (2012) use the term *instantaneity* to 'describe the drama of events unfolding, being recorded and reported online instantly', and claim it is 'the primary appeal of Twitter' (p. 273). Because Twitter is mobile-first and tweets are displayed in chronological order, time-sensitive messages should be sent as close to the moment of discovery as possible. The affordances of the platform thus condition timeliness, and those who stay true to this core value might have a higher likelihood of being resonant. Indeed, Nick T. contends that instantaneity breeds success: 'I will literally fire off a tweet in ten seconds: those are the ones that people engage with. First thought is always the best with Twitter'.

In general, the interviewees emphasised timing over discoverability, content, and style features in their explanations of tweet success. Mike, the Programme Director of Scientists for EU and Healthier In, provided a hypothetical example to illustrate the importance of timing over content:

If you have two tweets which are identical, you can't always say 'This one has done better because they put an extra word in', because this one might have been first and picked up by someone important, or this one might have been at a better time of day. There's big random factors in there in terms of the context.

Mike's example reflects two different definitions of timing—hour of day and novelty. As the former was included in the models, it will be discussed first along with day of the week.

5.4.1 *Hour and Day*

Interviewees were convinced that the time of the tweet made a difference, and their strategies included sending tweets at particular 'peak' hours, when more people were likely to be on Twitter. There were slightly different definitions of what those 'peak' hours are, however. The general consensus is that morning (7-11am), lunchtime (12-2pm), and evening hours are most effective—these represent commuting and break times, when people tend to browse Twitter. This somewhat contradicts the quantitative results, which illuminate that tweets sent between 2.30 and 3.30pm actually have the highest average retweet count, followed by tweets sent between 6.30 and 8.30am and between 6.30 and 8.30pm (Figure 4.14). The afternoon is a better

time for tweeting than most interviewees believe, which suggests that many people check social media at work, and lunchtime and the late morning are not particularly advantageous.

As the ‘Kiss of Death’ tweet demonstrated, the morning is especially important for *newsworthiness*. In Nick T.’s words, ‘People want their news in the morning or in the evening’. An added benefit of tweeting earlier is that it gives content ‘the lifecycle of the whole day’ to be seen. However, Elise contended that evenings are more reliable: ‘There was a clear time period when you were more likely to grab people in the evening. There’s still a debate about what’s the best time to tweet in the morning’. *Evening* can actually be broken into two categories: 5-6.30pm and 8-10pm. In the earlier interval, people are commuting home; in the later interval, people are at home watching news and political programmes on television. During the referendum campaigning period, a number of live debates were scheduled in the evening hours (Elgot, 2016). As noted in Subsection 5.3.3, tweeting during these debates was an important strategy.

The least ideal time is the interval between 10pm and 7am, as this is when people are sleeping; indeed, the regressions reveal that small hours is the least effective time of day category. Figures 5.13 and 5.14 provide empirical support for the ineffectiveness of late evening tweets. They contain two Vote Leave tweets that feature the exact same message and graphic and were even posted on the same day, but one was at 12.40pm and the second at 10.52pm. Although they were both successful, the earlier tweet had about twice as many retweets. As Matthew explained, ‘The lunchtime one—that has a shelf life of ten hours in a day, from noon until 10pm. The other one has a shelf life of at most an hour before you go to bed’.

FIGURE 5.13: Vote Leave Pledge Card Tweet (Midday)



Retrieved 2 October 2018 from Twitter. Screenshot by author.

FIGURE 5.14: Vote Leave Pledge Card Tweet (Evening)



Retrieved 2 October 2018 from Twitter. Screenshot by author.

Logan was the only interviewee who stated that 12pm was not a good time to tweet; he focused his 'tweeting energy somewhere between 2-6pm'. As he managed a student-focussed campaign, this interval was ideal. Although Paul believes that timing is important, he added a caveat that it could create an echo chamber effect: 'There's a self-fulfilling prophecy: you're reaching an audience that's available at those times'.

This effect also applies to *day of the week*, which was less emphasised by the interviewees, but still makes a difference in the context of politics and journalism. As Nick T. observed, 'the news cycle in Westminster is very attuned' to what happens early in the week. We Are Europe posted their key content on Mondays and Tuesdays for this reason. As a contrast to news cycles, Paul presented a lifestyle explanation for the effectiveness of Tuesday over Monday: 'People have just gone back to work on Monday, so it tends to be busier in your inbox... by Tuesday they're already fed up and they're looking at social channels more'. Wednesday and Thursday are both important for the political calendar because of Prime Minister's Questions—which according to Sam 'drives core political views'—and BBC Question Time, respectively. Friday was more controversial. On the one hand, interviewees said that it was not ideal because people are winding down and preparing for the weekend; on the other hand, it is when political magazines such as the *Economist* appear, and when people are searching for distractions after a busy week.

Although there was disagreement as to which day(s) of the week were most effective, the general consensus was that weekdays are better than weekends. Weekends are more

unpredictable, because people are following their own schedules. Sean also added that ‘sometimes people want a break from the news on the weekend’. The big exception, however, is Sunday. There are three hours of popular political shows in the morning: Sunday Politics (BBC), Peston on Sunday (itv), and the Andrew Marr Show (BBC).⁷² This contradicts the quantitative results, which illuminate that Saturday was more positively related to engagement odds than any other day of the week. It was also, however, the day on which the fewest tweets were sent, but not by a wide margin compared to Friday (frequencies are displayed in Figure 4.16). It is clear that less people were active on the platform on Saturday, but those who were active were more engaged. This makes sense in a campaigning context, as many events and canvassing sessions might have taken place on Saturdays.

5.4.2 *Timeliness and Novelty*

The importance of certain days of the week for political news cycles relates well to a more conceptual timing feature that emerged from the interviews: general *timeliness* or *topicality*, which relates to *newsworthiness*. As Sam declared, ‘It’s the story that drives the engagement for a tweet—how big, how relevant it is’. Because newsworthiness is dependent on ever-evolving external circumstances, it cannot be judged at the tweet level. It can only be determined through constantly monitoring the news and social media, and this is a strategy that We Are Europe employed with the help of digital marketing tools (Section 5.2).

One of the temporal control variables included in the models, the *number of days before the referendum*, captures a dimension of topicality. As the regressions reveal, its relationship with both retweets and likes is significant and quite strong: tweets sent closer to voting day had greater odds of success, given the increased timeliness of the referendum itself. Competition for attention, however, was extremely high—multiple interviewees observed that far more content was produced in the final days before the vote. As Paul declared, ‘If you looked at #Brexit the week of the referendum, it was continuously waterfaling down the page... It was really overwhelming trying to get a voice heard in that storm of tweets’. This is corroborated by the population-level dataset: over 12% of the tweets were from 22 June, and almost 10% were from 21 June. Thus, content from organisations was more likely to succeed because they could be heard above the noise. The most popular Stronger In tweet in the dataset was published only two days before the vote, and Will stressed the importance of this timing: ‘It was absolute peak. It was just before the BBC debate’ (Figure 5.15). The other major factor that explains the

⁷² As of September 2018, Peston on Sunday has been moved to Wednesday evenings as a result of too much competition with the Marr show (Hodson, 2018).

popularity of this tweet is the *celebrity image*—more shall be said on the effect of celebrities in the next section.

FIGURE 5.15: Stronger In’s Daniel Craig Tweet



Retrieved 2 October 2018 from Twitter. Screenshot by author.

The models do not take into account, however, a fuller definition of timeliness. Certain niche issues, collapsed under the ‘Other’ category of the issue variable in the sample-level models, became more salient in the final days before the vote. One such issue was Labour Leave, as so many members of the Labour party were undecided. Thus, according to Brendan, ‘Anything that said “we must get Labour to vote Leave” just did better’.

The concept of *novelty* accounts for why tweeting the same content is usually less and less effective, even if the content is initially resonant. In Mike’s words, ‘Sometimes if you have one successful tweet first, and then you repeat it, you’ve already exhausted your most dedicated network’. Nick T. observed this all the time because he promoted his articles with multiple tweets, and called it the ‘the law of diminishing returns’. Topicality has by definition a shelf life, and political news is especially time-sensitive. After a news item has come to the end of its shelf life, it is by definition no longer newsworthy, as Logan expressed: ‘Talking about events which are going on is always very useful because it’s new... if you’re talking about it after it’s been in

the news, then it's basically pointless'. It is no surprise, then, that one of Paul's personal tweets linking to a map from 2012 was unsuccessful—it is by definition not timely, and the text of the tweet makes this clear (Figure 5.16).

FIGURE 5.16: Paul's Map of EU Investment Tweet



Retrieved 2 October 2018 from Twitter. Screenshot by author.

One aspect of novelty is telling people what they do not already know. This is encapsulated by the concept of *originality*, which in turn is what makes a tweet unique or interesting. Many of the successful tweets from interviewees contained this feature. The most successful Green Party tweet is particularly emblematic because the originality is explicitly referenced in the tweet text by the words 'it's barely been mentioned' (Figure 5.17). The content may be able to speak for itself, but it is also effective to tell the audience that it is timely or novel. Moreover, Leni noted that 'they've omitted it' rhetoric resonates with people.


FIGURE 5.17: Most Successful Green Party Tweet



Retrieved 2 October 2018 from Twitter. Screenshot by author.

Another category of novelty is telling people what they *do* know, but from a different angle. In other words, being *counterintuitive* and *surprising*—contradicting what people assume or take for granted. When commenting on Healthier In’s most successful tweet (Figure 5.18), Mike called this technique a frame shift: ‘You thought it was that, but it’s this. You thought it was immigrants, but it’s underfunding’. The fact that the NHS was strained due to underfunding was only ‘known in the health community’; the general public believed that it was due to immigration, because that is what the tabloid press had been publishing. This is another example of Twitter being used to counter the mainstream media narrative.

FIGURE 5.18: Most Successful Healthier In Tweet



Retrieved 2 October 2018 from Twitter. Screenshot by author.

The Labour Leave campaign in and of itself represents counterintuitive novelty, because the mainstream media presented Labour as pro-Remain. This popular assumption gave all of their content an advantage. One of the campaign’s successful event tweets visually illuminates this (Figure 5.19). As Brendan explained:

The country was led to believe that if you were a Conservative you were for Leave, if you were UKIP you were for Leave, and if you were anybody else you were for Remain... And we wanted to have an image showing that actually what you're being sold is a load of nonsense. All political persuasions are divided on this; all areas of the country are divided on this.

Novelty is an effective strategy, but too much novelty can backfire, according to Nick T. There has to be a balance between the familiar and the new: ‘You need something that’s novel enough but recognisable enough. If it’s too novel, people can’t really relate’. This

underscores how difficult it is to measure this concept—it is not only context-dependent, but also highly subjective.

FIGURE 5.19: Labour Leave Camden Rally Tweet



Retrieved 2 October 2018 from Twitter. Screenshot by author.

5.5 Rhetorical Content

The interviewees did not use content features to explain tweet success as much as discoverability or timing features, which implies that they are of secondary importance: after all, if a tweet is not seen, its content does not matter. This is also reflected by the fact that specific questions sometimes had to be asked about these features to elicit strategies relating to them. Thus, on a high level the qualitative results reinforce the quantitative results: discoverability makes a larger difference than content. Nevertheless, content features represent important strategies that come into play once tweets *are* seen: they might help explain why some are more resonant than others.

Some interviewees, however, are convinced that content *does* matter—and also by implication *style*, which is the subject of the next section. As we have already seen, content is important in the context of live debates, as well as for accounts with many followers. Even outside of these contexts, where random chance has a bigger effect, content has significance.

Andrew made a marketing analogy when describing this significance: ‘Some of it is luck... But generally it’s things to do with the content: it’s how you frame it, it’s how attractive you make it. You’ve got to have something that’s worthy. You’ve got to have a product to sell’.

Which aspects of a tweet help to ‘sell’ it and hence make it more resonant? Mike’s remarks about message testing in the epigraph represent a strategy that some interviewees adopted to answer these questions. As Nick T. described it in the context of sharing his articles, ‘If I put this out with a slightly different bent... would it cause the same effect?’ This applies to tweets about successful articles, which immediately receive engagement. With unsuccessful articles, he tries ‘a totally different bent to see if the tweet changes the uptake’ (but often it doesn’t work, as resonance is usually ‘intrinsic’ to the article). The tweet represents a headline, and changing the bent means finding a different angle or perspective. Will also expressed the technique in this way:

You are trying different angles. Sometimes we’ll have used different texts with the same graphic, because we’re experimenting. We’re trying to see which version will take off. And that’s why we’ll brainstorm half a dozen.

Neither Nick T. nor Will was more specific than this when describing their experimentation with content. The rest of this section will review a selection of concrete content features that contribute to the expression of different perspectives. It begins with the four that were included in the models, and concludes with one that was not. The interview data help to contextualise the model results, and thereby further test hypotheses 3a, 3b, 4a, 4b, 4c, 5, and 6.

5.5.1 Emotions

Many interviewees did not make a distinction among different emotions, which justifies dichotomising the variable in the models. Alberto directly stated that it is the most effective strategy. Sean declared that in campaign contexts in particular, emotion is more important than rationality, which is why it informed We Are Europe’s three-pronged approach to content (which will be discussed in Subsection 5.5.4). Their most successful tweet (Figure 5.1) contains a vivid example of one particular emotion. As Harriet described it, ‘There’s a severe “Oh my God that’s disgusting”’. Thus, both Berger (2013) and Brader (2006) are affirmed by the qualitative data.


In light of this, one could argue that UKtoStay.EU had an inherent disadvantage, as the entire purpose of the campaign was to communicate well-researched facts as objectively (and hence authoritatively) as possible. Paul was well aware of the tension between this fact-based approach and emotion, and admitted that the former was often not compelling. What is

interesting, however, is that *emotional context can evolve*—UKtoStay.EU became more emotive after the referendum, and this was a deliberate shift in rhetorical strategy. As Paul explained:

We're not fighting a referendum to convince unconvinced middle-ground voters anymore. At this stage, it's dealing with the fallout of that referendum. The goal is even to change people's minds who voted no... We're finding more and more that being a little bit emotional is resonating extremely well.

There was less of a compromise between objectivity and emotion *after* the referendum, because the emotion was justified by the aftermath of the result, during which (in their view) many of their arguments were proven right. The emotion was specifically anger, which Paul stressed was 'born out of the fact that we repeatedly pointed out this stuff over and over and over again, and that we're still providing the research to back most of it up'. His argument that the anger was justified as opposed to unfounded reflects what multiple interviewees said about negative campaigning, which will be discussed in Subsection 5.5.4.

FIGURE 5.20: Scientists for EU Banner Tweet



Retrieved 2 October 2018 from Twitter. Screenshot by author.

Justified emotion is one thing, but facts and emotions in and of themselves are not mutually exclusive, as one of the successful Scientists for EU tweets demonstrates (Figure 5.20). Mike loved this tweet because of its emotional hit, which is conveyed through the image of people in a lab who 'care[d] enough to make a banner' with Scientists for EU on it (they had

sent the photograph to the campaign). Mike embellished the image with statistics from a diversity of publications, all indicating that ‘the science community as a whole is overwhelmingly for the EU’. Delivering the facts through the vehicle of an emotive picture allowed the information to be communicated ‘in a very human way’. Moreover, the reference to the 5,500+ scientists who had signed the letter to the *Times* in the tweet text establishes credibility and authority. This was the third most popular Scientists for EU tweet in the dataset.

Emotion was hardly ever discussed on its own—it was usually mentioned in conjunction with other content or style features. Sean specifically related campaigning to *issues* when discussing emotion: ‘If it’s an issues campaign, that’s your first go-to... I prefer to have the themes there so people will have the ability to check up and read those things, but the emotions come front and centre’. Labour Leave’s most popular tweet is an excellent example of the effectiveness of emotionally framed resonant issues (Figure 5.21). As Brendan explained, the words in the tweet ‘appeal to very core basic ideals: our NHS, our public services, democracy, inequality... You’re using that emotive language combined with things that concern people anyway, so you sort of give them a double whammy of feeding the emotion toward what you are saying’. This ‘double whammy’ of issue and emotion was more than a recipe for resonance: the tweet received 1,365 retweets, almost twice that of their second most popular tweet.

FIGURE 5.21: Most Successful Labour Leave Tweet



Retrieved 2 October 2018 from Twitter. Screenshot by author.

When speaking about the context of campaigning, Nick T. *did* distinguish between the effectiveness of different emotions. Enthusiasm works for other domains, but not for politics: ‘Generally if I say “This is a great thing”, people are like “Okay, whatever”’. Nick explained the effectiveness of negative emotion with the concept of *conflict*: ‘“This bad thing is happening; it’s really bad. Here’s why it’s bad”’. That’s what people engage with, in my experience... A story about people who just wander around and have a good time, it’s not really a story’. Anger and anxiety might inherently be more engaging than enthusiasm. This supports and qualifies

Brader's (2006) contention that emotion in general makes political advertisements more attention-arresting. It is also evidence against expressive theories, which view enthusiasm as more effective than anxiety because emotion is a coping mechanism (Roseman, Abelson, & Ewing, 1986). Such theories have limited relevance in the contentious domain of politics, where the ultimate objective is to win votes.

Roseman et al. (1986) also note, however, that anger might be appealing because it is associated with cognitions that presage victory. This is more relevant to campaigning contexts, and is supported by the interview data. Emphasising the strength of expression, Paul mentioned that 'hyperbolic or angry tweets always do great'. Nick T. provided an example of the effectiveness of hyperbolic anger. He had come across an account with 300 followers that sent a tweet that received 10,000 retweets: 'It is literally just him saying "Fuck Farage, fuck Rees-Mogg..." There's nothing else to it'. There is a trade-off, however, between short-term and long-term success. Despite the fact that his tweet was extremely successful, the user still has about the same number of followers: 'People just like the tweet. But they don't think this guy is a fountain of knowledge that I'm going to then go back to'. Nick T. cares far more about maintaining his reputation as 'a serious commentator' than he does about authoring rabid tweets that might be exceptionally successful.

FIGURE 5.22: Harriet's Anger-Driven Tweet



Retrieved 2 October 2018 from Twitter. Screenshot by author.

When describing the background of their tweets, some interviewees revealed that emotion is not always a conscious strategy. This supports Marcus, Neuman, and MacKuen's (2000) theory of affective intelligence, which posits that enthusiasm and anger are governed by the disposition system—a system defined by reliance on *habits* (Chapter 2). One of Harriet's successful personal tweets was driven by anger so overwhelming that she *had* to share it (Figure 5.22). She had read an interview with Arron Banks, and his comments were 'just so incendiary': 'When someone says something as arrogant as that I can't bear it, so I really felt that people

needed to read this'. Harriet did not explicitly reference anger to explain the success of the tweet (it received 1,055 impressions and 30 total engagements, including 13 link clicks, which was 'really strong' for her personal account), although she acknowledged that other people were probably 'as pissed off with that statement as [she] was'. She did, however, reference its 'strong call to action' implicated by the anger-driven phrase 'bore the electorate to death'. More shall be said on call to action in the next subsection.

Given Nick's comments about the resonance of negative emotions, some remarks must be made about a contextual aspect of emotion: Paul's observation that Leave campaigners had an *emotional advantage* during the referendum period. There were two reasons for this. The first relates to the context of Brexit campaign rhetoric: Remain supporters had more difficulty with expressing negative emotion, because the Leave side had framed them as 'Project Fear'. Thus, any pro-Remain tweets that exhibited anxiety reinforced this overarching narrative, and backfired as a result. Although Remainers could have constructed the same narrative about Leavers, they did not do it first—the Project Fear frame had been set very early on, and reinforced by multiple Leave groups.

The second reason relates to the broader context of British politics. Many pro-Brexiteers had campaigned against the EU 'for their whole political careers', as Paul put it. Leave campaigners were therefore able to harness a lot of existing arguments and extreme anti-EU sentiment. Several of the interviewees provided empirical support for this. As Brendan declared, 'There's always been a Eurosceptic element within the Labour Party'. In fact, these Labour Eurosceptics helped 'exert pressure on leadership to commit to a referendum' in the 'build-up to the 2015 general election'. The Freedom Association's Better Off Out campaign goes back even further: it was, according to Andrew, the 'John the Baptist of all this'. Launched in February 2006, it is the TFA's longest-running campaign as well as the mainstay that 'generated the most publicity' for the organisation, and interest and support from its members.

Conversely, Paul noted that 'there was no pool of pro-EU sentiment that the Remain movement could tap into'. This meant that positive emotions represented a less resonant strategy during the referendum period for Remain campaigners. Nick H. provided evidence for the lack of pro-EU sentiment, as he represents the London branch of 'the oldest pro-European movement in the UK'. Launched in 1949 under Winston Churchill's initiative, the European Movement wound down after the 1975 referendum, leaving behind a rump movement even in London. When reviving this rump in 2015, Nick effectively had to begin anew, as the members were dormant. Since the referendum, however, the emotional tide has turned—Paul contended that the government's push for a hard Brexit has provided 'a lot of emotional firepower' to

Remain supporters. Because the risks have become more tangible, Project Fear is less effective as a characterisation of their alleged scaremongering tactics. Emotional arguments have allowed the Remain movement to coalesce, just as they allowed the Leave movement to coalesce during the referendum campaign.

On the whole, the qualitative findings support hypotheses 4b and 4c and do not support hypothesis 4a: tweets containing anxiety and aversion are more successful, whereas tweets containing enthusiasm are not. They are somewhat at odds with the quantitative results: emotion was not significant at the sample level, and had a negative relation with retweets at the population level. This can partly be explained by the fact that interviewees generally did not view emotion as an independent strategy; it works best when conjoined with other features, such as issue. Moreover, given that Leave campaigners had an emotional advantage over Remain campaigners, the effectiveness of emotional appeals for the former might have balanced out the ineffectiveness of such appeals for the latter.

5.5.2 *Resolution*

In the models, resolution was not statistically significant at the sample level, and was negatively associated with retweet odds at the population level. This feature was, however, mentioned by many interviewees as having a positive impact on tweet engagement.⁷³ Indeed, actions represent the purpose of political campaigns: the most obvious being registering to vote and voting itself. As noted above, We Are Europe purposefully leveraged the anticipated popularity of the ‘Kiss of Death’ tweet to direct people to the voter registration page (Figure 5.1). Sean claimed that the text of this tweet is an example of ‘standard good Twitter practise: short and snappy, call to action’. When commenting generally on successful tweets, Elise also noted that briefer tweets that give an action are the most successful—*action* is as important as *brevity*, a style characteristic that will be discussed in Section 5.6.

Resolution was expressed in the name of the official Leave campaign, and this was a deliberate decision. When the referendum question was changed from ‘Yes or No’ to ‘Leave or Remain’ in September 2015, the campaign had to rethink its branding. As Matthew explained, ‘We liked the idea of “Vote Leave” because it was a positive action’. Consequently, many pro-Leave tweets at once invoked the name of the official campaign and their main call to action. Labour Leave’s most popular tweet showcases the effective execution of this strategy (Figure 5.21). On the other hand, ‘Stronger In’ is more descriptive, and thus passive.

⁷³ *Call to action* was used as opposed to *resolution* in the interviews given its greater transparency of meaning (it is a marketing term).

The exact type of resolution is important, as well as how it is expressed. One popular type relates to engagement on Twitter itself: a request to *share if you agree*. This explicitly associates a retweet with an endorsement, and is the clearest form of resolution. Andrew was convinced that it ‘always has a positive effect’, and in fact, ‘if you agree please RT’ defined his tweeting style. The most successful Better Off Out tweet demonstrates the effectiveness of this technique (Figure 5.23); coincidentally, the second most successful Better Off Out tweet was exactly identical to this tweet, and was sent two days prior. To further support Andrew’s point, the identical Vote Leave tweets in Figures 5.13 and 5.14 were the second and third most successful from the campaign, and also exhibit direct requests to retweet. Paul also affirmed that ‘“Share if you agree” is a very strong call to action’, and added that in turn it ‘would result in more people following [the] account’. Thus, asking for direct support is a strategy that boosts not only the tweet, but also the user—agreeing with the tweet implies agreeing with the messenger, after all.

FIGURE 5.23: Most Successful Better Off Out Tweet



Retrieved 2 October 2018 from Twitter. Screenshot by author.

In light of this, the population-level model results are surprising—resolution was represented by an even more explicit subset (including ‘please retweet’) in these regressions, due to constraints of the supervised machine learning classifiers (Chapter 3). Logan, however, cautions against explicit calls to action because they sound clickbaity. He made this comment in the context of links as opposed to retweets, but the same principle holds: ‘You don’t want to

literally just say “Click on the link”... You want to dress it up a little bit: “This is what’s behind here”’. One example of a more indirect call to action is ‘here’s our video summary’. It must be noted that this example would not have been labelled as resolution in the content coding stage, as it would have introduced a measurement issue—only more explicit calls to action were recognised, as they are more objective. Logan’s point, however, captures a fuller definition of this feature and corroborates the model results.

Although actions relating to registering to vote and voting were effective, less urgent actions relating to non-Twitter activities were not. This reflects the point about action-based campaigning made in Subsection 5.3.2: many actions are represented by hyperlinks, which take users off of Twitter. Two of Harriet’s least successful personal tweets are examples of this—they ask people to add a Twibbon to their profile picture⁷⁴ and to join a Change.org petition. As Harriet explained, both of these actions replaced the retweet: ‘People will do one thing; they won’t do two’. She also hinted at a second reason when commenting on the Change.org tweet: ‘If you’re going to fill out this petition, you’ll have the chance to tweet from the end of that’. Users might prefer to tweet that they have signed a petition as opposed to share a tweet imploring others to sign, as it is both more persuasive and self-promotional. It should be noted, however, that for these tweets URL clicks is a more appropriate measure of success. According to this metric, both tweets were effective: the Twibbon tweet received four link clicks; the petition tweet received five. These numbers are higher than usual for Harriet’s account.

The tension between URL clicks and retweets is the strongest explanation for the quantitative findings. Overall, the qualitative data provide support for and against hypothesis 6: *If a tweet contains resolution, it will be more successful*. Coupled with the quantitative findings, however, there is more evidence against the hypothesis.

5.5.3 *Relevant Issues*

The regression results for issue are somewhat complex: only the categories healthcare and sovereignty were significant at the sample level, but including them is negatively related to receiving a like (no categories were significant in the retweet model). On the other hand, the population-level models indicate that tweets that contain *any* substantive issue are marginally more likely to receive retweets than tweets that do not. The interview data support the findings from the population level but not from the sample level. In general, the five topics identified as relevant during the EU referendum—the economy, immigration, sovereignty, security, and

⁷⁴ A Twibbon is a microsite that allows social media users to add a visual icon to a profile picture indicating their support of a cause, brand, or organisation: <https://twibbon.com>.

healthcare—were viewed as beneficial for a tweet to include (of the five, security was mentioned the least). The most successful Labour Leave tweet invokes most of them (Figure 5.21). It reinforced the main concerns of Labour Leave supporters ‘just before the final televised debate’, during which these issues were sure to emerge: ‘There was obviously the big issue of the money for the NHS in the campaign, and a lot of our supporters were also concerned with the TTIP bill between the EU and the US, which threatened our public services and democracy’. Brendan’s description of the context of the tweet makes it clear that *timing* and *relevant issues* worked in conjunction—in this context, a simple expression of general principles was very effective.

Throughout the entire campaign, UKtoStay.EU heavily featured the economy and sovereignty in their tweets, as they did not believe people understood ‘how [the EU] made decisions, how the UK influences those decisions, how the democratic aspect of it works’, as well as the EU’s economic impact. InFacts also featured the economic argument, which in Sam’s opinion was ‘the strongest side for Remain’. Another theme that interviewees deliberately focussed on was immigration. Alberto’s two most resonant tweets were about Turkey in particular, as there was ‘a big debate around Turkey joining the EU during the referendum’. They both conflate resonant issue with *topicality*, as they are premised upon breaking news items. One responds to an MP’s comments about Turkey’s accession to the EU; the other quotes a tweet about visa free travel to the UK for Turkish people (Figure 5.24). These two tweets were, coincidentally, each shared over 2.5 times more than Alberto’s third most popular tweet.

FIGURE 5.24: Alberto’s Turkey Tweet



Retrieved 2 October 2018 from Twitter. Screenshot by author.

According to Nick H., migration represented a structural flaw in the Stronger In campaign: ‘they would never make the pro-European case’ for it. This created a novelty effect for other Remain campaigns tweeting about the theme, such as UKtoStay.EU. They shared a link to an *Independent* article, which gave ammunition to something that Remain supporters had been saying for awhile: how wrong Leave was on immigration (Figure 5.25). Like Labour

Leave's tweet, it was sent just days before the referendum, which made it timely as well. At this point, in Paul's judgement, migration was appropriate to mention: 'It's an issue that had been taboo to the Remain side up until the final weeks'. The topic was indeed resonant, as the tweet was shared almost nine times as much as their average tweet during the referendum period.

FIGURE 5.25: UKtoStay.EU's Immigration Tweet



Retrieved 2 October 2018 from Twitter. Screenshot by author.

Apart from these high-impact general issues, there were a number of specific issues that were quite resonant with certain subgroups, as they reflected their niche values and interests.⁷⁵ A good example of this is *education* for YEM's audience: their most popular tweet showcased a video of the EU's contributions to UK science, which included the fact that the UK receives much more funding than it sends out (Figure 5.26). As YEM's target demographic benefits directly from this funding, it is no surprise that the tweet did so well. Logan declared that had the campaign not discussed university educational funding, there could have been 'a very different overall result in terms of young people', 75% of whom voted for Remain.

⁷⁵ During content coding, they were all collapsed into the 'Other' category. In the population-level model, this category was merged with the five other issue categories to create a binary variable (contains issue vs. none).

FIGURE 5.26: YEM's EU and British Science Tweet



Retrieved 2 October 2018 from Twitter. Screenshot by author.

Another reason why niche issues work so well is because they are inherently more novel. This has already been demonstrated by the most successful Green Party tweet, which clearly states that their main argument has ‘barely been mentioned’ in the debate (Figure 5.17). As Leni asserted, ‘You would’ve had all the other parties talking about the economics side of it, or migration, whereas I’m not sure if the others explicitly were talking about the environment’. In a sense, the Green Party represents this niche issue in and of itself. Another example of a campaign symbolising a niche issue is Scientists for EU, which as discussed above received instant popularity on social media, followed quickly by mainstream media coverage. This is highly unusual for a grassroots account. Their equivalent on the Leave side was Fishing for Leave, which Mike himself noted. Their voices might be small offline, but they can be loud on social media.

As Nick T. noted, novelty can backfire—too much of it might not be relatable. This also applies in the context of referendum issues, as some specific issues were *too* niche. An example is given in Figure 5.27: Elise’s tweet promoting Scotland Greener In Europe’s Twibbon campaign. Although Scotland is relevant to Elise—she has Scottish heritage and supported their independence referendum—it is less relevant to her followers, many of whom are based in England.

FIGURE 5.27: Elise's Scotland Greener In Tweet



Retrieved 2 October 2018 from Twitter. Screenshot by author.

Finally, a qualifying consideration must be made with regard to this feature: sometimes *public* measures of engagement aren't likely not because content isn't resonant, but rather due to *lack of social cache*. While working for a sustainability platform, Harriet discovered that social messages were 'shared with impunity', whereas environmental messages were read for longer periods of time (a *private* measure of engagement), but not shared: 'There's obviously this association with being an environmentalist that people don't want to have associated with their personal brand, and yet they're concerned'. We are identified with, and thus carefully curate, what we post for our friends and followers on social media. Trendy topics are by definition dependent on social, political, and cultural context—in recent years, environmental issues have actually gained social cache because of increased evidence for climate change and the organic food movement.

Figures 5.5 and 5.6 illustrate this qualifying consideration: they are about LGBTQ rights and gender violence. Although Logan affirmed that the inclusion of an image was the most important differentiating factor, he also acknowledged that theme played a role: 'People on Twitter definitely don't talk about gender using hashtags, and people very rarely even discuss it explicitly, whereas people are very likely to discuss LGBT+ explicitly'. Even feminism, he said, is taboo on the platform. Harriet also sent an unsuccessful tweet about gender, specifically about women's rights. She noted, however, that there has recently been 'a watershed moment in terms of feminism'—we now 'have more language around this idea that feminism doesn't just apply to women', which means that the majority of users on Twitter are now comfortable talking about it.

All in all, the interview data provide weak support for hypothesis 5: *If a tweet contains relevant issues, it will be more successful*. It is difficult to make solid assertions about the effectiveness of any particular issue, because timing is an extremely important conditioning factor for this feature—different issues are more or less relevant at different times. Moreover, there are many subcategories of each issue, and some subcategories are more effective than others (e.g., Turkey is an effective subcategory for immigration). Finally, some niche issues are more effective than

the five main relevant issues. More shall be said on niche issues as a strategy for grassroots success in Chapter 7.

5.5.4 *Message Framing*

Message framing was an extremely controversial and nuanced topic. It was not explicitly mentioned by many interviewees, but rather had to be elicited. This indicates that it is perhaps too high-level of a concept—an assessment of tweet content in relation to the position of the sender is required. No firm conclusions can be drawn about whether positive or negative campaigning is more effective, which reflects the regression results: message framing was not statistically significant at the sample level; at the population level, all that can be said is that negative campaigning is less effective than not including either frame.⁷⁶ Finding evidence in support of or against hypotheses 3a and 3b is not the right approach, as many contextual factors need to be taken into consideration. Thus, the question should be *under what circumstances and from whom is negative or positive campaigning more effective?*

Before this question is addressed, it must be noted that many interviewees were uncomfortable with the term *negative campaigning*, due to its unsavoury connotations. The definition adopted for this research, which informed the content coding stage, is quite neutral: ‘anything negative said about anyone or anything related to the opposite side’ (Appendix C). This is not constrained to critiques of individuals, but also includes critiques of consequences resulting from a vote to Remain or Leave. Because many interviewees conceived of negative campaigning as unfounded name calling (an extreme version of the former), hardly anyone admitted to using it. Instead, interviewees emphasised the importance of *fact-based arguments*, because facts are more defensible than subjective opinions; indeed, both UKtoStay.EU and InFacts are premised upon this. Sean asserted that the best strategy ‘is to create friction but actually back it up’. Brendan’s advice was that if attacks were necessary, they had to be done ‘intellectually, not personally or professionally’.

While Sean upheld the importance of fact-backed arguments, he also acknowledged that it became divisive late in the campaign. At this point, according to Harriet, there was a war of ‘fact vs. facts; it was this big fact-off’—and the facts from the two sides were both high volume and mutually exclusive. Sean described the effect of this as *double-think*: ‘You start going “What is true? Is it just my opinion?” And it creates this unease’. Truth is even more obscured, as opposed to discovered. Harriet commented on the larger implications of this for Twitter: ‘Trust

⁷⁶ It cannot be compared with positive campaigning, which could not be included in the model due to low classifier accuracies.

in social just plummeted after that referendum’. These observations highlight yet again the importance of context for rhetoric: too much of anything can backfire; strategies have to be dynamic as opposed to static.


With this in mind, what are the contexts in which negative campaigning *is* effective? To Martha, if success is measured as retweets, negative campaigning is *always* effective: ‘It’s a fact that things that are negative get shared more than things that are positive’. This is neither supported nor contradicted by the quantitative results, as message framing was not statistically significant in the sample-level regressions. It is more useful to consider the effectiveness of negative campaigning with regard to specific audiences. Will affirmed that as a result of ‘huge amounts of public opinion research’, Stronger In decided to ‘emphasise the risks to people of leaving’, as these messages ‘resonated best with the British public’. He specifically meant the *undecided middle*, which was their target audience for this strategy: ‘The 30% in the middle were people whose hearts said they wanted to leave... But their heads said “Is it worth it?”’ Because this audience was emotionally predisposed against the EU, negative campaigning was more effective than ‘just talking about the benefits of being in the EU’. Will’s description also invokes a *fact-based* approach: appealing to heads, as opposed to hearts. The strength of his confidence in this strategy is reflected by his belief that had Stronger In ‘run an entirely positive campaign’, they probably would have ‘lost by more’. It should be noted that Will never used the term negative campaigning, but rather ‘the risk-based approach’. Elise also observed that negative framing is necessary for persuading the undecided: ‘If they see that the people they’re aligning with are saying dreadful things that are blatantly untrue, there’s more of a chance that they’ll question what else they’re being told by the Leave campaign’. The power of this strategy is its ability to undermine the credibility of the opposition.

Another factor to consider is situation: debates represent a context in which negative campaigning is both rampant and effective. The entire point is to not only make original arguments, but also rebuttals. We Are Europe’s second most successful tweet as measured by reach and engagement called out Michael Gove during a BBC Question Time debate; it received 73,260 impressions, 638 retweets, and 296 likes (Figure 5.28).⁷⁷ In Sean’s words, ‘It went mental’. As he explained with a war metaphor, relentless negative campaigning during live debates was part of We Are Europe’s strategy: ‘In the context of machine guns, you have to keep putting it out. When they would say something just completely daft, you would just quote it and go “nope!” But you were doing it live so as people were looking at Twitter, they see it and go “oh that’s funny” and retweet it’. Thus, this strategy was effective partly because such instant

⁷⁷ This tweet was not part of my dataset, but was mentioned during the interview.

rebuttals were playful. The intersection between humour and negative campaigning will be discussed in Section 5.6.

FIGURE 5.28: We Are Europe's Gove Rebuttal Tweet



Retrieved 2 October 2018 from Twitter. Screenshot by author.

Stronger In took planning to the extreme by orchestrating Twitter war rooms during debates. Tweeting content during these debates had high stakes, because a debate has to have a winner—and the winner is reported in the media. In Will's words, 'the media are looking at who has the best rebuttal'. He described the setup and purpose of the war rooms:


We invited people with Twitter accounts and significant numbers of followers—so politicians and other surrogates—to come into our offices to watch the debates, and the research guys would say 'Tweet X'. So let's say that Michael Gove had said something that was wrong; we would say 'Tweet Michael Gove is wrong for these reasons'... You're trying to win the battle of resonance, the battle of perception, about who has won that debate, who had the better arguments.

At the end of a debate, Stronger In's research team would email journalists their interpretation, which reflected in large part what they had been tweeting in real time during the debate. Matthew used his personal account to influence both media and public perceptions by sending 'verdict' tweets after debates. Regardless of how well Leave had done, it was important to express that they had won.

Beyond the context of debates, negative campaigning is effective for *rebuttals* in general. Graham emphasised the necessity of this: 'You've got to be negative if you're rebutting

something'. InFacts was actually set up as a rebuttal outfit—its founding purpose was to point out when Leave campaigners said something that wasn't true.⁷⁸ Will and Matthew both recognised that this technique is important for influencing mainstream media. Indeed, extremely successful tweets can actually become news stories in and of themselves, as Leni observed. She provided a recent example from the Green Party of this. On 2 October 2017, Caroline Lucas tweeted a photograph of herself and co-leader Jonathan Bartley laughing, in response to Michael Gove's claim that the Tories were the 'most ambitious green party' (Figure 5.29). The next day, her tweet was the headline of an article in the *i*, a British daily newspaper—and the photo was the main image of the story (Figure 5.30). Because the Green Party is not mainstream, it relies heavily upon social media to reach traditional media. In Leni's words, 'By putting out something that's engaging, we have the equal opportunity to actually go on there without having to rely on a journalist or an editor wanting to feature us'. This is also another example of an effective visual that literally speaks louder than words, and of the intersection of playfulness and negative campaigning.


FIGURE 5.29: Green Party Laughing at Gove Tweet



Retrieved 2 October 2018 from Twitter. Screenshot by author.

⁷⁸ Sam noted, however, that in light of Stronger In's negative campaigning, InFacts began to produce 'more positive content to try and make the case as well as rebut their case'. This relates to their strategy shift from focussing on influencers to targeting a more general audience.

FIGURE 5.30: Green Party Laughing at Gove / Article



Retrieved 2 October 2018 from *i* News. Screenshot by author.

It is no coincidence that the two examples given above and the example in Will's quote target Michael Gove. *Particular subjects* of negative campaigning are more effective: public figures and organisations that are unpopular with specific audiences. This contextual feature is difficult to measure, as it is inherently subjective and requires political and cultural understanding. Sam's attack on Douglas Carswell is a good illustration of the fact that he was unpopular among Remain supporters during the referendum period (Figure 5.31). Carswell had said in an interview that Remainers had 'nothing to offer but fear and paranoia', whereas Leavers were 'the sunny, optimistic, globalist vision of Britain'. Sam's mockery of these words worked 'mostly because a lot of people dislike Douglas Carswell'. This was not only his most popular personal tweet from the dataset, but was also seen by Carswell himself, who blocked Sam as a result. Similarly, Nick T. ascribed the success of his anti-Farage tweet to '“Farage has done something stupid”: That's what people are reacting to. It's people who don't like Farage... there's a definite audience for that'. Incidentally, although Carswell and Farage are both members of UKIP, they strongly dislike each other.

Nick T. also remarked that negative campaigning is effective if it's *relatable*—otherwise it backfires because it 'feels like scaremongering'. He critiqued Stronger In's execution of the strategy because in his view people could not connect with it. The numbers they used were too dramatic, and represented a metaphorical cliff edge:

The problem with the cliff edge is that it sounds very scary but it's also unknowable. It's like the abyss. No one really believes in the abyss until they look into it. They should have created more realistic scenarios that people could have immediately identified with. As soon as you say '50,000

people will lose their jobs’, no one can connect with that. But if you say ‘These specific 1,000 people who live in Ashfield, they will lose their jobs’—then suddenly people can connect with that.

Nick’s point about relatability qualifies the point about fact-based arguments made earlier: such arguments are not always more defensible. The type of fact matters; it has to be believable.

FIGURE 5.31: Sam’s Anti-Carswell Tweet



Retrieved 3 October 2018 from Twitter. Screenshot by author.

In general, interviewees who exclusively supported positive campaigning critiqued negative campaigning in their justification of their positive strategy. Brendan declared that negative framing *puts people off* and *lacks substance*, as it does not represent the campaign’s core values and beliefs: ‘You’ve got to first of all offer people a) hope that your argument is better and will offer something more positive and b) that you are in command of your own information’. Those who go on the attack might not know what they are saying. Labour Leave’s most resonant tweet is a classic example of positive campaigning; it effectively encapsulates the core principles of the campaign (Figure 5.21).

Brendan’s most successful personal tweet demonstrates one context in which positive campaigning is effective: live events. Tweets from live events are especially resonant if they contain a visual item, as displays of solidarity are powerful. Sent from a Labour Leave rally in Stoke, Brendan’s tweet demonstrated the magnitude of support for Brexit within the Labour Party—over 1,000 people were there, and this is evidenced in the photograph (Figure 5.32).

Brendan believed that ‘the reason it did better is because a lot of people at the event probably retweeted it’.

FIGURE 5.32: Brendan’s Labour Leave Stoke Rally Tweet



Retrieved 3 October 2018 from Twitter. Screenshot by author.

David also noted that negative campaigning ‘can often put people off’. Although it obviously breeds short-term success, it always comes back ‘to bite you in the arse’. Thus, negative campaigning has *short-term gain and long-term pain*, whereas the worst consequence of positive campaigning is being ignored. An example of this trade-off is the Vote Leave battle bus, which flaunted the message that the £350 million a week sent to the EU should fund the NHS. Although it was very effective during the campaigning period, since the referendum it has ‘created humongous levels of distrust, of anger’—and the questions it has generated about the NHS will never go away for Boris Johnson and Michael Gove. Therefore, David’s ‘general mantra of campaigning’ is to always represent the ‘benefits rather than the negatives’.

One of Logan’s justifications for positive campaigning is grounded in both political and temporal contexts. Because British millennials, the YEM target audience, grew up in the post-Thatcher milieu of polarised politics, they are not receptive to negative campaigning—they have heard it all their lives. Thus, for them positive campaigning has the effect of *novelty*. David made a very similar comment on the contemporary political context: ‘We’re now in a massive vacuum of positivity in politics, and it’s putting off people’. Vitriol was not always the norm, however: Logan pointed out that before the 1980s, there was an ‘era of consensus politics’. To people

born in this era, ‘negative campaigning seems very new and exciting’. This insight demonstrates that neither strategy is inherently better: it depends not only on timing, but also on the age of the audience.

A narrower temporal context comes into play as well: that of the referendum campaign itself. Sean observed that We Are Europe’s series of anti-Boris videos were published ‘too late in the day, when people were getting fed up of just laying shit onto people like Boris and Farage’. This qualifies the strategy of targeting unpopular figures for negative campaigning—this strategy might backfire if many others have already employed it. Once again, novelty plays a conditioning role.

Talking all of the arguments for positive and negative framing into consideration, it is no surprise that some campaigns believed that a blend of both is most effective. We Are Europe is one such campaign. As Harriet explained, ‘We had a three-pronged messaging strategy: inform and inspire, amuse or entertain, and cause friction’. The first two prongs relate to positive framing, the third to negative. Entertaining content drives people’s choices on social media, and is more likely to be shared; it was also important to feel really hopeful against the backdrop of negativity. At the same time, however, undiluted positivity is not believable—which was why the friction was necessary. Too much negativity, on the other hand, is not accessible. Like other interviewees, Harriet never referenced negative campaigning as such, but used the words *friction* and *pointy* instead. These words evoke its texture and effect.

Mike also advocated using a blend of positive and negative framing, but expressed the opposite focus: negative campaigning has more ‘immediate resonance and credibility’, but too much negativity burns people out. He associated positive campaigning with naïveté and marketing: ‘You just look like a happy clappy promo agent’. The implication of this is that positivity is instinctively questioned: ‘We have inclinations to feel that the negative is more authoritative than the positive. Positive cases are harder to build and keep robust’. Unlike We Are Europe, Mike viewed *negative* campaigning as entertainment. But if it goes too far, energy and the core values of the campaign are lost; this relates to Brendan’s point that negative campaigning lacks substance. Mike recommended a ‘this is good, this is good, this is good—oh my God stay away from this’ approach. To use a popular metaphor, the positive represents the carrot and the negative the stick that together ‘move the donkey’.

The effectiveness of dynamic strategies that integrate positive and negative campaigning is illuminated by the fact that unpredictable events have a large impact on which approach is more appropriate. The murder of Jo Cox on 16 June is one such event, as campaigning came to a standstill as a result. According to Harriet, negative campaigning was far less effective in

the wake of the murder. We Are Europe had a crisis meeting and decided to switch to a ‘message of hope and love’, as this is what people wanted to see. The positive message is embodied by a video produced by their supporters that they shared two days before the vote. This turned out to be their third most popular tweet, as measured by impressions: it was seen by 53,556 people.

Vote Leave had a similar strategy, in light of not only Jo Cox’s murder, but also Nigel Farage’s ‘breaking point’ poster that emerged on the same day. The poster was immediately reported to the police for inciting racial hatred (Stewart & Mason, 2016). According to Matthew, both of these events interrupted the positive momentum of their campaign, and caused some Leave voters to have second thoughts. To fight against this, they also tweeted a strong message of hope on 21 June: a clip from Boris Johnson’s speech during the BBC debate (Figure 5.33). This was the most resonant message sent from the Vote Leave account, and it was resonant precisely because of its circumstantial context. As Matthew asserted:

It was ‘take back control’ message, ‘this will be our independence day’, and it was really seminal to the campaign because there were swathes of people who perhaps had had slight doubts the week before, who then thought actually, I can go into the polling booth and vote Leave with a positive frame of mind.

Had the tweet been sent earlier, during their period of positive momentum, it probably would not have been as effective.

FIGURE 5.33: Vote Leave’s Message of Hope Tweet



Retrieved 3 October 2018 from Twitter. Screenshot by author.

To conclude this subsection, another contextual feature will be discussed: *the identity of the messenger*. This coincides with *ethos* in Aristotle's rhetorical framework. To further complicate the positive vs. negative campaigning debate, it must be noted that the effectiveness of either technique depends on who is delivering the message. Two main dimensions are at play here: whether the messenger is an *individual* or an *organisation*, and the messenger's *personality*.

Leni highlighted the first dimension when explaining the effectiveness of positive campaigning. It is not only more personal, but also defines the Green Party in the public's view: their general slogan, after all, is 'Greens for a Better Europe'. Their audience perceives them as 'a very optimistic positive party', and thus when they 'use a more negative attacky tone', they 'get told off'. Moreover, because they are an organisation, their language 'has to stay within certain limits' (i.e., be politically correct). The Party's co-leaders, Caroline Lucas and Jonathan Bartley, are quite active on Twitter, as is their MEP and Brexit spokesperson Molly Scott Cato. These individual accounts are able to be 'a bit snarky and more negative' by virtue of the fact that they are not organisations. An illuminating example of this is Cato's website 'Bad Boys of Brexit', which documents all funders of the Leave campaign and their suspicious activities. Although Cato's tweets about the website 'were the most successful of her whole Twitter history', when Leni shared it from the main Green Party account, people reacted unfavourably: 'Are you now resorting to this negative name-calling?' Because organisational accounts are judged and scrutinised more, the Green Party strategy is to use negative framing primarily from their leaders' personal accounts, and positive framing from the organisational account.

Stronger In was under similar scrutiny—as the core campaign, it effectively represented the Prime Minister. As such, Will adopted the 'Remainer-in-chief' role from his personal account: 'rallying the troops with a more positive and cheeky and irreverent take'. Independent cross-party organisations such as the Bruges Group are under even more constraint, because their messages have to be aparty political. Only in his personal tweets was Graham able to express UKIP policies. As mentioned earlier, UKtoStay.EU was under a reputational constraint: because it was branded as an objective, well-researched information campaign, its tweets had to adopt an authoritative tone. This authority was also the campaign's Achilles' heel—as Paul explained, 'social media is all about personalities'. To provide the organisation with a human face, he would occasionally retweet his personal posts that were 'a bit more emotional or more critical' from the UKtoStay.EU account, if they were relevant. Conveniently, this strategy improved the visibility of Paul's account as well.

Personality represents the second dimension of messenger identity. Matthew provided an example of this using Nigel Farage, who was a highly controversial figure during the campaign.

Farage's speech during the itv EU referendum debate 'was all pretty reasonable; he didn't do his big breaking point stuff'. In Matthew's opinion, had Michael Gove uttered the same words, he would have received much more approval because he is less divisive than Farage. This clearly demonstrates that 'it was a combination of having the right message and the right messenger'. In fact, Vote Leave was partly established to ensure that Farage did *not* become 'the voice and the face of the Leave campaign'—his years of saying negative and disgraceful comments rendered him less appealing to swing voters. Matthew also brought up a *personal style* difference between himself and Dominic Cummings, Vote Leave's Campaign Director. Matthew has a positive personality, whereas Dominic 'likes lashing out'. Thus, it is more effective for Dominic to exhibit negative framing.

Organisations have personalities too, and thus can also be differentiated along this dimension. Non-mainstream organisations can be more colourful than mainstream ones, as represented by the official campaigns and political parties. Unlike Stronger In, which was affiliated with the government, Scientists for EU could be 'more fun and cheeky and engaging'. In fact, they developed a reputation for this, and the very fact that science is *not* typically associated with such traits made it all the more effective.

5.5.5 *Celebrities*

The most-referenced content-related strategy not included in the models was including a *famous person* in the tweet. It has already been alluded to in the contexts of @mentions and targets of negative campaigning. Twitter allows fans to follow and interact with the public figures they admire; perhaps the best demonstration of this is when Justin Bieber trended for so long that the trending topics algorithm had to be changed (Dumenco, 2010). It is no coincidence that the most resonant tweets sent by both official campaigns featured two such figures: Daniel Craig (Figure 5.15) and Boris Johnson (Figure 5.33). These tweets have already been discussed in the contexts of timing and negative campaigning, but Will and Matthew also mentioned the celebrity factor as a reason for their success. Conscious of their cultural appeal, Stronger In built relationships with multiple non-political celebrities, including Craig. What was especially powerful about this tweet was that he 'had worn that t-shirt specifically for the campaign'—it was an original picture from James Bond himself. Although Johnson was not well-liked on the Remain side, he was 'loved' on the Leave side, so much so that he was 'worth a lot of likes and retweets' in his own right, as Matthew contended. The most successful Better Off Out tweet also featured a political celebrity, Dan Hannan (Figure 5.23). Andrew noted this before he mentioned the 'if you agree, please retweet' structure: 'One of the main attractions of it is the

man himself. Dan is very popular in Eurosceptic movements’. Returning to his marketing metaphor, Andrew claimed that any Hannan quote packaged in an attractive design ‘sells’. Conversely, several interviewees referenced prominent figures who were *not* resonant at the time of the referendum: Jeremy Corbyn and Tony Blair. Thus, a degree of political and cultural knowledge is necessary for the judgement of which famous people to include.

5.6 Rhetorical Style

If content is the *what*, style is the *how*. Andrew captured the difference between these two feature categories with his reflections on tweets that share media articles. Many news websites have a tweet button that generates an automatic ‘bog standard’ post based upon the headline of the article being shared. As Andrew observed, using this button is less effective than composing an original tweet: ‘No one’s interested in the title. What people are interested in is *how you share it*’ (emphasis added).

Style is arguably the most elusive dimension of rhetoric. As revealed by the last two chapters, it is incredibly difficult to operationalise, as it is not tied to keywords. Playfulness and message format were not included in the population-level models because accuracies for their machine learning classifiers were too low. Moreover, ‘good’ style is to a certain degree a matter of subjective opinion: a tweet can tell a joke, but how can the quality of the joke be systematically measured?

This section will review two stylistic features included in the quantitative analysis: playfulness and questions. Playfulness was not statistically significant; questions was significant in the like models, but was negatively related with like odds. It was also negatively related with retweet odds at the population level. Message format will not be discussed, as it was not highlighted to any substantive extent; this is reflected by the fact that it was not significant in any of the models. As such, only hypotheses 8 and 9 will be further tested. A third stylistic feature not included in the regressions but mentioned in the interviews, brevity, will also be discussed.

A qualifying factor must first be noted. According to Paul, account prominence has a substantial effect on strategy. Many of the discoverability and content-based strategies only apply to small accounts; larger accounts can be more focussed on style:

A lot of the stuff that I would have said is very rigid—always have hashtags, always have a call to action—is much more relevant when we were smaller. Now that we’re bigger, we can do what you see a lot of trending tweets do sometimes: if it’s a very witty tweet or a very clever remark, it can quite easily get loads of shares.

Mike voiced a similar opinion, stressing that random chance plays a far larger role for tweets from less popular accounts: ‘When you have accounts beneath a certain level of followers, then whether it catches or doesn’t is like trying to light fires if your material is slightly damp—it might take off or it might not. When you’re larger, the support for your tweets is more meritocratic’. The implication of *meritocratic* is that the quality of tweet content (and hence rhetorical style) actually matters for large accounts.


5.6.1 *Brevity, Clarity, and Simplicity*

Brevity, and the related concepts of *clarity* and *simplicity*, was the most referenced style feature not captured by the models. This comes as no surprise, as it literally defines the platform: up until 7 November 2017, tweets had an 140-character limit (Perez, 2017). As explained in Chapter 1, sending a tweet is supposed to be similar to sending a text message; Twitter’s affordances are premised on its mobile-first approach. The advantage of brevity has already been seen in videos (Section 5.3), and was alluded to in the subsection on resolution. Sean and Elise both consider short, action-oriented tweets to be best practise.

Why are short messages effective? Because they tend to be clear and thereby easy to understand. Alberto’s quote tweet about visa free travel, already discussed in the section on issues, is a good example of this (Figure 5.24). He acknowledged that clarity was a contributing factor to the tweet’s success. Sean’s comments on brevity operationalise the concept a bit further. On the one hand, sentences should ‘finish when they should finish—you don’t have extra words’. On the other hand, each tweet should have a single focus: ‘If you’ve got two messages, make it two tweets. So one message per tweet, please, keep it simple stupid’.

Logan recognised short tweets as one of two main stylistic strategies, and associated brevity with playfulness as opposed to being information-heavy and very detailed. Politicians tend to exhibit the latter strategy, and because they are serious accounts making serious points, people are ‘willing to read the whole thing’. This is another example of the importance of messenger identity. The former strategy, which focuses attention more on the messenger than on the substance of the message, is the ‘short, quirky, emojis, swearing approach’. It is demonstrated by a successful YEM tweet about Ryanair (Figure 5.34). As Logan observed, ‘There’s nothing in this tweet which is actually saying “vote Remain”. It’s just a way of saying “look at us” basically’. The message is at once playful because of the emoji and colloquial because of the language (‘nice one’)—this informal diction is far clearer than technical jargon.

FIGURE 5.34: YEM Ryanair Tweet



Retrieved 3 October 2018 from Twitter. Screenshot by author.

5.6.2 Playfulness

Nothing can be concluded from the quantitative analysis about playfulness, as the variable was not significant in any of the models. The qualitative data, on the other hand, affirm the importance of playfulness as a strategy for tweet success. Given the brevity of tweets, wit and humour are extremely characteristic of Twitter discourse—the connection between these two features is captured by Logan’s description of short tweets as a strategy, and illustrated by Figure 5.34. As Harriet asserted, ‘Twitter’s about showing that you’re clever’. In her view, the best Twitter influencers are ‘the ones who say the unsayable and say it really sarkily—it just flies’. Denis drew a parallel between tweeting and public speaking with regard to this strategy: it is effective because ‘when you get people laughing with you, smiling with you, then they’ll listen to some of your more serious points’. Thus, humour improves a messenger’s *ethos* by allowing for a rapport with the audience to be established.

Nick T. provided an example of a Twitter account that in and of itself demonstrates the effectiveness of playfulness as a strategy. On election night in 2015, someone spontaneously created an account called Paddy Ashdown’s hat, and tweeted ‘I’m Paddy Ashdown’s hat’

followed by a handful of other quippy remarks. Within 24 hours, the user had 20,000 followers—and then tweeted something that perfectly captured the essence of the situation: ‘20,000 followers for this? This is why nothing good ever happens’. Indeed, such parody accounts are quite common on the platform.

Not all accounts can be equally playful, however, and this helps to explain the inconclusiveness of the quantitative results. As observed in Section 5.5, grassroots organisations can be more fun and cheeky than established organisations, and Scientists for EU is an example of this. Occupying the niche of ‘being entertaining and humorous’ was their conscious strategy, as it was not only engaging, but also a technique that Stronger In could never fully adopt—their attempts at humour ‘didn’t work because everyone understood that they didn’t have the liberty to just mouth off whatever they thought’. Conversely, Scientists for EU was a ‘disposable campaign that could be cheeky’, and this cheekiness was authentic. It must be noted, however, that the most resonant Stronger In tweet in the dataset had playful text due to its alliteration (Figure 5.15). Although Will mentioned that the tweet is ‘not trying to bombard you with facts and figures or complicated arguments’, he did not explicitly reference its quippy tone—timing and celebrity were more important features.

Mike also declared that personal accounts can be even more playful than organisational ones, and that this playfulness is more effective: ‘People like supporting a witty individual who just said something witty. Whereas if you’re retweeting a campaign, you’re aware that you’re supporting a campaign message. It’s got a slightly different feel’. Harriet echoed Mike’s comments in her observation that individual accounts ‘can be a bit sardonic in a way that perhaps organisations can’t be’. Being playful is part of Martha’s personal brand, and allows her to establish her ‘own reputation of credibility’. Importantly, playfulness can be used to deliver a serious message, and should not compromise politeness: ‘It’s coming at it with a slightly more light touch. Being cheeky, being provocative, but respectful. Not just here’s another really boring statistic’.

The best example of a playful tweet from the interviewee datasets came from Scientists for EU (Figure 5.35). The tweet was resonant precisely because it communicated one of the main referendum issues, peace and stability in Europe, in a non-preachy (and thereby non-condescending) manner. As Mike explained: ‘So you had this “war war war war war war” going down—it pulls people’s attention in, and then it seems funny that it should be this long chunk of war, and then just “bananas”’. It was not only the most popular Scientists for EU tweet, but also one of the most popular tweets in the population dataset, albeit in a slightly different form from a personal account. This variant will be discussed in the next chapter.

FIGURE 5.35: Scientists for EU War and Bananas Tweet



Retrieved 4 October 2018 from Twitter. Screenshot by author.

Playfulness also relates to message framing. Mike argued that humour cannot be used in the context of positive campaigning: 'You never joke about saying something is wonderful, because then it's not a joke'. However, humour is especially effective in the context of negative campaigning—it allows negative campaigning to be delivered with positive energy, and is thus highly shareable. As Mike explained, 'Humour is good because it has the credibility of cutting something down because you see further than you are better, but at the same time it doesn't have that mean miserly spirit'. Thus, an effective strategy is to deliver a critique in a playful manner; the tweets discussed in the next chapter provide empirical evidence of this.

FIGURE 5.36: Robert's NewsThump Tweet



Retrieved 4 October 2018 from Twitter. Screenshot by author.

This strategy can, however, backfire. Figure 5.36 contains an example of humour and negative campaigning that did not succeed because people did not understand the joke. As Robert explained, 'I guess the irony doesn't necessarily become clear unless people know that NewsThump is a satirical website'. This illuminates the subjective nature of this feature and why

it is impossible to systematically assess the quality of a playful tweet. Ultimately, the subjectivity of humour best explains the inconclusive quantitative results—playfulness is an effective strategy, although it is difficult to execute well. Hypothesis 9 is supported by the qualitative data, but in light of the quantitative results, this support is weak.

5.6.3 Questions

The quantitative analysis indicates that including a question in a tweet is not an effective strategy: questions was negatively related with like odds at the sample level and negatively related with retweet odds at the population level. The interviewees who mentioned this feature did, however, believe that it was effective. Ruthi claimed that questions increase engagement because they ‘elicit reactions from people’ and are more provocative than statements. Elise commented on the fact that they arrest attention, and are therefore useful to include at the beginning of tweets.

Although Sean recognised questions as an effective strategy because they draw people in, he also observed that some questions work better than others. Rhetorical questions are not as successful as direct questions, which engage with the user. One of We Are Europe’s resonant tweets included a direct question: ‘Are you one of them?’ (Figure 5.37). This question encapsulates the campaign’s personal approach. Logan, however, declared that ‘rhetorical questions are always good’. Psychologically, these pseudo-questions trigger self-reflection: they ‘[get] the person reading the content to think about themselves instead of the person that’s telling them the stuff’. In Logan’s view, *any* type of question is personal.

FIGURE 5.37: We Are Europe’s #POSTIT Campaign Tweet



Retrieved 4 October 2018 from Twitter. Screenshot by author.

Matthew and Paul did not explicitly mention questions as a technique, but implicitly invoked the technique by incorporating it into hypothetically reworted tweets. One of

Matthew’s unsuccessful tweets about an article was an uninteresting statement. When asked whether the subject of the article or the content of the tweet contributed to its lack of resonance, he presented a more engaging version: ‘If I’d said something like “Whose side are you on? Nigel Lawson, Norman Lamont, and X or George Osborne, Gordon Brown, and Tony Blair?” ’ This version elicits an answer from the audience, which in turn might incentivise them to share or like the tweet. Paul rephrased his unsuccessful tweet about the 2012 EU investment map that was not newsworthy (Figure 5.16): ‘ “Ever wonder where your money towards the EU goes? Here’s where just some of that investment goes...” ’ This makes the tweet more relatable and attention-grabbing: ‘They see the question and they stop and think about it and then go “I do want to know more about that question” ’. Paul cautioned, however, that if questions are overused or used in a painfully obvious way, they come across as too clickbaity.

Finally, Harriet pointed out that *timing* also matters for questions: ‘It’s usually when you hit that question that everyone’s kind of thinking, but perhaps not everyone’s expressing—that’s usually the time you get a really big jump’. This relates the feature to *novelty*, being the first to express something.

On the whole, the interview data support hypothesis 8 but contradict the quantitative data. Once again, the subjectivity of style is at play: questions can be an excellent strategy, but only if executed well. The regression results indicate that most questions are not executed effectively; in light of Paul and Harriet’s comments, they might be ill-timed or sound too promotional.

5.7 Discussion

This chapter has illuminated not only the strategies that Twitter users consciously employed to increase the resonance of their tweets, but also the broader context behind each strategy: viz., under which circumstances and from whom they are most effective. This is a complement to the quantitative analysis based on the manual coding of tweets presented in the last chapter, which sacrificed local nuance in order to represent global trends. One important conclusion is that *tweets from different contexts cannot be compared*: those sent during live events and those that are more conversational, for example, represent special subcategories with unique characteristics. Moreover, *balance* is key—too much of any one feature is not effective, even if the feature is perceived to be.

Despite these qualifications, it is evident that discoverability and timing features (which comprise tweet context) were emphasised over content and style features. This reinforces the

main conclusions of the last chapter. In fact, message format was hardly mentioned. Nick T. did observe that lists were resonant, but only at the level of his articles; moreover, his example was from a pop culture context (“The five worst Kiss lyrics ever”). When asked directly about lists, Leni said that they make visuals more digestible—this is again not at the level of tweet text.

It is important to stress that the interviewees featured in this chapter are not at all a representative sample of users who tweeted during the referendum period. In fact, they are extremely *un*representative, as many are organisations, influencers, and experts. However, as the chapter has evidenced, this subset of users is important because they are deeply familiar with Twitter and were extremely motivated to tweet strategically during the referendum period. The fact that so many used war metaphors (expressed in words like *battle* and *ammunition*) reflects this. One meta-strategy is represented by this thesis in and of itself: many interviewees deliberately tested their messages in order to determine how to achieve resonance, and examined both successful and unsuccessful tweets. As Nick H. explained, ‘You think “Why didn’t that one hit a chord?” It’s helpful for you to know what concerns people’.

More shall be said on the implications of the interviewees’ tweeting strategies, and how they complement and contradict the quantitative results from the regression models, in Chapter 7. The next chapter will explore exceptional tweets and users—i.e., those that break the rules or cannot be explained—from both quantitative and qualitative perspectives.

6. Exceptional Tweets and Users

*I'm fully aware that things that resonate and become real hits are the exception to the rule,
so much so that I've wired myself for failure.'*

—Ron Perlman

There are two levels of popularity: receiving at least one retweet or like, and receiving many more. Success is defined in the regression models according to the first level, as one engagement is above the average in the population of EU referendum tweets. But as the previous chapter has highlighted, just getting one like or share is a failure for some users, particularly in the context of a campaign. Moreover, Twitter is designed to act as a megaphone—the platform's broader function is emblematised by the retweet feature itself. This chapter focuses on the second level of resonance through an examination of the most popular tweets in the dataset: of the minority of tweets that *are* retweeted or liked, what are the characteristics of those that are tremendously successful? This is where content and style features might play a pivotal role. It should be noted that extreme success is not the main focus of this thesis, even though it is associated with the research questions. It is also a matter of substantive interest to the interviewees: when speaking about tweet success, the *scale* of visibility was implicit in some definitions—the more people who see the tweet, the better. This was especially true for organisations and their representatives, and less so for individuals.

This chapter also examines exceptional tweets defined in a broader sense: those that *defy the expectations* of the models or the interviewees. Thus, both quantitative and qualitative approaches to exceptionality are explored. In the context of the models, *which organisational tweets or tweets with visual media did not do well? Which tweets without media items from non-organisations did do well?* In the context of the interviewees, which tweets were surprisingly—and inexplicably—successful or unsuccessful? The objective is to understand what characterises these exceptions and their users, and thereby contribute evidence towards a framework for a new theory of resonance on new media. As noted in the previous chapter, several of the interviewees represent exceptional users in and of themselves (e.g., Nick T., Scientists for EU).

A selection of the most resonant tweets in the population dataset are examined in Section 6.1, along with the users who sent them.⁷⁹ Section 6.2 scrutinises quantitative exceptions, as defined by influential observations. Section 6.3 is the qualitative complement to this, as it focuses on tweets that interviewees could not explain. Summary reflections are made in Section 6.4.

⁷⁹ *Note:* the screenshots of case study tweets were taken on 4 October 2018, so retweet and like counts may be slightly different than they are in the dataset.

6.1 Most Popular Tweets

What are the most retweeted tweets in the full dataset? What characterises them, and who are they from? Due to the missing data issues for like count explained in Chapter 3, the most liked tweets are not examined separately, but there is significant overlap between these two measures of engagement at the second level of popularity—nine of the ten most shared and liked tweets are the same. Because the tenth most liked tweet is also the eleventh most shared tweet, it is included in the analysis. All of the tweets were shared at least 4,378 times. Table 4.1 puts their popularity into perspective: the average retweet count at the population level is 4.22, and the standard deviation is 48.03. This means that the most popular tweets are at least *91 standard deviations* above the mean.

Moreover, given the undeniable impact that follower count has on tweet success, it is interesting to observe that only one of the exceptionally popular tweets is from one of the top ten most followed users in the dataset (Table 6.1). As almost half of the popular tweets are from users with less than 7,500 followers, and the most popular user has over 22.7 million followers, this is astounding.

TABLE 6.1: Most Followed Users in Population Dataset

SCREENNAME	MEAN FOLLOWER COUNT	MAX FOLLOWER COUNT
BBCBreaking	22,752,666	22,752,666
TheEconomist	14,546,898	14,546,898
victoriabeckham	10,117,754	10,117,754
richardbranson	7,809,547	7,817,571
hootsuite	7,696,057.50	7,714,389
jk_rowling	7,490,150.50	7,500,795
BBCNews	6,776,080.44	6,841,151
lilyallen	5,841,086	5,841,086
RitaOra	5,654,326	5,654,326
tyleroakley	5,239,935	5,256,078

Mean and maximum counts are presented because many users published multiple tweets in the dataset. The mean and maximum follower count values are identical for BBCBreaking, TheEconomist, victoriabeckham, lilyallen, and RitaOra because the dataset only has one tweet from these accounts.

Five of the most followed accounts, including the top two, are news organisations; the other five are celebrities. Recognising the disproportionate advantage that organisations and celebrities have over individuals (as revealed by the previous two chapters), the popular tweets will be grouped according to these user categories.

6.1.1 *Tweets from Individuals*

Popular tweets from individuals represent the most surprising, and thereby interesting, group. As most individuals are not as influential as organisations, and do not have the same financial and temporal resources, they are at a significant disadvantage when it comes to tweet success. In light of this, it is remarkable that four of the five most shared tweets in the dataset are from individuals. None of the accounts is verified, which indicates that the users have not achieved a level of public prominence. Their descriptions indicate that they represent a Grimsby Town Football Club supporter, a member of the Army Catering Corps, a computer science student, a sports journalist, and a music producer.

The most popular tweet is no longer available, as the user account has been deleted since the referendum. It received 14,258 retweets and 14,241 likes, and was sent at 14.42 the day before the vote. The user only had 1,822 followers at the time, but was quite prolific: he had published 11,188 tweets over the course of 3.23 years. Nothing about the tweet text helps to account for its extreme success; it does not contain any obvious rhetorical content or style features, apart from brevity: ‘#VoteLeave starter pack’. It does, however, include three images. The popularity of the tweet implies that these images were quite powerful; given its proximity to the vote, they probably contained timely content. Nothing more can be said about this tweet because its visual content cannot be viewed.

Figures 6.1-6.4 represent four of the eleven most shared tweets. They were sent on very different dates and times: 20.27 on 20 June, 10.53 on 17 May, 23.14 on 29 April, and 14.52 on 23 May, respectively. At the moment the tweets were published, the users had 2,290, 629, 20,422, and 842 followers—this is quite a range, but three of the four are significantly below the mean follower count of 4,600.39 in the population dataset (Table 4.9). It is noteworthy that the user with the least number of followers also had the least Twitter experience: he had sent 333 tweets over 3.35 years (the other users had sent between 10,248 and 25,120 tweets). Despite this, he is responsible for the fourth most resonant tweet in the dataset (Figure 6.2).


All four tweets display a wide range of rhetorical content and style strategies, and the strategies are quite similar. Figures 6.2 and 6.4 include native images that both complement and complete the tweet text; Figures 6.1 and 6.3 include multiple line breaks that extend the length of the text. As Sean noted, ‘playing around with line breaks’ is an alternative approach to visual content: ‘you’re making a visual almost with words’. This creative approach is not captured by tweet metadata and has the attention-arresting effect of an image, in that it takes up more space on the Twitter feed. Only Figure 6.2 uses other attention-arresting features: it contains two hashtags and six @mentions, all of campaign groups and prominent users.

FIGURE 6.1: 'Dear England' Tweet



Retrieved 4 October 2018 from Twitter. Screenshot by author.

FIGURE 6.2: 'Define #Brexit' Tweet



Retrieved 4 October 2018 from Twitter. Screenshot by author.

FIGURE 6.3: War and Bananas Tweet



Retrieved 4 October 2018 from Twitter. Screenshot by author.

FIGURE 6.4: Vote Leave Sign Tweet



Retrieved 4 October 2018 from Twitter. Screenshot by author.

Going beyond discoverability, the overwhelming characteristic that dominates all four of these tweets is playfulness: Figure 6.1 is sarcastic (and also personifies England); the other three are both sarcastic and humorous. Moreover, all four are examples of negative campaigning. They are from Remain supporters and castigate Leave in creative ways: a pseudo letter, a Google search revealing that Brexit is commensurate to insularity, the juxtaposition of war and bananas, and a visual sign that changes the context of a Vote Leave banner. The Google

search image is particularly creative, in that it does not reflect the actual results for ‘define Brexit’—the image had to be edited. This provides empirical support for Mike’s observation that humour is an effective vehicle for divisive rhetoric. It should be noted that the war and bananas tweet is almost identical to the most resonant Scientists for EU tweet (Figure 5.35). Given its local and global popularity, the content of this tweet must account for most of its success.

Other content and style features are present as well. Three of the four tweets contain relevant issues—healthcare and immigration (Figure 6.1), security (Figure 6.3), and economy and labour (Figure 6.4). The humorous aspect of two of the tweets is established through antithesis: the contrast between war and bananas (both semantic and visual), and between the Vote Leave banner in the photograph and its tongue-in-cheek response (‘...if you want to cut workers’ rights’). The humour in the war and bananas tweet is also enhanced through a list—the word ‘war’ running down the screen. Aversion is expressed by the ‘Dear England’ and ‘define #Brexit’ tweets. The two tweets that do not include images have clear and succinct sentences, and communicate one simple message. All of these features enhance the effectiveness of the tweets, but are not as salient as playfulness.

It is interesting to observe that the five tweets from individuals received similar amounts of retweets and likes, except for the Google search tweet (Figure 6.2). This tweet was shared over 10,000 times, but liked only 16 times. What could account for the significant discrepancy? On the one hand, the content of this tweet was clearly judged to be more shareworthy; users wanted it on their own profiles, perhaps because the image was carefully edited (and thereby novel). On the other hand, the user only had 629 followers. As likes represent private as opposed to public interaction, perhaps they are more associated with the popularity of users as opposed to the content of tweets (i.e., they might have more of a social support function). The next subsection provides evidence that corroborates this claim.

6.1.2 *Tweets from Celebrities and Established Organisations*

Three of the top eleven most popular tweets are from representatives of established organisations: Graham Norton, a BBC TV and radio presenter, Richard Branson, the founder of the Virgin Group, and Labour MP Jo Cox (Figures 6.5, 6.8, and 6.10). The sixth, seventh, and eighth most popular tweets are from celebrities: Welsh actor Taron Egerton, Gary Lineker, a former professional footballer and sports broadcaster, and Billy Bragg, a singer-songwriter and political activist (Figures 6.7, 6.9, and 6.6). All of these accounts are verified, and apart from

Jo Cox⁸⁰ have substantially more followers than the individual accounts described above—ranging from Egerton’s 221,582 to Branson’s 7,809,910 (he is the user who is also one of the top ten most followed in the dataset).

FIGURE 6.5: Graham Norton Tweet



Retrieved 4 October 2018 from Twitter. Screenshot by author.

FIGURE 6.6: Billy Bragg Tweet



Retrieved 4 October 2018 from Twitter. Screenshot by author.

⁸⁰ As will be discussed, Cox’s tweet represents a special case.

FIGURE 6.7: Taron Egerton Tweet




Retrieved 4 October 2018 from Twitter. Screenshot by author.

FIGURE 6.8: Richard Branson Tweet



Retrieved 4 October 2018 from Twitter. Screenshot by author.

FIGURE 6.9: Gary Lineker Tweet



Retrieved 4 October 2018 from Twitter. Screenshot by author.

FIGURE 6.10: Jo Cox Tweet



Retrieved 4 October 2018 from Twitter. Screenshot by author.

Celebrities and established organisations are grouped together in this analysis for two reasons. Firstly, in light of interviewee comments in the previous chapter about the dimension of *prominence*, they are more similar to each other than established organisations are to grassroots organisations, or celebrities are to non-celebrity individuals. Secondly, the line between celebrity and organisation is particularly porous with regard to these six users, as the organisational representatives are also personalities in their own right (Norton, for example, hosts an eponymous comedy chat show). It must be noted that the organisations that these individuals represent—the BBC, the Virgin Group, and the Labour Party—did not author the tweets. This distinction, overlooked in the regression models, is an important one. As Chapter 5 revealed, supporting a person is different from supporting a group, even if the person is the figurehead

of the group. Moreover, celebrity is a content feature that is positively related to engagement, and it can only be associated with individuals.

Paul and Mike affirmed that content matters far more for accounts that are prominent—their tweets are more meritocratic and less governed by random chance, because they do not have to struggle to be seen. But does this hold true in practise? Do celebrity and organisational tweets employ more rhetorical strategies than those from individuals, and are their strategies more sophisticated?

These popular tweets do not provide evidence in support of Paul and Mike's claims. To illuminate the contrast between tweets from individuals and tweets from celebrities or organisations, they can be grouped into two categories. The first, represented by Figures 6.5 and 6.6, reflects the strategies discussed in the previous section: playfulness conjoined with negative campaigning. Both are anti-Farage, but Bragg's tweet has a stronger and more creative expression that utilises wordplay and antithesis: 'Not every Leave voter is a racist, but every racist will vote Leave'. Like two of the tweets from individuals, it includes a striking image to complete and complement the text—the text, in fact, explicitly references the image. But unlike the images in the tweets from individuals, it is not original. The theme of immigration is also central to the tweet, and is expressed by the word 'racist' as well as the infamous 'Breaking Point' poster in the image.

Although Norton's tweet employs the same strategies of playfulness and negative framing, it is arguably less well-executed than Bragg's tweet or any of the individual tweets. The tweet is quite wordy and represents a run-on sentence, and there is no punctuation. Moreover, there is no visual content—not even from text broken up into multiple lines. However, it is the third most shared tweet in the dataset, and the most liked: it received 11,798 retweets and 18,107 likes. In light of the fact that Norton is a famous TV personality (he had 1,203,956 followers at the time of the tweet), this is not surprising. As suggested by the Google search tweet in the previous subsection, likes are associated with social support, and in particular fan support. In Mike's words, 'You know what gets very few retweets and a shedload of likes? Celebrities talking about their kitten sneezing. A few people will retweet it, because they're weirdos, but a lot of their fanbase would just go "oh isn't that sweet" and then like'. If this holds true, the ratio of retweets to likes for celebrity tweets can serve as a litmus test for the quality of their content.

The second, and larger, group of celebrity and organisational tweets is represented by Figures 6.7-6.10. These tweets exhibit quite a different rhetorical strategy from those in Figures 6.1-6.6: they are all characterised by enthusiasm and positive campaigning. Egerton and Cox's tweets are playful, but the other two are not—they are straightforward and emphatic

expressions of support for Remain. The humour in Egerton and Cox's tweets is enhanced by personal photographs, and Branson's endorsement of Remain is amplified by a short video of him reflecting upon how the EU has benefited the Virgin Group. It is important to observe that the images in the individual tweets are not of the individuals themselves—as they are not famous, personal photographs might not have been as effective for them (unless there was something unique or novel about the photograph).

Timing is another factor that helps to explain the success of these tweets. Egerton's was sent the day before the voter registration deadline (and hence also contains resolution); Lineker's was sent the day before the vote; Branson's was sent two days before the vote. Because they are influential users, they can take advantage of timeliness without the danger of being drowned in a sea of other voices: they are heard above everyone else. There are no consistent time of day patterns; the tweets were sent at various hours in the morning, afternoon, and night (with the exception of small hours).

All four tweets received more likes than retweets. The disparity for Egerton and Lineker is especially large, as they received almost *twice* as many likes as shares: 15,774 vs. 8,796 and 15,092 vs. 8,672, respectively. This is further support for the fan function of likes, and indicates that the tweet content could be improved. Indeed, Lineker's tweet is not really saying anything: it alludes to 'reasons to remain' on a high level, is repetitive, and lacks visual media.

Cox's tweet is a special case because she was killed the day after she sent it. As mentioned in the previous chapter, this was a significant event in the run-up to the referendum that halted campaigning from both sides. It was, as a matter of fact, the last original tweet she sent before her devastating death. As evidenced by the date of the comments, much of the engagement occurred on and after the 16th, as opposed to on the 15th, when the tweet was sent. Thus, the engagement was driven primarily by a desire to show support for Cox, and is less connected to the inherent appeal of the tweet. It is no coincidence that the twelfth most shared tweet in the dataset is also from Cox.

6.2 Outliers (Quantitative Exceptions)

The popular tweets from individuals represent one category of outliers: as the regression models have shown, tweets from organisations have a 684.7% greater odds of being shared and a 473.4% greater odds of being liked. The models also reveal that visual media is the second most influential variable for tweet success: in the sample-level regressions, including visual media increases retweet odds by 116.6% and like odds by 79.7%. Differences between different media

item categories are illuminated by the population-level models. Video is undeniably the most resonant media item: including a video increases retweet odds by 162.6%, whereas including images and GIFs increase retweet odds by 64.8% and 25.9%, respectively.

This section examines other examples of outliers that defy the trends summarised in the models—not only surprisingly successful tweets, but also surprisingly *unsuccessful* tweets. Which tweets from organisations or with videos did *not* resonate? Which tweets from individuals without media items *did* resonate? What characteristics, if any, do they share—to what extent can they be explained? To capture the second level of resonance, the outliers are calculated from a multivariate linear regression at the population level, which maintains retweet count as a continuous variable (Appendix H).

Because it combines leverage and residuals, Cook's D is the best measure of influence on model fit. It was thus used to identify the quantitatively counterintuitive tweets discussed in this section: as they had the largest divergence from model trends, they also had a large influence on its fit. It is important to note, however, that tweets with high influence can also significantly *reinforce* model trends (many include videos, for example). Thus, the tweets with the greatest Cook's D values were manually inspected, and only those that diverged from the model were selected for this section.

6.2.1 *Surprisingly Successful Tweets*

The first category of influential tweets is surprisingly successful—these tweets are from users who not only do not represent organisations or celebrities, but also have extremely low follower counts; three are represented by Figures 6.11-6.13. The sender of the tweet in Figure 6.11 had seven followers, the sender of the tweet in Figure 6.12 had six, and the sender of the tweet in Figure 6.13 had an astounding zero at the time they were published.⁸¹ They received four retweets and three likes, 25 retweets and 27 likes, and six retweets and one like, respectively. Given the extremely low chance that these tweets were even seen, what could account for their resonance?

⁸¹ There are no other examples of tweets from these users in the dataset.

FIGURE 6.11: Vote Leave £350 Million Tweet



Retrieved 4 October 2018 from Twitter. Screenshot by author.

FIGURE 6.12: Vote Leave Pledge Tweet



Retrieved 4 October 2018 from Twitter. Screenshot by author.

FIGURE 6.13: Imagine If We #VoteLeave Tweet



Retrieved 4 October 2018 from Twitter. Screenshot by author.

The feature that all three tweets ostensibly have in common is their pro-Leave stance: they contain the hashtag #VoteLeave, and two of the three include images from the Vote Leave official account. Thus, the visuals are not original, and the text of all three tweets also lacks originality—the tweet with no image expresses the popular £350 million a week argument that was one of Vote Leave’s official taglines, the tweet with the pledge graphic copies the text of the graphic, and the tweet with the passport image has an extremely high-level expression of positivity. This lack of novelty makes the success of these tweets even more extraordinary.

The most rational explanation is that because the tweets included the #VoteLeave hashtag, they were seen by a pro-Leave user who happened to be following the hashtag stream at the time they were sent—or automatically shared by a bot. As illuminated by the previous chapter, the Leave network on Twitter during the referendum period was far stronger than the Remain network. In fact, a fourth surprisingly resonant tweet is from the bot account @iVoteLeave: ‘Our message is loud and clear: Britain is #BetterOffOut and we must #VoteLeave and #TakeControl! Join us to spread this message!’ This is as generic and unoriginal as the other three tweets (it is also used as the account’s description), and it does not come with a visual item. Moreover, @iVoteLeave had zero followers at the time it sent the tweet: 12.39am on the 23rd of April. This falls into the small hours category, which is when tweets are least likely to be seen. Despite all of these negative factors, the tweet received 525 shares, and is thus the most surprisingly resonant of all. A certain degree of automation or pre-planning with regard to retweets must have occurred, leveraging the pro-Leave hashtags for discoverability.

6.2.2 *Surprisingly Unsuccessful Tweets*

The second category of influential tweets is surprisingly unsuccessful—most come from prominent organisations and include visual media, the two factors with the strongest relation with engagement in the regression models. Three with images are presented in Figures 6.14, 6.15, and 6.16, respectfully: they were sent from @dabeard, an Executive Editor at Public Radio International (who was verified during the referendum), freelance journalist @shehnazkhan (who since the referendum has been verified), and @ceffylglasgafr, a Welsh user. Their accounts had 18,286, 5,890, and 3,732 followers at the time their tweets were sent. Like many of the examples discussed in the previous section, all three tweets contain images and playful language. Given these extremely favourable factors, it is remarkable that @dabeard’s tweet received no engagement, @shehnazkhan’s received two likes, and @ceffylglasgafr’s received one like. What could have happened?

FIGURE 6.14: 'White Guys in Ties' Tweet



Retrieved 4 October 2018 from Twitter. Screenshot by author.

FIGURE 6.15: 'Steal Your Job' Tweet



Retrieved 4 October 2018 from Twitter. Screenshot by author.

FIGURE 6.16: 'English at Play' Tweet



Retrieved 4 October 2018 from Twitter. Screenshot by author.

Closer scrutiny of tweet content is necessary to answer this question. The phrase 'going beyond the white guys in ties' relates to the article that is shared in @dabeard's tweet, which coincidentally is about the founders of the We Are Europe campaign. The picture in the tweet, however, is of a reporter at a desk and does not seem to have any overt connection with the campaign; moreover, the reporter is not a recognisable figure. Apart from this disjunct, it is important to note that the user is not British. Judging from his timezone and user description, he is based on the East Coast of the United States. This explains the fact that his tweet was sent at 2:14am BST, which is equivalent to 9.14pm EDT. His American followers would have seen it, but the EU referendum might not have been a particularly resonant topic in the United States in May 2016—the more relevant political matter was the presidential election in November.

@shehnazkhan's tweet contains a photograph of a man holding up a sign with small text, which needs to be read for the tweet to be understood. The sign is difficult to read, especially on mobile devices; zooming in is necessary. Moreover, even if the sign were read, its humour is extremely dark: 'An immigrant with a degree is not coming over here to "steal your job"... when all you have is a standard grade in P.E. and an S.T.I.' Sarcasm does have its limits, and the implication that Leave supporters are not only uneducated, but also have sexually transmitted infections, is deeply uncomfortable. The fact that there are policemen in the picture deflates the humour even further.

Similarly, the humour in @ceffylglasgafr's tweet is also too uncomfortable. It arises from the contrast between the tweet text and the picture, but the contrast is too extreme: the

phrase ‘The English at play’ is being stretched to its limits by what appears to be an image from a terrorist attack, with people and police running on a chaotic street. It is difficult to joke about terrorism, especially in light of the attacks across Europe in the years leading up to the referendum. Thus, both this tweet and @shehnazkhan’s are examples of playfulness gone too far—not many users are keen to share content that is unduly offensive or belittles extremely sensitive topics.

The fourth and fifth examples of surprisingly unsuccessful tweets are from the same user, Welsh Labour and Co-operative Party MP Stephen Doughty (Figures 6.17 and 6.18). Although Doughty had 10,079 followers when these tweets were sent, and although they both contain videos, neither received any shares. The reason for their lack of resonance becomes clear when the greater context of the tweets is considered: they are two from a series of seven sent between 8.38pm and 8.56pm on 10 June. All seven tweets contain hand-recorded video clips of interviews with local Labour Party members and Remain supporters. Each of the other five videos received at least three retweets; the first received seven and the last 12. These two were the third (Figure 6.17) and the fifth (Figure 6.18) that were sent. The most important difference between them and the resonant tweets is not video quality, but length: the videos in the resonant tweets are about 10-15 seconds each, allowing the interviewees time to state their main reason for voting Remain. By way of contrast, these tweets are three and four seconds, during which the interviewees could only quickly declare their name, political party, and support of Remain. There is literally no substantive message.

This example enriches the conclusions from the previous section by demonstrating that even for influential users, content is important—while standards are lower than they are for non-influential users (especially for receiving likes), public figures and organisations cannot assume that everything they tweet will be shared, even if it includes a video. A minimum threshold of quality has to be reached, and the tweets in Figures 6.17 and 6.18 did not meet this threshold.

FIGURE 6.17: Welsh Labour Vote Remain Video Tweet I



Retrieved 4 October 2018 from Twitter. Screenshot by author.

FIGURE 6.18: Welsh Labour Vote Remain Video Tweet II



Retrieved 4 October 2018 from Twitter. Screenshot by author.

6.3 Inexplicable Tweets (Qualitative Exceptions)

Despite all of the discoverability, content, and style features explored in this thesis, and despite the social media and campaigning expertise of the interviewees, the success or failure of tweets can be inexplicable. Some interviewees recognised the inherent unpredictability of success on Twitter. As Ruthi straightforwardly declared, ‘I think a lot of it is just random’. Observing that sometimes a tweet that resonates might not resonate when tweeted on another occasion (or vice versa), Sean emphasised the importance of persistence: ‘Sometimes it’s a shot in the dark and you’re trying your best until it sticks, until it catches fire. You can do it perfectly and it won’t get any retweets; you can do it really badly and it can still work quite well... Sometimes you need to just keep chucking at the wall, and eventually one of the tweets will stick’. This is especially true for the second level of popularity, even though Sean also noted that extreme success can be orchestrated via influencer outreach, as demonstrated by We Are Europe’s most resonant tweet (Figure 5.1). In his opinion, to attain this level of success ‘you have to catch a wave and get lucky as well’. One cannot ‘make a viral video’, but one can ‘make a really good video and hope it goes viral’.

Indeed, as noted in Chapter 5, ‘message testing’ via tweeting the same content with variation is a common strategy. It is especially common in journalism: Robert, Katie, Alberto, Sam, and Nick T. all tweet multiple links to their articles, with different angles or quotes each time; editorial practise at InFacts is to schedule three tweets for every article. As mentioned in the previous chapter, Nick T. increases variation in his tweet text if the blog post is intrinsically unpopular—sometimes a more enticing bent can help attract people to the post. Despite this message testing, however, it is not always possible to determine why some tweets resonate and others do not. Although the general consensus is that context features (namely timing) account for the success of imperfect tweets and the failure of perfect tweets, sometimes the ‘right’ time (e.g., 5pm) does not work, or the ‘wrong’ time (e.g., 11pm) does work.

The counterintuitive nature of tweet success characterises Nick T.’s Twitter experience. In his view, his most unique and insightful tweets receive two or three shares, whereas his ‘completely puerile’ tweets receive 100. As he remarked, ‘It’s slightly depressing in that sense... The most retweeted and engaged tweets are not the ones that I think “Oh I really got to the bottom of something there”’. The parody account named Paddy Ashdown’s hat mentioned in Chapter 5 serves as an excellent case study of resonant puerility. This puerile factor relates to the playfulness feature and one of Elise’s points: ‘All the research out there about how you get messages across says that you’ve got to appeal to a *Daily Mail* reader who’s about twelve years

old'. Yet puerility is subjective and not necessarily appropriate in a referendum context, and twelve year olds are difficult to please.

This section complements the previous section by focusing on interviewee tweets that were inexplicably successful or unsuccessful. The interviewees expected these tweets to be more or less resonant than they actually were, and could not explain the unexpected resonance or lack thereof. Viz., from the perspective of their senders, the success or failure of these tweets could not be rationally explained. From an external perspective, however, could there be a rational explanation?

6.3.1 *Unexpectedly Successful Tweets*

Given the fact that some of the interviewees knew each other, it is possible to begin with an example of a successful tweet that could not be explained from the sender's perspective, but *could* be explained from another interviewee's perspective (Figure 6.19). This was Andrew's second most resonant tweet in the dataset, as it received 115 retweets and 102 likes (his average received 14 shares and likes). Unlike his other popular tweets, it does not include an image or a call to action—it is a simple statement that he has been campaigning in Hull, where Labour support for Brexit is 'very strong indeed'. It is also a quote of Brendan's tweet, which is itself a quote of Conservative MP Steve Baker's tweet, which is a quote of a Sky News tweet. Thus, to obtain the full context many clickthroughs are necessary—and clickthroughs by definition lead people away from the original tweet, decreasing the chances of engagement. Andrew simply could not understand what made the tweet so resonant: 'Well it's Friday the 10th of June 2016, at 5.27pm. Does that make any difference? I wouldn't have thought so... I haven't a clue. I am beaten by that one, I can assure you... That should have been lucky to have a retweet or two'.

FIGURE 6.19: Andrew's Labour Leave Campaign Tweet



Retrieved 4 October 2018 from Twitter. Screenshot by author.

FIGURE 6.20: Brendan's Labour Leave Campaign Tweet



Retrieved 4 October 2018 from Twitter. Screenshot by author.

As Brendan was another interviewee, and as his tweet that Andrew quoted was also in the dataset, his reflections on its resonance help explain the success of Andrew's tweet (Figure 6.20). Coincidentally, this was Brendan's second most popular tweet—it received 61 shares and 43 likes (his averages were five shares and six likes). He reinforces the point made in the tweet that he quotes from Steve Baker: the BBC has 'totally ignored' the fact that some members of the Labour Party are pro-Brexit. Baker's tweet references a senior Labour MP, John Mann, who is one such member. According to Brendan, the fact that the BBC has chosen not to report on Mann's decision to support Leave is shocking, given Mann's prominence and the size and duration of the Labour Leave campaign. They are only presenting a 'blue on blue' battle: i.e., pro-Leave vs. pro-Remain Conservatives (blue is the colour of the Conservative Party). Brendan's tweet was resonant because it focussed on an overlooked party: 'Certainly Labour was not heard of on either side of the campaign'. As discussed in the previous chapter, this relates to the contextual feature of *counterintuitive novelty*. Labour support was important for Leave's victory, and many people shared the tweet because they 'realised the importance of needing to win the Labour vote', and the importance of shattering the myth that 'not all Tories are Brexiteers—anyone can be a Brexiteer'. Because the Labour Party was not associated with Leave in the mainstream media, the visibility afforded by social media for Labour Leave was practically important.

Thus, it can be concluded that Andrew's tweet, albeit expressed less forcefully than Brendan's, was resonant because it highlighted the strength of the Labour Leave vote. As the success of the Labour Leave campaign demonstrates, pro-Brexit members of the Labour Party are quite active on Twitter (and used the platform to build community). Interestingly, even when asked directly about the Labour Leave theme, Andrew himself did not recognise this.

On a similar note, the Better Off Out tweets about the UK fishing industry were also resonant because of their theme; they were the third and fourth most successful campaign tweets in the dataset. Fishing for Leave was an extremely popular grassroots campaign, which suggests that workers in the fishing industry were active on Twitter during the referendum. Although one of the tweets was well-designed, with a direct call to action and a quirky cartoon graphic (Figure 6.21), the other was poorly designed, with a descriptive statement and an unattractive map that does not effectively illuminate the statement (Figure 6.22). Andrew had difficulty explaining its resonance, and ascribed it to luck.⁸² In his opinion, the map is aesthetically repulsive: ‘It’s not attractive; it’s not standing out at me. What the hell’s this? It looks like it’s got a handle that you pull. The font, the font colour is just... it’s sort of an ill-looking colour’. Moreover, the message is quite bland. Given the unattractive image and text, the fishing theme is the only feature that might account for this tweet’s unexpected resonance. As Andrew noted, ‘common fisherman’s policy has been a big thing’ for those in the fishing industry. On the whole, however, he maintained that the tweet should not have been as successful as it was.

FIGURE 6.21: Better Off Out’s Fishing for Leave Tweet (Cartoon)



Retrieved 4 October 2018 from Twitter. Screenshot by author.

⁸² It should be noted that Andrew did not send this particular tweet, even though he contributed to the account.

FIGURE 6.22: Better Off Out's Fishing for Leave Tweet (Map)




Retrieved 4 October 2018 from Twitter. Screenshot by author.

A third example of an inexplicably successful tweet is from the We Are Europe campaign (Figure 6.23). With 309 shares and 115 likes, this tweet is their third most resonant, but it does not align with their creative and edgy voice. It contains a *Financial Times* bar graph that visualises the UK's misperceptions of the EU, and its message has three clauses: 'This is important', '@FT has busted those #brexit myths', 'everyone share this far & wide'. Sean struggled to identify the features that account for its success, as he believes it is 'badly written' (normally We Are Europe tweets have fewer clauses), and in house visuals are better than those from external sources.⁸³ He surmised that it was inspired by the resonant rebuttal tweets with graphs they had sent during live debates.⁸⁴ Unlike the rebuttals, however, this tweet does not have a sardonic or playful tone. In Sean's opinion, despite its high levels of engagement, it represents 'one of the weakest things' that We Are Europe did with regard to their objective—inspiring young people to vote Remain. He does not believe that its success can be replicated: 'Sometimes things just go really well, and it just happens, and then you can't recreate it. I'm pretty sure that if we tried this again it wouldn't work'.

⁸³ Sean did not personally author this tweet, but he reviewed and approved of its content.

⁸⁴ The most resonant of these tweets (Figure 5.28) was actually published two days after this one, on 15 June.

FIGURE 6.23: We Are Europe's *Financial Times* Graphic Tweet



Retrieved 4 October 2018 from Twitter. Screenshot by author.

All of this being said, Sean observed that one particular user type on Twitter responds favourably to graphs: activists. He hypothesised that many of the accounts that shared the tweet are probably of this type, and one probably had a very large number of followers. In light of the style features reviewed in the last chapter, it can also be argued that the opening of the tweet, ‘This is important’, might have contributed to its success. Although it sounds overly simplistic and obvious, it is also attention-arresting due to its clarity and urgency. People on Twitter like to be told what to focus on, given the vast volume of content on their feed. The ‘this is’ clause is reminiscent of Alberto’s second most resonant tweet (Figure 5.24).

The fourth and final inexplicably successful tweet is perhaps the most inexplicable of them all. YEM published two register to vote tweets with the same native campaign graphic and a link to the same Facebook video; the tweets also share an enthusiastic tone. One was published at 2.15pm on 28 May (Figure 6.24), the other at 5.46pm on 21 May (Figure 6.25). The former received 1,941 impressions, 11 retweets, six likes, and five link clicks; the latter received 210 impressions, no public engagements, and two link clicks. From Logan’s perspective, the successful tweet should not have been, because it ‘has a lot of blue text’, ‘it was posted [a week] afterwards’, and ‘2pm is not a good time’. The success of this tweet is difficult to rationalise given its similarity to another tweet with more favourable features that failed. Logan could not even attempt to provide an explanation: ‘I can only say it’s probably down to random chance.

It flummoxes me'. The more resonant tweet does, however, have two favourable features. Although it is not as novel as the earlier one, it is more timely because it was sent closer to the voter registration deadline. Accordingly, it is also expressed with more overt enthusiasm ('TODAY!!!').

FIGURE 6.24: YEM Register to Vote Tweet (Successful)



Retrieved 4 October 2018 from Twitter. Screenshot by author.

FIGURE 6.25: YEM Register to Vote Tweet (Unsuccessful)



Retrieved 4 October 2018 from Twitter. Screenshot by author.

6.3.2 *Unexpectedly Unsuccessful Tweets*

Just as the success of tweets can be inexplicable from the sender's point of view, the failure of tweets can also be inexplicable. Like the previous subsection, this subsection begins with an example that two interviewees discussed: what is flummoxing to the sender might not be to a knowledgeable observer. Ruthi published an Oxfordshire Greens tweet that linked to a popular Monty Python Brexit parody video, and it received no engagement (Figure 6.26). Given the effectiveness of humour on Twitter, she had expected a better performance, and could only guess that it failed due to a lack of novelty: 'I think a lot of it is just random. Somebody tweeted something similar, so our followers retweeted that'.

Elise, however, could provide a rational explanation grounded in language (i.e., rhetorical content). The link leads to a Monty Python video, but the words 'Monty Python' are not in the tweet. As this was a 'hugely popular' television programme that many people in Britain grew up with, explicitly including its title would have 'appeal[ed] to a certain demographic'. Moreover, it would have appealed to Leave voters in particular, as many of them 'were of the Monty Python generation'; in Elise's view (but not Ruthi's) this was the Oxfordshire Greens target audience. To further enhance the tweet, Elise affirmed that 'Monty Python' should have been at the very beginning, in order to 'hook people': 'Monty Python sums up the EU...' She also noted that there is no need to explicitly indicate that the video is funny, because those familiar with Monty Python *know* that it is funny. In this case, unlike that of the We Are Europe tweet in the previous subsection, telling people what to think might not be the best strategy—it is too heavy-handed.

FIGURE 6.26: Oxfordshire Greens' Monty Python Tweet



Retrieved 4 October 2018 from Twitter. Screenshot by author.

The second inexplicably unsuccessful tweet is also playful, and is another quote tweet from Andrew (Figure 6.27). It received no engagements, which is surprising in light of its humour, the fact that it mentions a popular political figure, and the 'good photograph' in the quoted tweet. This photograph is only visible as a thumbnail in Andrew's tweet, so perhaps the

tweet would have been more successful if he had tweeted the photo directly (but this would be a copyright violation). Interestingly, however, the quoted tweet only received one like and one retweet from the same user, despite having the full photograph and equally playful text. In the context of this chapter, one like and share might not seem like much, but it does represent the first level of success defined by this thesis.

FIGURE 6.27: Andrew's Steve Turner Quote Tweet



Retrieved 4 October 2018 from Twitter. Screenshot by author.

The third and final inexplicably unsuccessful tweet is from Robert, and shares a link to his most popular piece for the *Conversation*, which at the time happened to be ‘the most shared *Conversation* article on social media’. He tweeted multiple times about the piece, using a different pull quote on each occasion. The quotes represent in Robert’s view ‘the most surprising and strongest arguments’ in the article, which are its ‘key argument[s] against Lexit’. While one of these quotes was remarkably resonant (Figure 6.28), another was not (Figure 6.29). According to Robert, the less successful one is ‘as shareable as the other one’ because it contains a surprising and thereby attention-grabbing quote: ‘For those in favour of Lexit it would be very provocative’. The fact that it was less successful is ‘a complete coincidence’—at that point, he only had 1,199 followers. The chances of any of his tweets being ‘completely unnoticed are quite high’; all it takes is one share to ‘massively [enlarge] the potential audience of people who could retweet’.

Robert’s comments reinforce the importance of the first level of success: whether a tweet is shared at all. It is interesting to examine possible reasons why the tweet did not receive even one share—all of Robert’s tweets linking to this article were handicapped by the fact that he only had 1,199 followers at the time he sent them, and yet some *did* receive shares. In light of the interviewee insights from the last chapter, two hypotheses can be made about the failure of this tweet. The first (and more obvious) relates to timing: it was sent only one minute later than the successful tweet about the same article, on the same day (12.21pm vs. 12.20pm on 19

June). Thus, Robert's followers who were on Twitter almost certainly saw the earlier tweet first; it was more novel. The second hypothesis relates to the content of the tweets—the successful tweet is arguably clearer and simpler, as it contains concrete nouns: *protections*, *workers*, *consumers*, *environment*. The quote in the unsuccessful tweet contains more high-level and conceptual nouns: *integration*, *corrective*, *neoliberalism*. It is thus less relatable and semantically transparent, especially in light of Twitter audiences and practises. While Robert's followers might not be 12-year-olds who read the *Daily Mail*, it is undeniable that the successful tweet is easier to digest and understand.

FIGURE 6.28: Robert's Lexit Article Tweet (Successful)



Retrieved 4 October 2018 from Twitter. Screenshot by author.

FIGURE 6.29: Robert's Lexit Article Tweet (Unsuccessful)



Retrieved 4 October 2018 from Twitter. Screenshot by author.

6.4 Discussion

This chapter has explored various categories of exceptional tweets and users: those that achieved extreme success, and those that defied statistical and interviewee expectations with regard to both success and failure. It is clear that although organisations have a significant advantage over individuals for the first level of success (achieving at least one like or share), at the second level individuals can be just as, if not even more, successful if they publish exceptional content. Four of the five most shared tweets in the dataset are from non-influential users. Thus, these individuals were even more visible than celebrities and organisations.

What made their tweets exceptional? Although half contain images, the most important feature is their usage of creative humour to deliver divisive rhetoric. In the context of a contested political campaign, this strategy is incredibly resonant. The popular tweets from celebrities and organisations exhibit more positive campaigning and are not as creative (several are not playful at all). For these influential users, being *personal* is just as effective as being clever or original, because the public is interested in them as individuals. This is why a distinction must be drawn between organisational accounts and the accounts of their figureheads—they both have an advantage, but when it comes to the second level of success, the figureheads have even more of an advantage. What is interesting is that in this dataset, playfully divisive tweets from non-influential users are more resonant than either personal expressions of positivity or playful expressions of negativity from influential users. This underscores the importance of content for extreme popularity.

In general, when the most resonant tweets are taken into consideration, it is clear that @mentions, hyperlinks, and multiple hashtags are not important, emotiveness trumps fact, time of day does not really matter (as long as the tweet is not sent during small hours), and video is not necessary, while images are helpful. Interestingly, although the most popular tweet is pro-Leave, the ten others are pro-Remain.

Support of Leave does, however, characterise the tweets that were surprisingly successful in the context of the models. These tweets were sent from users with extremely low (or zero) follower counts, and were retweeted at least several times (in one case 525 times). This suggests that Leave had an extensive Twitter network that leveraged hashtags to disseminate content. Due to this network and automation techniques, unoriginal tweets from unpopular users that included #VoteLeave were not only noticed, but also shared. Bot accounts played an active role on Twitter during the referendum period, as the algorithmic amplification of content on social media has become an indispensable political campaigning strategy (Woolley & Howard, 2018).

The surprisingly unsuccessful tweets in the context of the models illustrate that playfulness has its limits: it can backfire if it is insensitive or too extreme. Moreover, although organisations have an advantage over individuals, they still have to maintain a minimum level of quality for their tweets to be shared—but not necessarily for their tweets to be liked. Likes appear to serve more of a social support function than retweets, which are more related to tweet content. The followers of influential users might support them, but do not blindly share everything that they tweet with their own followers (because they are presumably unimpressed by the content).

The surprisingly successful and unsuccessful tweets from the interviewees' perspectives reveal that while it is easy—indeed instinctive—to ascribe resonance or failure to random chance or timing, an explanation can usually be constructed. As users are biased about what they publish, it is often helpful to obtain insights from another knowledgeable user. This external perspective allows for more context, and the best example of the benefit of added context is Andrew's tweet (Figure 6.19). Without the knowledge from Brendan that Labour Leave is a resonant topic on Twitter because many members of the community are on the platform, its success does not make any sense. Resonance, therefore, is a function of both luck and skill (represented in this case by tweet content)—neither can be zero. Chance helps a tweet obtain more visibility, but without inherently appealing content, it will not obtain any visibility at all.

Taking all of these insights into consideration alongside those from the previous two chapters, the next chapter revisits the theoretical framework presented in Chapter 2. Through the integration of these insights, the key factors that influence tweet resonance are established, and nuanced distinctions between the two measures of resonance are made. In doing so, the foundation for a new theory of resonance on new media is put forward.

7. Discussion and Conclusion

'Societies have always been shaped more by the nature of the media by which men communicate than by the content of the communication.'

—Marshall McLuhan

This thesis has explored the resonance of persuasive tweets through a multi-strategy approach: content coding (both manual and automatic) followed by regressions, semi-structured interviews, and close analyses of exceptional tweets and users. The conclusions drawn from each strategy are integrated in this chapter. No single method can explain resonance on its own, as it is a highly complex concept. Moreover, given the volume of data collected, both macro and micro-level investigations are necessary for the fullest insights.

The regression models have provided an effective overview of global trends. The overview is strengthened by the fact that the sample-level regressions were replicated at the level of the population from which the sample was drawn, due to the efficacy of supervised text classification algorithms. The population-level results reaffirm the sample-level results, and illuminate the differences between visual media types. As comprehensive as the models are, they are necessarily constrained because they present a static, one-size-fits-all summary of the characteristics associated with successful tweets, averaged across all user and tweet types. Thus, they reveal more about Twitter's algorithm and offline hierarchies than they do about tweet content and meaningful differences in this content. The regressions are also constrained because what can be measured from Twitter data and metadata is extremely limited. Indeed, the low pseudo- R^2 values of even the full models indicate that substantial explanatory features for tweet success are missing.

The interview data served as an invaluable complement to the quantitative analysis. It illuminated not only the strategies that Twitter users deliberately implemented to increase tweet success, but also the local context of each strategy: under which circumstances and from whom it is most effective. What is important for tweet resonance is not which strategy to adopt, but rather *when* and *how* to adopt it; the presence of a strategy—even as intuitively appealing as strong emotion—means little in and of itself. Many of the strategies overlap with the features included in the models, but strategies representing features that were not included were also mentioned; they pertain to context, content, and style, and their detection requires both human judgment and information external to Twitter. Thus, they are beyond the scope of the quantitative models designed for this thesis, but they indicate ways in which the fit of the models can be improved.

The third results chapter synthesised quantitative and qualitative data in its examination of exceptions to the model trends and user strategies, as well as tweets that attained the second level of success: extreme popularity. Given the diverse motivations and practises of users on the platform, tweets are quite idiosyncratic—even the most poorly designed messages might succeed, or the most carefully crafted messages might fail, under certain conditions. Exceptional popularity is by definition impossible to predict, but patterns can be detected among the most resonant tweets. Understanding these conditions and patterns exposes the possibilities and limits of the factors that contribute to resonance on Twitter. In particular, they reveal how influential users or tweets with visuals can fail, and how non-influential users or tweets without visuals can succeed.

In delineating the factors that influence tweet success, Section 7.1 will reassess the interdisciplinary theoretical framework presented in Chapter 2 (which is premised upon discoverability, content, and style). All of the hypotheses are summarised in Table 7.1. They will be reviewed in light of both quantitative and qualitative findings, and factors that emerged from the qualitative analysis that expand the framework will be discussed. To conclude this section, a revised schema for the resonance of digital cultural objects that is derived from Griswold's (2013) cultural diamond will be presented. Based upon these conclusions and the results of the outlier analysis, basic strategies for how tweets from 'ordinary' users can achieve success will be proposed in Section 7.2. The operationalisation of success will be revisited in Section 7.3—differences between the two measures used in this thesis will be discussed, as well as their limitations and other potential measures. Broad comments will be made upon the extent to which Twitter mirrored or impacted the actual results of the referendum in Section 7.4, along with other implications of this thesis for the intersection of politics and marketing and platform design. Section 7.5 discusses a variety of further directions that can be taken to enhance the findings, and final reflections are provided in Section 7.6.

TABLE 7.1: Hypotheses

1a	If hashtags are included in a tweet, it will be more successful.
1b	If user mentions are included in a tweet, it will be more successful.
1c	If a tweet contains visual media, it will be more successful.
1d	If a tweet contains hyperlinks, it will be more successful.
2	If a tweet is from an organisation, it will be more successful.
3a	If a tweet contains identification rhetoric, it will be more successful.
3b	If a tweet contains division rhetoric, it will be more successful.
4a	If a tweet contains enthusiasm, it will be more successful.
4b	If a tweet contains anxiety, it will be more successful.
4c	If a tweet contains aversion, it will be more successful.

5	If a tweet contains relevant issues, it will be more successful.
6	If a tweet contains resolution, it will be more successful.
7a	If a tweet contains a list, it will be more successful.
7b	If a tweet contains a contrast, it will be more successful.
8	If a tweet contains questions, it will be more successful.
9	If a tweet is playful, it will be more successful.

7.1 Factors that Influence Tweet Resonance

Given the analyses in the previous three chapters, the main research questions of this thesis can now be addressed. *What makes a message resonate with its audience on social media? How can the message be expressed and delivered most effectively? What strategies do individuals and organisations use when they send a tweet and why?*

The overarching conclusion from both the regression models and the interviews is that in the fast-paced and high-volume landscape of social media, context is more important for resonance than content or style. Most content and style features are not significant in the models; indeed, the two that are have a *negative* impact on tweet success. Although the interview data reveal that content and style are important in certain contexts, discoverability features were more emphasised overall. Increasing the likelihood that a message is *seen* matters more than *what* is said or *how* it is said, because the vast majority of tweets are not seen (regardless of their quality). Schudson’s (1989) concept of *retrievability*, which is one of his dimensions of the potency of a cultural object, has thus been affirmed in the landscape of new media.

The framework of the attention economy is extremely relevant for understanding resonance on Twitter (Wu, 2016b). Multiple interviewees described the purpose of their strategies as arresting attention, and recognised both the platform’s underlying business model and user interface. Including visual media is the most effective approach because it literally takes up more space on the Home feed. It is also algorithmically favoured to appear more often than text, as several interviewees with digital marketing expertise noted. The population-level regressions reveal that video is the media item with the strongest relation to the likelihood of a retweet, and this aligns with Paul’s comment that it accounts for the largest proportion of advertising revenue (video advertisements are more profitable than image or text advertisements). To capitalise on more of this revenue, social media platforms heavily favour video content, and have created the infrastructure to behave like ‘conversational television channels’. Thus, on a conceptual level, this thesis has demonstrated that the rhetoric of platform design—which is partly driven by financial incentives—is more powerful than the rhetoric of content. User interfaces are persuasive, and content cannot be viewed in a vacuum.

An even more potent attention-arresting factor that is not within user control is the tweet source: tweets from organisations and their representatives have a significant advantage over tweets from non-influential individuals. Organisation was by far the variable with the strongest impact on tweet success in all of the regressions. The interviews brought out an important distinction between organisations and individuals that affirms the advantage that the former have: organisations invest more time and financial resources into Twitter, and are more concerned with tweet success.

A third factor that contributes to both discoverability and arresting attention is timing, and it is within user control. It is represented in the models by three empirically measurable features: time of day, day of the week, and number of days before the referendum (all three served as control variables). Only the third of these features represents an aspect of a broader concept that the interviewees referred to as timeliness, which is strongly related to newsworthiness. Another more conceptual definition of timing is novelty, which connects to originality. This broader definition of timing, along with visual media and whether the tweet is from an organisation, constitute the three main features that influence tweet success. All three features will now be discussed in more detail, along with the other discoverability, content, and style features derived from the theoretical framework. Comments shall be made upon the applicability of each component of the framework to new media.

7.1.1 *Organisational Power*

As organisation was the most important factor for tweet resonance in the quantitative analysis, it will be discussed first. The strength of this feature relative to that of the other predictors is extreme, and is supported by the qualitative findings. Multiple interviewees who represented organisations (most of which were campaigns) affirmed that they spent countless hours on Twitter—some, indeed, practically lived on it. They were employed to do so, but they were also politically motivated, given the significance of the referendum.

Compared to individual interviewees, organisations thought much more carefully about tweet strategy and planned out their content almost scientifically; they were more concerned with tweet success. Their Twitter accounts were often managed by multiple people, and relied upon a vast range of digital marketing tools—We Are Europe is a particularly good example of this, as they used four different tools for different activities. This allowed for tweet engagement to be orchestrated to a certain degree—and on one occasion for exceptional success to be orchestrated (Figure 5.1). Moreover, the interviewees who managed both personal and organisational accounts focussed more on the organisational.

All in all, hypothesis 2 has been more than affirmed. Two qualifications, however, must be made about this feature. Firstly, *prominence* is a conditioning factor: established and grassroots organisations do not have the same degree of advantage, although both have an advantage over non-influential individuals. Influential individuals (i.e., celebrities) are more similar to established organisations than they are to non-influential individuals. This is demonstrated by the fact that mentioning a celebrity is a strategy that interviewees used to increase the appeal of their tweets, and the fact that three of the eleven most popular tweets are from celebrities.

Secondly, a distinction must be made between organisational accounts and the accounts of their representatives—they both have an advantage, but the representatives might have even more of an advantage. Interviewees commented upon this distinction: Twitter is ultimately about *personalities*, and organisational representatives are able to say things (from their personal accounts) that organisations themselves cannot. Certain forms of playfulness and negative campaigning are inappropriate from organisational accounts; more shall be said on this in Subsection 7.1.3. The main advantage that grassroots organisations have over established organisations is that they are able to express more personality. Scientists for EU is an excellent example—they developed a reputation for being cheeky, while still making serious points about science.

Given the advantage of personal accounts, it is perhaps no coincidence that none of the extremely popular tweets came from organisations themselves, even though the two most popular users in the dataset are media organisations—BBC Breaking and The Economist (Table 6.1).⁸⁵ It is more fulfilling to engage with other people, as opposed to with relatively faceless companies or campaigns. Users might follow the BBC and the *Economist* to identify with them and see their updates, but they might prefer to like and share tweets sent from journalists themselves.

Media organisations are of particular interest to this thesis because they relate to gatekeeping and agenda setting theory. To what degree is Chaffee and Metzger's (2001) prediction about the disruptive impact of new media evidenced by this thesis? Has there been a shift from 'what issues the media tell people to think about to what issues people tell the media they want to think about' (p. 375)? Put more broadly, has the rise of what Castells (2009) terms *mass self-communication* challenged organisational power?

Given the demonstrated influence of organisations on new media, the answer is not to a large degree; this echoes Sayre, Bode, Shah, Wilcox, and Shah's (2010) findings. In spite of

⁸⁵ BBC News is a separate account from BBC Breaking and is the seventh most popular user, with over 6.7 million followers. Both belong to the same parent organisation.

the qualifications noted above, organisational tweets receive far more engagement on average than non-organisational tweets, and their accounts have more followers. The mean retweet counts for organisations and individuals are 43.26 and 2.97, respectively; the median counts are 10 and 0. This extreme difference aligns with that in average follower count, which is 33,459.86 for organisations and 3,679.27 for individuals (median follower counts are 7,882 and 435). Offline power structures are reflected and reinforced by social media, rather than collapsed. The era of tweets has not heralded the end of what Chaffee and Metzger (2001) call mass communication: users have more control than ever before about where they obtain their news, and there is an ever-increasing diversity of sources, but mass media organisations have not lost their influence. In fact, given the vast number of low-quality sources and tweets from political bots, mainstream news sources might appear all the more trustworthy and appealing.

There are, however, opportunities for individuals and less mainstream groups to influence the news agenda in the new media landscape. Interviewees who represented such groups provided illustrative examples of this: the Green Party triggered a news article with a reactive tweet to a member of the political elite (Figures 5.29 and 5.30), and Scientists for EU actually wrote and promoted an article on behalf of mainstream media, thereby defining the media narrative (Figure 5.3). In general, the Green Party leverages social media because it allows them to voice what is not expressed in mass media; their tweets are, accordingly, mostly reactive. Scientists for EU, on the other hand, leveraged social media to establish *community credibility*, which in turn attracts attention from mainstream news outlets—and coverage in mainstream media serves to reinforce credibility on social media. Thus, the more ‘democratic’ credibility exists in a positive feedback loop with the more ‘authoritative’ credibility; social media and mass media are part of the same news ecosystem and can reinforce each other, as opposed to controvert each other.

The Green Party’s reactive use of Twitter reflects a deeper trend that qualifies the power of organisations on the platform. The BBC is the most prominent media organisation in the UK and, accordingly, the most followed user in the dataset (Table 6.1). It is also, however, a resonant target of negative critique from Leavers and Remainers alike. Multiple interviewees expressed their deep dissatisfaction with BBC coverage. Ironically, Paul’s critique of the BBC is that it strove too much for ‘balance between the different campaign groups at the expense of objective fact-checking’. UKtoStay.EU was founded partly because of the misrepresentation of the truth in mainstream media. This relates to Robert’s motivation to actively use Twitter in 2013, as he lived near Tahrir Square during the military coup. He saw a need to not only fill in the gaps in mainstream news coverage, but also to *counter its narrative* by presenting a more

accurate version of what was happening. However, counternarratives and reactive (even vituperative) criticism by no means diminish the influence that organisations exert or undermine their authority. Thus, this thesis lends support to Weimann, Weiss-Blatt, Mengistu, Tregerman, and Oren's (2014) claim that agenda setting has not lost its relevance, but needs to be 'readjusted to some degree to reflect changes brought about by the patterns of flow, structure, access, and ownership of new media' (p. 821). Importantly, the '“old” issues of power and resistance, and structure and ownership' still persist (p. 822).

7.1.2 *Tweet Attributes (Discoverability)*

Four tweet attributes that contribute to discoverability were examined in this thesis: hashtags, @mentions, hyperlinks, and visual media (images, videos, GIFs). Only visual media was positively related to both measures of success in the models—it was the second strongest factor, after organisation. The population-level regressions reveal that video is indisputably the most resonant visual item. Images are the second most resonant, followed by animated GIFs. It should be noted, however, that videos and GIFs are extremely rare media types: only 0.51% and 0.38% of tweets in the full dataset include them, whereas 22.03% contain images. The interview data corroborates these statistics—as Paul remarked, creating an original video requires significantly more effort than taking a picture. And unlike images, videos are not as easily borrowed from other sources; this raises a copyright and Twitter etiquette issue. If a user is willing to invest time into producing them, however, they are an effective strategy for tweet resonance and account growth. Hardly any interviewee tweets contained GIFs, and only one interviewee mentioned them, but his comment was critical.

Why are images and videos so effective? Apart from the fact that they arrest attention and are favoured by the Twitter algorithm for business model reasons, they also extend and enhance the message of a tweet. Numerous examples of this are provided in Chapters 5 and 6; it is no coincidence that six of the eleven most popular tweets contain images, and one contains a video. The interviewees commented on how visuals can add humour, emotion, and authenticity. They also serve an invaluable PR and marketing function for campaigns, as they influence the public *perception* of how active and successful the campaign is. Photographs from events are particularly helpful for this, and Brendan's most successful personal tweet serves as an excellent example (Figure 5.32). Taken at a Labour Leave rally in Stoke, it included a photograph of an extremely full room. Given the fact that most people did not associate the Labour Party with Leave, this was all the more effective (*novelty* will be discussed in Subsection 7.1.5). The added benefit of photographs over other types of images (e.g., graphics, charts) is

that they have a connection to the real world. This connection is one of the reasons the ‘Kiss of Death’ tweet from the We Are Europe campaign was so resonant—it included a photograph not just of the mural, but of Harriet and another person standing in front of the mural (Figure 5.1). The two figures underscore the fact that the mural physically exists; it is not just a digital graphic.

Thus, hypothesis 1c has been more than affirmed by both the quantitative and qualitative data. It is important to note, however, that visuals are not unconditionally effective, even if they are sent from influential accounts. As the surprisingly unsuccessful tweets examined in Chapter 6 reveal, certain types of media items are not resonant: images that express crude language and violence (in the context of a playful tweet), and videos that do not communicate a message (Figures 6.15-6.18). Other attributes of ineffective images and videos identified by the interviewees are lack of visual contrast, incorrect dimensions, and excessive length.

What is more interesting than the effectiveness of visual media is the *ineffectiveness* of hyperlinks and hashtags, and the fact that @mentions seem to have no significant relation to tweet resonance. All three factors improve discoverability, for the reasons described in Chapter 2. Despite this, the sample-level regressions indicate that every hashtag that is added to a tweet decreases retweet and like odds; including a hyperlink also decreases both measures of resonance. The population-level regressions reinforce these findings. @mentions are the only non-statistically significant discoverability feature in the models. It should be noted that all tweets in the dataset contain at least one partisan hashtag. The results, therefore, do not indicate that hashtags are ineffective—rather, including more than one is ineffective, despite the added discoverability that they afford.

What could account for the ineffectiveness of these tweet attributes? As noted in Chapter 2, there is an important distinction to be made between discoverability and arresting attention. Increasing the former does not guarantee the latter—just because a tweet has a higher chance of being seen does not mean it *will* be seen, much less engaged with. Visual media increase discoverability and are attention-arresting because they take up space on the Home feed; hyperlinks, hashtags, and @mentions are not as favoured by the Twitter algorithm (and thus do not increase discoverability as much), and are also not as attention-arresting. Hashtags and @mentions are part of tweet text, and so do not extend the space a tweet occupies on the feed. As the interview data illuminates, hyperlinks do not always generate previews (Twitter Cards); even if they do, not all previews come with a large image. That being said, hyperlinks are a standard digital marketing strategy, and define the tweeting style of many users. Because limited information can be communicated in a tweet, links are important for more substantive

points. If the links are from well-regarded sources, they can boost both credibility and authenticity. Journalists and bloggers in particular use Twitter to share links to their articles; each article is often shared multiple times.

In marketing and political contexts, hyperlinks are also important because they are connected with *action-based campaigning*. The EU referendum clearly represented one such context—actions included registering to vote, attending an event, and donating to a campaign. We Are Europe leveraged the anticipated popularity of the ‘Kiss of Death’ tweet to direct their supporters to the register to vote page (Figure 5.1). Ironically, links that relate to campaign actions are more closely connected with *persuasion*, the ultimate metric of success, than retweets or likes; such actions have real life impact, whereas retweets and likes are by definition restricted to the Twittersphere. As Alberto pointed out, there is actually a tension between URL clickthroughs and retweets, because link clicks lead users off of Twitter. This helps to explain the negative relation revealed by the regressions. Because they are not provided by the public Twitter API, link clicks are not used in this thesis as a metric of success, but they are an important metric for some users.

Hashtags also represent social media best practise, especially for new accounts. They are generally perceived to be helpful, in that people often search for them, and they are associated with trending topics—if enough people use a hashtag, it will trend (thereby increasing the discoverability of tweets that contain it). Another important dimension for the effectiveness of hashtags is *time sensitivity*—hashtags are especially useful, indeed indispensable, during live events. A prominent subcategory of tweets was sent during televised debates, all of which adopted a specific hashtag (e.g., #BBCDebate, #ITVEURef). Users who tweeted during these debates consistently used their hashtags, as this significantly increased the discoverability of their messages. More people follow time-sensitive hashtags than evergreen hashtags such as #Brexit and #EURef; their entire point is to act as channels for real-time commentary. Tweets sent during live events should be treated differently from other tweets, and will be discussed more in the next subsection.

Hashtag effectiveness can also be engineered to a certain degree, as the examination of exceptional tweets has shown. All of the surprisingly successful tweets include #VoteLeave and are from users with extremely low follower counts (Figures 6.11-6.13). The hashtag is the only feature that these tweets have in common, and so it is highly likely that it contributed to their success. The pro-Leave community on Twitter actively chose to follow this hashtag, and to share supportive tweets that contained it. Thus, the effectiveness of a hashtag is directly correlated with the effort its proponents exert to promote it.

All of this being said, several interviewees maintained that hashtags do not work. Mike claimed that people do not click on them, although they serve an aesthetic function because they appear as blue text. Sam declared that they are clunky and come across as attention-seeking. Reinforcing the model results, Leni and Logan commented on how including too many hashtags (in Leni's opinion, more than two) is ineffective because it looks spam-like and overinflates the ratio of blue to black text.

@mentions were the only non-statistically significant discoverability feature. This aligns with the interview data—one of the feature's main functions is to credit or support another user. This usage is viewed more as Twitter etiquette than as a deliberate technique for tweet success, even though it can attract engagement as a side effect (usually from the person being credited). Mentioning users to ask for support *is* a strategy for tweet success, and the prominence of the user is positively correlated with the extent of tweet success, but it is only effective if the user knows the author of the tweet. It takes time and effort to develop interpersonal relationships, either online or offline. Many campaigns that were interviewed had an influencer strategy that they developed alongside tweet content; important messages were systematically communicated to top influencers before they were published on Twitter. When they were published, these influencers were often *tagged in images* as opposed to @mentioned in tweet text—this strategy notifies the user without taking up any of the 140 characters. As it is not captured in publically available tweet metadata, it could not be included in the models.

Thus, hypotheses 1a and 1d have not been supported, and nothing conclusive can be said about hypothesis 1b. In light of this, it is interesting to observe that of the most popular tweets examined in Chapter 6, only one uses @mentions, one uses a link, and two use more than one hashtag.

7.1.3 *Rhetorical Content*

The overarching conclusion from the quantitative analysis is that content is not as important as discoverability. None of the content features are significant in the sample-level retweet models. Only in the like models are several content features significant: the issue categories *healthcare* and *sovereignty*. However, although they were salient at the time of the referendum, they have a moderate *negative* relationship with whether a tweet is liked. The qualitative analysis broadly supports the quantitative findings: interviewees did not focus on content as much as context; some directly stated that they did not believe content was important (unless the account had many followers). However, the interviews also illuminate that content *is* important in certain contexts.

One such context is live events. As noted in the previous subsection, this context is one in which hashtags are indispensable, as they unite the conversation—bringing together a vast spectrum of voices in the same Twitter stream. Live events were especially important during the referendum campaigning period because there were a number of televised debates in addition to commentary provided on recurring political shows. Tweets sent during these popular events represent a special subcategory that should be treated differently, because their volume is far higher. Thus, many unsuccessful interviewee tweets fell into this subcategory. Despite the increased competitiveness, live events constitute a more meritocratic environment for tweet content because everyone is looking at the same hashtag stream at the same time. Timing, then, is arguably more important in this context than account prominence—reacting to something in the moment in which it is said is essential. But spontaneity can be planned, and planned spontaneity is an excellent strategy for tweet resonance during live events. UKtoStay.EU and the Green Party executed it well, with their libraries of well-researched content and graphics.

The *logos* and *pathos* dimensions of Aristotle's conception of rhetoric are conflated in both the quantitative and qualitative results: *logos* as defined in this thesis is message framing (identification vs. division), following Burke's (1969) approach, and *pathos* is represented by three politically relevant emotions derived from Marcus, Neuman, and MacKuen's (2000) theory of affective intelligence (enthusiasm, anxiety, aversion). All subcategories of both variables are at least moderately collinear; the strongest correlations are between identification and enthusiasm and division and aversion (Appendix G). This collinearity is reflected in the fact that interviewees discussed emotiveness in general and did not distinguish between different emotions, *except* in the context of message framing: anger is directed against members of the opposition, and so denotes negative campaigning, whereas enthusiasm denotes positive campaigning. It is interesting to observe that emotion was hardly ever mentioned on its own, and was also conflated with issue. This is because pathos requires a target—anger is not generally expressed in and of itself; 'I'm angry' is far less common (and resonant) than 'I'm angry at the EU'.

Nick T. did, however, differentiate between the emotions with his assertion that in politics, enthusiasm doesn't work. Anger and fear appeals are more resonant because people engage more with *conflict*. This is evidence against affective theories, which view emotion as a coping mechanism—and hence enthusiasm as more effective than anxiety (Roseman, Abelson, & Ewing, 1986). Neither emotion nor message framing is significant in any of the models, but all of the most popular tweets examined in Chapter 6 are negatively or positively framed. More

shall be said on these tweets in the next section on grassroots success. Overall, only qualitative evidence has been provided in support of Berger's (2013) claim that content containing emotions is more likely to be popular than content that does not, and Atkinson's (1984) finding that positively and negatively framed messages are effective.

It is tempting to ask whether positive or negative campaigning is more effective in a contentious context like Brexit. Indeed, this is the question upon which the literature on positive vs. negative campaigning is premised. But as the qualitative analysis has shown, it is not the right question to ask, because its answer depends entirely on context; this reflects the mixed findings from Lau, Sigelman, and Rovner's (2007) meta-analysis of 111 studies of negative political campaigns. A better question to ask is *under what circumstances and from whom is negative or positive campaigning more effective?* Debates represent a context in which negative campaigning is both rampant and resonant. Their entire point is to bring opposing viewpoints head to head with each other; rebuttals are as important as original arguments. Indeed, the power of negative campaigning is its ability to undermine the credibility of the opposition. Positively framed live tweets from campaigning events are also effective, as displays of solidarity are powerful.

Two broader circumstantial factors that influence the effectiveness of message framing are the socio-political and temporal contexts of the message. As Chapter 5 has demonstrated, particular targets of negative campaigning are especially resonant, like Douglas Carswell and Nigel Farage. These prominent politicians are unpopular with Remain campaigners because of their controversial opinions and brash personalities; in order to recognise their unpopularity, knowledge of contemporary British politics and culture is necessary. Knowledge of the broader temporal context of UK politics is also important. As Logan pointed out, negative campaigning is not as effective with younger generations as it is with older generations because the political landscape has been highly polarised since Margaret Thatcher was Prime Minister; before the 1980s, however, there was an era of consensus politics. Thus, for millennials positive campaigning is novel and exciting, whereas for their parents the same can be said for negative campaigning.

The second dimension that influences the success of message framing is the identity of the messenger. Chapter 5 has illuminated that two aspects of identity are important: whether the messenger is an individual or an organisation, and the messenger's personality. Because the Green Party is an optimistic organisation, it is rebuked whenever it sends negatively framed messages. Its leaders, however, can be negative—this is best demonstrated by Molly Scott Cato's tweets about her 'Bad Boys of Brexit' website, which were more successful than any of her previous tweets. When the Green Party's official account shared this website, on the other

hand, their followers reacted unfavourably. For the second level of success, Chapter 6 has demonstrated that *user influence* is another aspect of identity that must be taken into account. The analysis of the most popular tweets in the dataset reveals that negative campaigning can be highly effective from non-influential users, whereas both positive and negative campaigning are highly effective from influential users. Interestingly, there are no examples of positively framed popular tweets from non-influential users. More shall be said on this in the next section.

The interviewees' general recognition of relevant issues as resonant supports Ansolabehere and Iyengar's (1995) assertion that effective advertising speaks directly to voter concerns and West's (2005) contention that which topics to mention in advertisements is a key campaign decision. In light of this, it is surprising that two of the issues—healthcare and sovereignty—are negatively related with likes, whereas the rest are not statistically significant. None of the issues are significant in the retweet models. Once again, a contextually-insensitive question has been posed: it should not be *whether referendum-related issues are resonant*, but rather *under what circumstances and from whom are specific issues resonant*? West (2005) and Brooks and Geer (2007) found that issues-oriented messaging is effective in negative campaigns. The qualitative analysis supports this, as interviewees emphasised the importance of *fact-based arguments* in negative campaigning. It has also illustrated that the contextual factor of timing conditions thematic effectiveness. Novelty, an intangible aspect of this feature, is essential because the issues were discussed throughout the entire campaigning period. As the same arguments were tweeted over and over from multiple accounts, a unique or topical angle had to be adopted to arrest users' attention and attract their engagement. Alberto's tweet about Turkey being granted visa free travel to the EU (Figure 5.24) was successful not because it was about immigration, but rather because it was about a newsworthy and highly controversial aspect of immigration. Novelty is also relative to the identity of the messenger: immigration was a far more unique topic for Remain campaigners than the economy because Stronger In focussed on the latter and almost completely ignored the former. Timeliness plays a role here as well: immigration was most resonant for Remain campaigners in the days leading up to the referendum, as this was when the topic was no longer taboo for their audience.

Some issues are unique by definition, because they are niche and not as widely discussed in mainstream media. These issues are particularly resonant among specific groups, and are represented by special interest campaigns: Fishing for Leave, Scientists for EU, YEM, and the Green Party are excellent examples (they focussed on fishing, science, education, and the environment, respectively). The voices of these groups might be small offline, but they can be extremely loud on social media. Thus, a substantive distinction must be made between the five

more mainstream issues identified as relevant (the economy, immigration, security, sovereignty, and healthcare) and the less mainstream issues that were collapsed into the ‘Other’ category. This category is merged with the five mainstream issues in the population-level regressions, but in the sample-level regressions where it is maintained as a separate category, it is not significant. This indicates that not all niche issues are by definition resonant—as Logan and Harriet observed, gender violence and women’s rights did not have social cache during the referendum period.

Resolution, the final content feature, has no significant relationship with either measure of success—but some interviewees were convinced that it did. This feature was part of Schudson’s (1984) dimensions of the potency of a cultural object, and its effectiveness needs to be qualified. It is undeniable that the EU referendum was an extremely action-oriented event; moreover, Sean and Elise identified action orientation as a core element of Twitter best practise. Twitter-specific resolution defined Andrew’s messaging style: he included ‘please RT’ in many of his tweets, and it proved to be quite effective. Other types of resolution, however, might not be as effective; the success of some might neutralise the failure of others, so having an action in and of itself is not enough. The qualitative analysis has revealed that less urgent actions that take users off of Twitter—e.g., adding a Twibbon or joining a petition—might be successful for link clicks, but are not successful for retweets or likes. However, urgent actions that take users off the platform (registering to vote and voting itself) can attract likes and shares due to their tangible impact on the referendum outcome. Although the effectiveness of resolution related to retweets (‘please RT’, ‘read and RT’, and ‘spread the word’) seems obvious, it is interesting to observe that in the population-level regression, in which only this subset of the feature is captured, resolution decreases retweet odds (Table 4.13). Because the variable is not significant in the sample-level regressions, this negative relation cannot be corroborated.

All in all, nothing conclusive can be said about the hypotheses that relate to message framing and emotions: 3a, 3b, 4a, 4b, and 4c. Contextual factors play too significant a role in the resonance of these features. Evidence has been provided against hypothesis 5; hypothesis 6 is supported only by the qualitative results, not by the quantitative.

7.1.4 *Rhetorical Style*

The quantitative analysis has indicated that style is even less important than content: not only are most of the features not significant for either retweets or likes, but they also add complexity with the least improvement of fit. In fact, the fit of the retweet model with style and control features is *lower* than the fit of the model with control features alone. This result can be explained

by the fact that style cannot be isolated from content, as it represents the way in which content is expressed. It can also be explained by the fact that unlike content, style is more concerned with *execution* as opposed to presence. Crafting the perfect witty remark, antithesis, or rhetorical question is difficult, and operationalising the quality of execution beyond a very limited degree (e.g., the presence of typos) is even more so. Qualitative analysis is necessary for this, and it is represented by Chapter 6. The popular tweets examined in Section 6.1 are exceptional examples of well-executed playfulness; the surprisingly unsuccessful tweets presented in Section 6.2 are excellent examples of poorly executed playfulness.

Thus, although it is not significant at the global level captured by the regressions, playfulness can be extremely effective, especially when it comes to the second level of success. The fact that most of the extremely popular tweets contain both playfulness and negative campaigning aligns with Baym's (1995) findings from her Usenet group study. As Mulkay (1988) observed, hostile humour is particularly relevant to politics because of its adversarial nature and interpretative diversity. Mulkay also observed that humour is a means of 'conveying serious information without appearing to do so', and the popular tweets in Figures 6.1-6.4 are examples of this—they comment upon four of the five mainstream issues relevant to the referendum (p. 69). Playfulness is not at all incompatible with substantive issues, and in fact can be an effective vehicle for their delivery.

More generally, the interview data supports the effectiveness of playfulness. Harriet asserted that the performance of wit is fundamental to Twitter. Indeed, parody accounts are abundant on the platform, and the user who personified Paddy Ashdown's hat on election night in 2015 is a prototypical example. The overnight success of this account also reveals an interesting tension between puerility and playfulness. Although humour can be used to deliver substantive messages, it is associated with immaturity as well, and Nick T. recognised this dimension as more resonant—his tweets with the least insight receive the most engagement. Thus, for accounts concerned with establishing serious reputations, certain forms of playfulness might be damaging. As noted in Subsection 7.1.1, playfulness is also a far less effective technique for organisational accounts than it is for personal accounts. Almost by definition, organisations are more concerned with reputation than private individuals.

The usage of questions is the only style feature that is statistically significant in the like models; no style features are significant in the sample-level retweet models. In the context of all predictors, it has a moderate *negative* relationship with whether a tweet is liked. The population-level full model reveals that it also has a negative relationship with whether a tweet is shared. The interviewees expressed mixed opinions about the resonance of questions. While some

maintained that they are effective because they are attention-arresting and provocative, others noted that if they are overused or used in a painfully obvious way, questions come across as too promotional and clickbaity. None of the most popular tweets contain this feature; instead, they consist of emphatic statements. On the whole, the quantitative and qualitative results indicate that there are more poorly executed questions than well-executed ones, and that questions are not an important factor for the second level of success. This is at odds with the extensive literature on the effectiveness of rhetorical questions and the cognitive response approach (Frank, 1990; Koshik, 2005; Ahluwalia & Burnkrant, 2004; Zillman, 1972; Mothersbaugh, Huhmann, and Franke, 2002). However, it should be noted that a more general definition of question was captured in this analysis—the detection of rhetorical questions requires more context than is provided by a tweet, and is also inherently subjective.

The two stylistic devices that Atkinson (1984) identified, lists and contrasts, are not only insignificant in the models, but also were barely mentioned by any of the interviewees. This signifies that features effective in political rally speeches delivered in front of live audiences are not effective in political tweets; Atkinson's theory does not translate to new media settings. Indeed, the resonance of these devices is in part grounded in their acoustic delivery—they signal to the audience that the speaker is coming to the end of a phrase, which is when clapping is expected. However, interviewees did discuss contrast in the context of images. *Visual contrast* is created through the juxtaposition of colours; foreground elements are emphasised through brighter colours, whereas background elements are more muted. Moreover, several of the most popular tweets in the dataset contain examples of both contrasts and lists; the best example of these features in combination is in the tweet about war and bananas (Figure 6.3). This suggests that Atkinson's theory may be applicable to the second level of success.

Brevity (and the related concepts clarity and simplicity) is the most highlighted style feature from the interviews that was not included in the models. It applies to text and visuals alike, and can be defined in all media as expressing only one message. The affordances of Twitter encourage brevity in tweet text: at the time of the referendum, the length constraint was 140 characters. More effective tweets do not contain unimportant filler words. Length is also important in the context of videos; the best videos are short soundbites that capture the most poignant moment.

All things considered, what can be said about the resonance of style features? Hypothesis 9 has been strongly supported by the qualitative results, but not by the quantitative. Hypotheses 7a and 7b have been neither proven nor disproven, and hypothesis 8 has not been supported. Brevity represents a style strategy that was not measured by the models, but its

relationship to tweet success is strongly supported by the qualitative data. It is interesting to consider this subsection and the previous subsection together, in light of Atkinson's (1984) comments on the small number of content and style features that characterise resonant political messages. This narrowness is accounted for by audience size and diversity—the larger and more heterogeneous the audience, the more difficult it is to arrest attention and communicate ideas. Attendees of party political rallies are far more homogeneous than Twitter users, who represent a wide range of demographics (Greenwood, Perrin, & Duggan, 2016). Thus, the fact that none of the content and style features are positively related with tweet resonance in the models is not surprising. The fact that the qualitative results illuminate mixed opinions about many of the features is also not surprising—different features appeal to different Twitter users; the various subcategories of each feature add a further layer of nuance. There is no one size fits all approach, and no single aspect of content or style is universally resonant for all users. Rather, the success of these features needs to be contextualised and tailored to specific audiences, taking messenger identity into consideration.

7.1.5 *Timing*

As mentioned in the introduction to this section, timing represents the third main feature that contributes to tweet success, and it is also the most multifaceted. Three measurable facets are captured in the regression models and treated as control variables: time of day, day of the week, and days before the referendum. The models, however, only explain a fraction of the variation in tweet success; it is clear that significant explanatory features are missing. This is mostly a measurement problem—many of these features are quite high-level and require context that cannot be obtained through Twitter alone.

The context-dependent facets of timing that emerged during the interviews represent such features, and can be separated into two categories: timeliness or topicality, which relates to newsworthiness, and novelty, which relates to originality—telling people what they do not know. These categories also overlap, in the sense that newsworthiness by definition has a shelf life. The number of days before the referendum captures only a narrow definition of timeliness. To more fully identify whether a tweet contains timeliness and/or novelty, an awareness of current events as well as social, cultural, and political knowledge are necessary. Thus, Twitter needs to be analysed relative to the mainstream media, as it exists in an ecosystem not just with other social media platforms, but also with news outlets. A tweet is timely if it reflects news headlines, or more conceptually provides information at the moment in which it matters; a tweet is novel if it is the first to express something, or provides information that is outside of general

knowledge. Another aspect of novelty is being counterintuitive—telling people what they *do* know, but from a different angle.

As discussed above, timing and novelty play an especially important role for issue. Tweeting about immigration at some point during the referendum campaigning period is not effective. Tweeting about visa-free travel to the EU for Turkey at the moment in which the debate arose in mainstream media, however, *is* effective.

Timeliness and novelty apply not just to content, but also to *strategy*. The most effective strategies are dynamic as opposed to static; too much of any feature is not effective. Even visual media have reached a saturation point, according to Sean: ‘Longer forms of media are now becoming more popular on other channels. I feel like we might have hit trough attention span’. What’s more, strategies have to be executed at the right time: negative campaigning in the wake of Jo Cox’s death was extremely unsuccessful (by virtue of being inappropriate), whereas negative campaigning during a live debate is quite resonant (by virtue of being appropriate). Contextual timing is arguably more important than time of day; the key to tweet success is to know not just what to say and how to say it, but also what to say *when*, or which strategy to use *when*. This flexibility is represented by We Are Europe’s three-pronged approach: to ‘inform and inspire, amuse or entertain, and cause friction’. All three prongs are important in different contexts, and together constitute a robust messaging strategy. An illuminating example of the importance of contextual timing over content is the fact that UKtoStay.EU tweeted material that they had tweeted during the referendum after the referendum. It was more resonant after the referendum due to the different temporal context: Remain groups were far more active on Twitter, as they were driven by anger.

Strategies should also shift in accordance with user-level attributes: account influence (as measured by follower count) and increased Twitter experience. Larger accounts do not have to rely as much upon discoverability features as smaller accounts; they can focus more on crafting compelling content and style. Increased Twitter experience allows for strategies to diversify even further, and for new techniques to be tested (such as retweeting one’s own tweet or tagging users in images). Continuous experimentation represents a strategy in and of itself, as it breeds novelty—Paul was actually inspired by the interview itself to do this, as he reflected in an email:


Reviewing everything in our discussion has also motivated me to keep on trying new things with the UKtoStay.EU twitter account. You might find a new experiment I’ve started interesting: links posted in a Facebook group I run (called brexit news) are now automatically tweeted out. The Facebook content is completely crowdsourced, which adds an entirely new dimension to the usual twitter activity. Twitter automation, but with people instead of bots. Going great so far! :)

This novel strategy is particularly interesting, as it blurs the boundaries between people and bots, as well as between Twitter and Facebook.

7.1.6 *Revised Cultural Diamond for New Media*

In light of these results, it is evident that a new framework is needed to understand the resonance of digital cultural objects. Griswold’s (2013) schema of the cultural diamond is a useful starting point, but has to be revised and reweighted. This thesis has demonstrated that five factors account for cultural power in the attention economy-governed landscape of new media: the producer, the receiver, the cultural object itself, the social context in which it is created, and the algorithmic context in which it is created. *Algorithms* are the central mediating force and influence all of the other factors: they control what content is discoverable in the first place, and reinforce (or amplify) existing human biases and hierarchies. The interrelationships between the five factors are visually depicted in Figure 7.1.

FIGURE 7.1: Revised Cultural Diamond for New Media



Updated schema for understanding the resonance of a cultural object in new media.

Producers are given extra visual weight (as indicated by the sizes of the rectangles) because it has been illustrated that organisational authority overwhelms that of individuals in new media; this factor represents the creators and deliverers of cultural objects, and is more influential than the cultural objects themselves. Receivers, on the other hand, are given reduced visual weight to emphasise their diminished agency: new media is characterised by fast-moving, high-volume content, and their attention is manipulated against their will. Visual rhetoric is extremely helpful for the resonance of digital cultural objects, as images and videos are naturally attention-arresting as well as favoured by algorithms. Timeliness and novelty are also important contributors to resonance, and they are anchored in the social world, where mainstream media plays a significant role (as social media is defined relative to traditional news sources).

Thus, Castells' (2009) contention that individuals are granted 'unprecedented autonomy' in the era of new media is controverted (p. 135). Even though in theory any receiver can also be a producer, mass culture theory has been upheld: culture is still strong and receivers are still weak, relatively speaking. The conclusions of two studies from the information science literature have been supported—Weng, Flammini, Vespignani, and Menczer (2012) demonstrated that meme popularity is derived entirely from competition for users' attention and the structure of Twitter; Weng, Menczer, and Ahn (2013) found that the future popularity of a meme can be predicted from the number of communities it penetrates. Message content does not have to be considered.

Taking McLuhan (1964) one step further, it can be said that in new media, the algorithm is the message. Algorithms are not static, however: they are intelligent and dynamic. This dynamism undergirds the entire cultural diamond. Content production strategies also have to be dynamic to adapt to Twitter's ever-evolving design, which is a direct response to users and their interests and needs. Strategies shift, and affordances shift as well in a continuous feedback loop. The most significant design shift was the doubling of the character limit on 7 November 2017, as the 140-character limit had defined the platform from the very beginning (Perez, 2017). This clearly has an impact on content strategies, and might diminish the importance of visual media because more can be expressed with text (more space can also be consumed by text on Twitter timelines).

Another feature that emerged after the referendum that also expanded the limits of expression is the thread. Defined by Twitter as 'a series of connected tweets', it is visually denoted by a line on the left margin of its constituent messages (Twitter Help Centre, 2018b). The tweets in a thread can be published all at once, or tweets can be added to existing threads at any time (by the author). By stitching content (both text and multimedia) together, threads

allow for the communication of detail and substantive depth—quite a contrast to what can be expressed in 140 characters. As Robert observed, ‘It’s not really tweeting anymore; it’s like writing. If you write a thread with 20 tweets, with 140-280 characters each, then you almost have a text’. Indeed, threads are especially useful for journalists and bloggers and define their tweeting style; Robert suspects that for political journalist Jon Worth, they attract more of an audience than blog posts. Paul remarked that Twitter threads from experts directly after events contain brilliant analysis: ‘They’re summarising an article they’re in the process of writing... It has to be boiled down into really short tweets, so it has to be really concise’. Because they are insightful and quick to read, threads are highly shareable. It is important to note that this feature already existed informally before Twitter introduced it: previously, users had replied to themselves to connect their tweets, and numbered them to keep everything in order. This is yet another example of Twitter adapting in response to user behaviour.

Other platform design shifts directly influence the dependent variables in this thesis. Bookmarks were officially launched on 28 February 2018, to address user requests for a way to privately save tweets to read later (Perez, 2018). Indeed, several interviewees used the Like button for this purpose, even though it overtly communicates endorsement due to its heart shape. Apart from being confusing, the Like feature has been so controversial that Twitter is considering its removal. In the words of their Communications team, ‘We are rethinking everything about the service to ensure we are incentivizing healthy conversation, that includes the like button’ (Lunden, 2018). Likes have been abused ever since Twitter shifted its algorithm to include them in notifications and timelines—thus, if many users like a malicious tweet, it can spread far more than it used to. This relates to Twitter’s decision to adopt a Facebook-based approach for displaying content, which takes engagement (likes and shares) into account. The most engaged with tweets are shown first because they are deemed to be the best and most relevant, but this is a stark contrast to Twitter’s more democratic chronological timeline, which displays the most recent tweets. Due to user demand, an option to switch back to the chronological view was introduced in September 2018 (Schofield, 2018). More design shifts relating to the Like feature and timeline algorithms are surely on the horizon, and tweeting strategies will have to shift in response to them.

7.2 Grassroots Success in New Media

Given the overwhelming influence of organisations and public figures in both old and new media, how can non-prominent individuals and groups achieve success on Twitter? Figure 7.1

captures the general dynamics of reception in the new media landscape, but exceptions can be made. The potential for the disruption of offline hierarchies is always present through social media—it just needs to be activated. The question of *whether* social media allows non-influential individuals and groups to challenge organisational power, or even *to what extent* social media allows this to happen, is thus fundamentally limited. A better question is *under what circumstances* are non-influential individuals and groups most likely to activate their potential to reach a wider audience—and succeed? Two strategies for grassroots success will be presented in this section, derived from empirical examples from Chapters 5 and 6 and the insights presented in the previous section.

7.2.1 *Playfulness and Negativity*

The first strategy is derived from analysis of the exceptionally popular tweets in Chapter 6. It is remarkable that four of the five most resonant tweets in the dataset are from non-influential individuals; given the vast difference in their numbers of followers, these users are at a severe disadvantage compared with organisations and celebrities. When it comes to high levels of success, content and style clearly play more of a role: several of the popular tweets contain no media at all, and only one contains a video.

Playfulness and negative campaigning are the predominant characteristics of tweets from non-influential individuals. Witty and humorous personalities are celebrated on Twitter, and in the highly polarised political context of the referendum, negative campaigning was inevitable. What is interesting is that critical humour seems to be more effective than positive humour from private individuals—none of the popular tweets from these users contained identification rhetoric, whereas the majority of those from organisations and celebrities did. In fact, several of the popular tweets from celebrities are straightforward (i.e., non-humorous) endorsements of Remain. Because of their fame and status, it is sufficient for users like Richard Branson and Gary Lineker to express their support in a matter-of-fact and non-creative manner; this expression of support might even constitute a news item (and thus be newsworthy). After all, celebrity is in and of itself a feature that increases tweet engagement. Had any of the non-influential individuals published straightforward endorsements, they probably would not have received much engagement. It is of little consequence if someone without a public voice supports Leave or Remain; what matters more for these users is creativity, and in particular with a critical bent.

Creativity can be expressed not only through content, but also through typography. Including multiple line breaks has the same attention-arresting effect as including an image, as

it consumes space on Twitter timelines. This technique also enhances the message itself, as evidenced by the tweet about war and bananas in Figure 6.3. The repetition of ‘War’ on separate lines enacts the endless episodes of bloodshed in European history, and creates a sharp (and humorous) contrast with ‘Arguments about bananas’. No picture could have captured the impact of both stylistic dimensions as effectively as the text; in fact, images of war might have detracted from the tweet’s resonance by being too graphic. It is difficult to joke about such a serious topic, and evidence of the failure of overly graphic images is provided by the surprisingly unsuccessful tweet about the ‘English at play’ (Figure 6.16).

Even the second level of tweet success does not necessarily lead to user success, however. As Nick T. noted, the account that tweeted ‘Fuck Farage, fuck Rees-Mogg...’ received 10,000 retweets (despite only having 300 followers), but no additional followers. Similarly, one of the non-influential users who published an exceptionally popular tweet only gained about 100 followers in the two and a half years since its publication (increasing from 842 to 941); another has actually *lost* almost 300 followers (dropping from 629 to 351).⁸⁶ One-off extreme success does not have long-term impact—for this to happen, and for user influence to grow, a brand must be developed through the consistent production of creative and critical content.

7.2.2 *The Novelty of Niche Interests*

The second strategy for grassroots success is drawn from examples from the interview data. Scientists for EU, We Are Europe, and Labour Leave are non-influential campaigns that were exceptionally resonant. The feature that they have in common was already highlighted in Subsection 7.1.3—they target niche issues that are not represented in mainstream media, and thereby audiences who care about these issues (scientists and academics, millennials, and Labour Party members who support Leave). Niche issues are by definition novel, and these audiences are, importantly, active on social media. Because no previous campaigns had represented their interests and concerns, when these campaigns were launched, they had an instant critical mass of supporters. Over the course of one weekend, Scientists for EU organically acquired 2,500 Facebook followers, and 1,500 Twitter followers quickly ensued. This was enough to catch the attention of the *Times*, which further propelled their growth.

Nick T. represents a successful non-influential individual, and he has captivated a group of politically interested people who ‘live in some unfashionable bit of England’—the complete antithesis of the SW1 elite. His tweets (most of which link to his blog posts) do not focus on a specific niche issue, but they do provide an alternate perspective on political events that is not

⁸⁶ The follower counts were noted on 3 December 2018.

reported in mainstream media. The most successful of these tweets during the referendum period voiced a deeply uncomfortable truth about regional inequality. It was resonant because it captured exactly what his followers were thinking and feeling, but dared not express out loud. Nick has blogged and tweeted his alternate insights so consistently that some of his followers are ‘extremely rabid fans—people who just read literally everything’. If he publishes a book, 100 copies of it will be sold the first day because of this fanatical core.

The strategy, then, is to *find a niche, and represent its interests well*—or, alternatively, to systematically share insights that are not reported in the news media. This requires awareness of what is in the mainstream media; niche interests are only niche relative to mainstream interests. In a sense, the audience is already there—people use Twitter to obtain information that cannot be obtained from the news. This is one of Robert’s key motivations for being on the platform, and he follows subject matter experts to learn about underreported topics, such as the coup in Zimbabwe (as noted in Chapter 5, he began tweeting to serve as a source of such information in a different context: the military coup in Egypt). During the referendum, Robert used Twitter to find out about complex legal issues that were not widely discussed: ‘What kind of arrangements for the future are possible, what Article 50 really means’. Ironically enough, although Twitter is not known as a venue for substantive information, it is where specific details about the legal proceedings are divulged. Such details are not included in news articles by virtue of being too specific. It should be noted that journalists who represent mainstream media are also a subset of the experts who provide embellishing details on Twitter—the line between alternative and mainstream is blurred by representatives of news organisations. Nevertheless, one does not have to be affiliated with an organisation to be an expert on a niche subject.

7.3 The Measurement of Success

Two measures of success have been used in this thesis: direct retweets and likes. As appropriate as they are for capturing the concept of tweet resonance, this research has also illuminated their problems and limitations. Conceptually speaking, likes and shares are an impoverished operationalisation of success, as they do not capture *influence* over ‘what people think and how they act’ (Schudson, 1989, p. 153). As Chapter 3 has revealed, a whole range of other metrics are available to users about their own tweets that are potentially richer measures: link clicks, replies, and follows are but a few examples. These expanded metrics are not available from the public API, however. Retweets do capture *diffusion*, which is how success is typically defined in the information science and communication literature (Maybaum, 2013; Vosoughi, Roy, & Aral,

2018; Romero, Meeder, & Kleinberg, 2011). Yet a more accurate representation of diffusion is *reach* (measured by impressions)—how many people saw the tweet. Some interviewees preferred this measure over retweets, because they were only concerned with spreading their tweets as widely as possible. Although tweets that are shared more have more reach, reach is not necessarily directly correlated with retweets—it depends on *who* has retweeted. Like the expanded tweet metrics, impressions cannot be collected from the public API, so were not included in the analysis.

Another problem with retweets and likes is that strategies most effective for maximising them—e.g., adding videos, expressing vituperative anger, including clickbaity headlines—are not necessarily beneficial to pursue, because they damage an account’s reputation in the long term. They are at best crude measures of success that are focussed on short-term (i.e., tweet level, as opposed to user level) gain.

On an even higher level, retweets and likes are not always conducive for social good. In recent years, as Twitter has expanded in scale and significance, it is more often quite the opposite: a lot of misinformation is liked and shared; sensational news is more successful according to these metrics than well-researched news. Perhaps the strongest evidence for this is the recent MIT study of all controversial news stories on Twitter between September 2006 and December 2016 (Vosoughi, Roy, & Aral, 2018). The authors found that false news is far more shareable and likely to attain virality than real news, especially when the topic is political (and by implication divisive). To be more precise, tweets bearing false stories reach 1,500 people six times faster than tweets bearing true stories: ‘Twitter users seem almost to *prefer* sharing falsehoods’ (Meyer, 2018). As Lohr (2018) remarked, ‘Polarisation has turned out to be a great business model’.⁸⁷

Twitter itself recognises this, and has recently declared that it is ‘rethinking everything about the service to ensure [it is] incentivizing healthy conversation’ (Lunden, 2018). To do so, an understanding of what accounts for the spread and shareability of false news is necessary. Both context and content features are at play here. Bots represent the main context feature: they can easily take advantage of the system and its affordances by automatically tweeting or retweeting vast volumes of content, but many can be computationally detected (from, e.g., tweet source or the volume of original tweets vs. retweets). The content features are far more difficult to resolve, as they characterise real news as much as fake news: Vosoughi, Roy, and Aral (2018)

⁸⁷ Facebook, coincidentally, is facing an identical quandary: its advertising model favours controversial content, and thereby makes it cheaper; the main like and share engagement metrics incentivise vitriol (Martínez, 2018a and 2018b).

found that *novelty* and *emotion* explain why the latter is more engaging than the former. Moreover, novelty is almost impossible to detect without human understanding.

What, then, represents a better quality measure of success? One metric that is especially relevant given the expanded tweet character count (and introduction of threads) is the time taken to read the post. Borthwick (2015) expressed its rationale quite eloquently:

Unless we measure and value attention—time spent reading, listening, or viewing versus the raw click volume—we aren’t going to build things that are actually of interest to humans... the social networks are going to have to think about how to expose attention metrics and help publishers understand outside of the torrent of shares what are people actually reading.

The distinction between spending time on a tweet and sharing a tweet is clear: spending time is more directly connected to resonance, to what is *actually of interest* to the user. In the attention economy of social media, time is money—so lingering on a tweet is equivalent to paying more for a product. Sharing a tweet, on the other hand, does not indicate that it has been thoroughly digested (or even read at all). Indeed, many people retweet based on a cursory examination of tweet content; if they support the user, they might retweet on the basis of this support alone. As Andrew acknowledged, ‘I always agree with Nigel Farage, so I won’t even bother reading what Nigel Farage has said’. Twitter does not yet provide time spent metrics for tweets (either privately or publicly), but if it does, a far richer operationalisation of resonance will be revealed.

7.4 Implications

Three broad implications derived from the results of this thesis will be discussed in this section. The first, the impact of Twitter on the real-life referendum result, has been alluded to multiple times and is of significant practical interest. The second is concerned with the intersection between politics and marketing—in the new media landscape, the two are often inextricable from each other. The final implication is policy-related, and is derived directly from reflections made in the previous section about the manipulation of Twitter’s algorithm. There is a strong need to incentivise healthy conversation on the platform, and policy intervention might be necessary.

7.4.1 *Impact of Twitter on the Referendum Result*

Can tweets about the EU referendum be used to explain the Leave victory? Viz., Leave might have won the vote, but did they also ‘win’ on Twitter? Given the size and scope of the dataset collected for this thesis, these questions are both tempting and inevitable. As emphasised throughout earlier chapters, however, they cannot be answered from the data. Moreover, they

are not the focus of this thesis. The research questions are concerned with persuasion on Twitter in and of itself, as opposed to the campaign definition of persuasion. The former is connected to tweet engagement, whereas the latter is associated with changing opinions and affecting real-life behaviour (in this case voting decisions). There is a link between the two, but it is not direct or causal.

Why can the impact of Twitter on the referendum result not be determined from the data? Because the platform simply is not representative enough. Blank (2017) has demonstrated that in the UK, Twitter users are disproportionately elite—they are ‘younger, wealthier, and better educated than other Internet users, who in turn are younger, wealthier, and better educated than the offline British population’ (p. 679). Ironically enough, these are the very groups that were more supportive of Remain than Leave; as Mike observed, the entire higher education community was against Brexit. Another problem with representativeness stems from self-selection bias: those who tweet about the EU referendum are by definition more politically active than the average British citizen (Jungherr, 2016). These are two of the critiques that Gayo-Avello (2012) leveraged against electoral prediction from Twitter studies in his literature review of the topic; he affirmed that the predictive power of tweets has been grossly exaggerated. Other issues with the studies relate to methodology: most are *post-hoc* analyses as opposed to predictions, there is no standard way of ‘counting votes’ on the platform (tweet volume, number of unique users, and sentiment analysis are three different options), and there is also no standard way of measuring reality (polls and the vote itself are two different options). On top of all this, the effect of propaganda and misinformation is ignored. Jungherr, Jürgens, and Schoen (2012) directly demonstrate that election results cannot be predicted from Twitter data in their response to Tumasjan, Sprenger, Sandner, and Welpé’s (2010) attempt to do so.

All of this being said, general comments can be made about the relationship between Twitter and the referendum outcome. As Tiwari and Venkataramakrishnan (2018) contend, ‘social media alone cannot change the election result, but it may represent a swing in the barometer of the broader British digital mentality’. Given the high stakes and once-in-a-lifetime nature of Brexit, more citizens participated in online discourse during the referendum campaigning period than they did during campaigning periods for general elections. Although representativeness is still an issue and the analysis is necessarily *post-hoc*, it is interesting to see whether or not patterns of activity in the dataset reflect the real-life result. Moreover, in their analysis of a panel dataset of political tweets in the 2015 and 2017 UK general elections, Bright et al. (2017) conclude that campaigning on Twitter *does* help to secure votes, although its impact

is ‘small in absolute terms’. Other communication channels mediate the effects of Twitter, and hence the broader media ecosystem needs to be taken into account for predictive models.

Firstly, remarks will be made upon whether Leave or Remain published more resonant tweets. The population-level crosstabulation of the position variable with the binary version of retweet count reveals that Remain tweets were shared at slightly higher frequencies than Leave tweets, and that both categories were more often shared than ambiguous tweets (Appendix F). This result is contradicted by the regressions, however—only Leave is statistically significant, and in the context of all predictors the odds of pro-Leave tweets being shared are 42.5% *lower* than the odds of ambiguous tweets being shared. When retweet count is treated as a continuous variable, Leave tweets have a slightly higher mean than Remain tweets, and both have substantially higher means than ambiguous tweets (the means are approximately 5, 4, and 1; see Figure 4.17). What is even more interesting is the fact that when it comes to the second level of success, Leave is significantly weaker: 10 of the 11 most popular tweets in the dataset are pro-Remain.

Nothing conclusive can be observed from the tweet resonance metrics, partly because resonance can itself be measured in multiple ways. A better reflection of the British population’s voting inclinations might be volume of content. Leave tweets were more pervasive in the dataset collected for this thesis: half of the tweets in the sample (50.71%) support Brexit, 39.18% support Remain, and 10.11% are ambiguous. Population frequencies broadly reflect sample frequencies: 53.53% support Leave, 40.88% support Remain, and 5.59% are ambiguous. Even though many of the pro-Leave users might be bots, there is clearly more Leave activity on Twitter than Remain activity, and such a difference in activity might reflect the fact that Leavers were more effective overall in campaigning.

The qualitative data provides evidence that Twitter is more closely intertwined with real-life campaigning than Gayo-Avello (2012) suggests. Interviewees were convinced that the platform made a difference—they were unable to quantify what that difference was, but maintained that it had an impact. In Andrew’s words, ‘It’s an echo chamber. But it will have helped’. Sean vehemently affirmed that quantifying the impact was immaterial: ‘The stuff that we do, day in and day out, even if it’s not tangible—I think sometimes we get too obsessed with the numbers... It cannot *not* have had an impact of some sort’. This being said, he also noted that 1,263 people clicked on the register to vote link in the ‘Kiss of Death’ tweet, which reached an international audience through mass media coverage. The campaign hashtag #infor appeared 11,750 times across social media channels, and the campaign itself reached 10.5 million people despite being almost exclusively digital. These numbers are impressive, but it must be noted

that We Are Europe was supported by institutional funding and involved a large team, including three digital consultants. UKtoStay.EU, on the other hand, was a digital campaign that had no funding at all and consisted of three volunteers. During the final week of the referendum alone, they managed to reach an audience of 450,000 through their website, Twitter, Facebook, and email newsletters. Paul declared that although he did not know if any of these people were persuaded to vote Remain, the number of impressions in and of itself justified the campaign. Twitter played an important role in that it was where their content could be encountered by those who had not declared support for the campaign—many senior officials actually shared some of their graphics, which ‘must have had some sort of an impact’.

The impact of a tweet is often indirect, but its ultimate result can be significant. Matthew remarked that tweets were useful ‘in terms of shaping the media debate’, as they were picked up by broadcast and print journalists. Logan described an even more elaborate butterfly effect:

Maybe we looked very active on Twitter, and maybe journalists saw that, maybe they took YEM seriously, maybe they went to interview a YEM person later on and that then made a story, and people then read that story and became convinced, or decided to get involved and that person became an activist and convinced loads of people, and made us loads of votes.

Logan’s example is illuminating yet hypothetical. Brendan, on the other hand, described the concrete knock-on effect of Twitter for Labour Leave. It was not only a messaging tool, but also a practical tool for recruiting volunteers. Supporters would contact the campaign directly on Twitter because communication via the website was too time-consuming (tweeting at the campaign triggered a more immediate response due to phone notifications). Quite a few hundred volunteers were recruited in this way, and assisted with distributing leaflets or speaking at events. One particular volunteer from Liverpool coordinated other local volunteers to spread 200,000 leaflets across the entire city. Thus, although it was a small campaign, Labour Leave had a large impact through its vast network of supporters. Its impact was so large that David Davis, the former Secretary of State for Exiting the European Union, said on public record that Labour Leave ‘made the difference between defeat and victory for the Leave side’. Twitter was certainly an important part of their efforts.

Finally, it is interesting to make a contrast between two grassroots campaigns organised by the same individuals: Scientists for EU and Healthier In. Mike affirmed that the former was effective enough on Twitter to have some impact on the result; the latter was not. This is because Scientists for EU was launched early enough and acquired a large enough base of supporters to ‘own the territory completely’. Healthier In, on the other hand, was a response to Vote Leave’s NHS campaign. Its arrival was too late and its voice was too small for territorial ownership to be established.

All in all, the quantitative data have demonstrated that the referendum result is reflected (but not caused) by tweet volume metrics: the proportion of pro-Leave tweets in the sample and population datasets was over ten percentage points higher than the proportion of pro-Remain tweets. The qualitative data have shown that although Twitter's impact cannot be scientifically measured, it was an indispensable component of many campaigns.

7.4.2 *Politics and Marketing*

The second implication of this thesis relates to the overlap between political communication and marketing. In the context of a referendum campaign on new media, digital marketing and political communication techniques collide. Penney (2017) recognises this intersection in his *citizen marketer framework*, which is defined by peer-to-peer influence through the spread of content. As he contends, contemporary socio-political movements can be better understood through the lens of marketing theory. Ordinary citizens use the same techniques employed by professional marketers—the circulation of symbols through different media—to influence their peers. Framing citizens as marketers enhances the analysis of why some political messages become extremely successful, as it focuses on the persuasive dimension of these messages.

Indeed, marketing techniques are part of the framework established by this thesis for tweet resonance: resolution, questions, lists, contrasts, and literary figures (a component of playfulness) are all used in advertising. Andrew explicitly framed the tweet as a commercial object: 'You've got to have something that's worthy. You've got to have a product to sell'. What's more, the dependent variables in this thesis are measures of success in digital marketing frameworks. They are also metrics of commercial value to Twitter, as they reveal more information about its users. Thus, Marwick's (2013) contention that new media technologies teach users 'to be good corporate citizens... by harnessing marketing techniques to boost attention and visibility' is remarkably apt (p. 12). One must not forget that tweet-level success for the user is also financial success for the platform; Twitter is designed to amplify the effectiveness of already effective messages. Its business model and commercial incentives are inseparable from the study of its content.

To conclude this subsection, two qualifying observations will be made. Firstly, several interviewees cautioned that marketing techniques are not effective if they are too overt (e.g., a painfully obvious question or call to action). Balance is the key: if marketing techniques are too heavy-handed and not connected with high-quality content, they will backfire. Secondly, it is important to recognise that as much as Twitter is about community, it is also deeply focussed on the self (or the organisation). Self-marketing emerged as a subtheme from the interview data;

it is not lost in a campaign context. Some interviewees who managed organisational and personal accounts used the former to augment the latter. All of the interviewees who wrote news or blog articles promoted their pieces with multiple tweets. Because such a high concentration of journalists and politicians actively use Twitter, it is an indispensable networking tool for individuals from these communities. Sam's sobriquet for the platform captures this perfectly: 'LinkedIn for political nerds'.

7.4.3 *The Regulation of Platform Design*

Finally, some brief reflections shall be made upon one of the overarching implications of this thesis: given that the rhetoric of platform design is more persuasive than the rhetoric of content published on the platform, design must be altered in some way to encourage the production of higher quality content. The explosion of propaganda and misinformation on Twitter during the EU referendum and US presidential election in 2016 illuminates how conducive the platform is to spreading extremely distorted truths and outright falsehoods. These falsehoods are far more damaging to democratic discourse than any marketing technique—not just because they are false, but also because they are attention-arresting and highly sharable. Due to their resonance, they can spread far and wide even without being amplified by bots.

Of course, every media system has been manipulated and hacked (Borthwick, 2015). The difference now is one of scale and speed—because of the size and affordances of Twitter, falsehoods can spread faster and farther than ever before. What defines all social media platforms is that they allow for receivers to be producers. The increased scale of the network that this affords is at once its strength and its weakness.

A return to Chaffee and Metzger (2001) is appropriate here. They predicted that in the era of media communication (as opposed to mass communication), research paradigms would shift from the analysis of content and its effects on the audience to 'user interface design and information search strategies' (p. 373). This thesis has provided empirical evidence of the necessity for this shift—content cannot be viewed in isolation from contextual factors. How, then, might design be improved and regulated? This is connected to a larger-scale societal question that has recently been posed by Lazer et al. (2018): *'How can we create a news ecosystem and culture that values and promotes truth?'* (p. 1096; emphasis added).

It is clear that incentives have to shift. Retweets and likes could be replaced with healthier metrics of success, such as time spent on a post. As noted in the previous section, this metric is more directly connected to what is of interest to the user, and it does not facilitate the spread of misinformation. Perhaps visual content should *not* be favoured by the algorithm that

governs the Home feed, and perhaps there should be a broader range of emotional responses to messages, as there is on Facebook—liking and sharing can only communicate so much. Such changes might motivate users to focus on producing insightful rather than puerile content, and to be more creative with text. On a higher level, regulation can be put in place to disincentivise companies from creating algorithms that amplify the spread of content without checks and balances. The problem is about scale and speed, not about unhealthy content in and of itself—misinformation and propaganda have always existed, and always will exist. Depending on how they are designed, platforms like Twitter can either escalate or inhibit the popularity of such malicious content. There is an ethical element of design that is in tension with commercial imperatives, but in the post-truth aftermath of Brexit and Trump, this ethical dimension is increasingly being recognised.

7.5 Future Directions

There are numerous ways in which this thesis can be enhanced, and a select few will be described here. One immediate avenue to be explored is the use of rhetorical strategies in a smaller, more specific context: that of live events (e.g., debates, demonstrations, political programmes). The qualitative analysis has revealed that for this unique category of tweets, hashtag usage is essential, timing is especially important, and content features are also more significant. Another regression model on this subset of tweets would reveal what the effective strategies are during live events, and a comparison can be made to the effective strategies for the average tweet during the campaigning period (as illuminated in Chapter 4).

Another way in which this thesis can be refined is through modifications to the regression models. Given the qualitative findings, interaction effects can be explored between certain independent variables (e.g., negative framing and playfulness). This would illustrate whether there are combinations of rhetorical strategies that are effective, even if the strategies are ineffective on their own. Testing interactions of all of the content and style features with whether the user is an organisation would be particularly interesting.

Different approaches to modelling the data can be tested as well. Count regression models can be used, as they are specifically designed to account for the fact that the dependent variable cannot be less than zero, has to be an integer, and has no upper limit. What's more, distributions of count data are often positively skewed—many observations have a value of zero. Different models can be tested: the full distributions of retweet and like counts can be retained, or they can be collapsed into fewer categories: e.g., zero, one, two, and three or more.

To better understand the differences between receiving one like or share and receiving more, a more limited linear regression can be performed on the subset of tweets that attracted engagement.

Regardless of the model used, the timing variable can be further contextualised through a longitudinal, or time-series, approach. As campaigning strategies are always changing in response to real-world events (two prominent examples being Barack Obama's visit and the murder of Jo Cox), it is extremely important to take temporal dynamics and external information into account. As reaffirmed in multiple chapters, the question is not whether a specific strategy is effective, but rather under what circumstances the strategy is effective. Tweets need to be considered relative to the news media; incorporating data from the latter would help in the measurement of timeliness and novelty. This contextualised timing would almost certainly increase the amount of variation that the models can explain.

The measurement of content and style features at the population level can be enhanced through a more advanced SML for text classification approach. The algorithms deployed for this thesis utilised a traditional bag of words (BoW) model—the frequencies of specific words are counted, and the words are treated as independent units (one hot representations); tweets are represented as tf-idf vectors. This model is sufficient for the identification of position and topic, but the identification of features such as emotion, message framing, and playfulness are more complex and context-specific. As such, it would be fruitful to try an alternative and newer technique: *word embeddings*, which take semantic and syntactic relations between words (viz., context) into account. Words are mapped in a continuous vector space; those that are similar in meaning are closer to each other (e.g., 'good' and 'great'). The key advantage of this over the BoW model is that words that do not appear in the training data can be classified due to semantically similar words being present. Given the diversity of vocabulary and style represented by tweets, as well as the various ways in which tone can be communicated, this advantage is quite important. Rudkowsky et al. (2018) have demonstrated that word embeddings are a robust methodology for the classification of sentiment (negativity) in Austrian parliamentary speeches, and may improve BoW results. The state of the art library for word embeddings is Facebook's fastText,⁸⁸ which uses neural networks and has pre-trained models for 294 languages based upon Wikipedia data. Its embeddings are based on the skipgram model—each word is represented as a bag of character n-grams.

Finally, for future investigations of extreme success, it is essential to consider the role of gatekeepers via a network structure approach. Multiple interviewees affirmed that

⁸⁸ <https://fasttext.cc>

exceptionally successful tweets attained their high level of resonance not just because of their content, but also because they were shared by influential individuals. In their investigation of virality, Nahon and Hemsley (2013) illuminate the importance of gatekeepers and their control over information flow; they are central nodes in the network. Chadwick's (2013) concept of the hybrid media system is relevant here: journalists, celebrities, and institutional actors are involved in the dissemination of messages on social media alongside ordinary citizens. It is important to note, however, that there are two ways in which a tweet can go viral. The *shallow* route is represented by retweets from high-profile users; the *deep* route is represented by a chain of retweets from unpopular users (Vosoughi, Roy, & Aral, 2018). Both might ultimately reach the same number of retweets, but their network propagation patterns are different; the shallow route is much faster. It would be interesting to compare the types of tweets that attain virality via each route—'fake news', for example, becomes popular via the deep route (ibid). Are certain context or style features more associated with either? What does this illuminate about human behaviour?

7.6 Final Reflections

Twitter is a platform where opposites collide: Remain and Leave campaigners, humans and bots, individuals and organisations, mainstream and alternative groups. It has recently become associated with abusive content and propaganda, but it is also about honesty, authenticity, diversity, and exposure to insightful individuals. These qualities of the platform are best illuminated by comments from Harriet and Denis:

Twitter is my favourite out of any platform. It's great; it's where a lot of political discussion happens. It's also brutal, so you get a really good cross-section. If [you] want to know what people really wish they would say to each other in pubs, go on Twitter... it's where a lot of honesty happens. It's a good place for listening, but it's also a good place to kind of be.

For me it was like a public conversation. I know what I think when I've discussed it with other people, and Twitter is an odd shorthand way of being in touch with lots of very smart people... many much cleverer than me.

As such, even though it is not representative of the offline population, Twitter is the perfect window into how rhetorical content works in new media. It is hoped that this thesis can not only inform theories of reception, political communication strategies, and digital marketing approaches, but also serve to inspire more studies of 'big data' rhetoric that consider the fullest complexities of content. The intersection of natural language processing and qualitative analysis of social media text is extremely rich; much remains to be discovered through the synergies

between both strategies. As Ahmed Deedat once proclaimed, ‘Language is the key to the heart of people’. Social media provides an unprecedented opportunity to understand people at scale without the sacrifice of local nuance.

8. Appendices

Appendix A: List of Hashtags Used in Search Query on 22 June 2016

#2012gir2016, #23june, #23june2016, #43brokenpromises, #afterbrexit, #badbrexit, #bankerslovebrussels, #beaverin, #beleave, #beleaver, #belfastgo, #betonbrexit, #battereurope, #betterin, #betterineu, #betterineurope, #betteroffin, #betteroffout, #betterout, #bettertogether, #bolstertheborders, #borxit, #bremain, #brexin, #brexit, #brexit2016, #brexitandchill, #brexitbustour, #brexiteer, #brexiteffect, #brexitfacts, #brexitfantasy, #brexitfears, #brexitlies, #brexitnow, #brexitref, #brexitrisks, #brexitthefilm, #brexitthemovie, #brexodus, #brin, #britainfirst, #britainin, #britainineurope, #britainout, #britexit, #britin, #britinds4in, #britsdontquit, #brout, #businessgo, #byebyeuu, #camforeu, #campaigntoremain, #cats4brexit, #cats4britain, #catsagainstbrexit, #catsforbrexit, #cloutnotout, #conservatives4britain, #conservativesforbritain, #conservativesin, #davesdodgydeal, #dogs4brexit, #dogs4britain, #dogsagainstbrexit, #dogsforbrexit, #dontfuckmyfuture, #dontwalkaway, #englishexit, #eudebate, #eudebateni, #eudebates, #eudecide, #eugood4uk, #euin, #euinorout, #eustheproblem, #euleave, #eunegotiation, #euout, #eupoll, #euref, #euref16, #euref2016, #eureferendum, #eureferendum16, #eureferendum2016, #eurefin, #eurefni, #eurefout, #eurefready, #euremain, #eurenegotiation, #euref, #eurosceptic, #euroscepticism, #eurosceptics, #eustay, #exiteu, #fishingforleave, #flexcit, #fucktheeu, #getbritainout, #goactionday, #gobritain, #gogogo, #goodbyeuu, #gosupersaturday, #grassroots_out, #grassrootsout, #grassrootsoutgo, #greenerin, #healthierin, #hearothervoices, #hugabrit, #imleavebecause, #incampaign, #infor, #inorout, #inout, #intogether, #jun23, #june23, #junereferendum, #justnotintoeu, #labourgo, #labourin, #labourinforbritain, #labourleave, #leadingnotleaving, #leadnotleave, #leave, #leavealliance, #leavechaos, #leaveeu, #leaveeuofficial, #leaveeurope, #leavetheeu, #leaveuk, #leavingeu, #leftleave, #letstakebackcontrol, #lexit, #london4europe, #loveeuropeleaveeu, #lsebrexitvote, #makebritaingreatagain, #mutts4remain, #mydailybrexit, #no2brexit, #no2eu, #nobrexit, #noeu, #notbetteroffout, #nothankseu, #notintoeu, #notobrexit, #notoeu, #otherref, #outcampaign, #outeu, #pets4brexit, #pets4britain, #petsagainstbrexit, #petsforbrexit, #petsforbritain, #pleasevoteleave, #proeu, #proeurope, #proeuropean, #projectbullshit, #projectcheer, #projectfact, #projectfantasy, #projectfear, #projecthope, #projectidiocy, #projectintimidation, #projectnasty, #projectpanic, #projectreality, #projectsmear, #projecttruth, #projectwhinge, #referendum, #referendum2016, #referendumdebate, #remain, #remainandgain, #remainer, #remainers, #remaineu, #remainiac, #remainiacs, #remainian, #remainians, #remainin, #remainineu, #risksofstaying, #saferout, #sagedebateeu, #sageeudebate, #sayno2eu, #saynotoeu, #sayyes2eu, #sayyestoeu, #standupforeurope, #stayeu, #stayin, #stayineu, #stayineurope, #stoptheeu, #strongerin, #strongerout, #studentsgo, #studentsin, #takecontrol, #theleavealliance, #uk4eu, #ukandeu, #ukin, #ukineu, #ukineurope, #ukinorout, #ukout, #ukouteu, #ukref, #ukreferendum, #uktostay, #unisforeu, #viewsoneurope, #vote2stay, #vote_leave, #votebeaver, #votein, #voteleave, #voteleaveeu, #voteout, #voteremain, #votestay, #votetoleave, #votin, #weareeurope, #wewantourcountryback, #wewantourfishback, #whyvoteleave, #womenforbritain, #womenin, #yes2eu, #yeseu, #yestoeu, #yourvoteeurope

Appendix B: Examples of Spam and Unrelated Tweets

Prominent categories of these tweets include automated news digests from Paper.li, #InFor being used by software companies, music-related messages, sex-related messages, Britain First/British National Party content, and #bettertogether being used in other contexts.

1. Crash Kings - Mountain Man <https://t.co/6gIXnAueLv> #TakeControl
2. Moving to #Infor #Cloud saves Concordia Plan Services \$400,000 in annual operating expenses. #ensw <https://t.co/QeHFGChOzW>
3. I'm so drained these days. Wish to have a webcam chat or meet? Find link in bio. #banen #business #voteleave #fcbvbv <https://t.co/4KzUuLj0aU>
4. Quick someone check @ChuckGrassley's record on taking \$\$\$ from the @NRA #enough #voteout #NRApuppets #DisarmHate #senatevote
5. ? I'm very hot tonight and ? I want to show you my boobs! ? #brexitornot #remain
6. . . ready for the event. #anifest2016 #event #weekend #pfear #projectfear #junkyarts... <https://t.co/PjHqxYboiJ>
7. When people try & do cross fit @ the Alaska club... #StopNow #Please #Leave
8. #Atlanta WE LOVE YOU!! #TeamPOYNTE #POYNTE #TakeControl #DiscreetEnemy #WeAreComing... <https://t.co/0lx3cWAYVH>
9. Only TWO PER CENT Of Men Aged 18-24 Feel Masculine <https://t.co/iz4eslq4jX> #LOL #MakeBritainGreatAgain
10. We're in the middle of the transition to the 4th Industrial Revolution right now... #HootConnect #bettertogether
11. Pick a card, any card... and you'll get world class service from any of our account managers! #BetterTogether <https://t.co/41zaTFxm1U>
12. The One Thing You Should Never Do With Your Boarding Pass <https://t.co/Sq5ejdy7pF> #BetterTogether via @HPLifestyle
13. Ready for #bettertogether block party!! #churches #partners... <https://t.co/6eiyyI4pEB>
14. All UK pensioners deserve taxpayer funded speed & quality of treatment as #mrsWindsor #birthday #BirthdayBash #republic #Grin #Remain
15. California Increases Amount #Paid for Family, #Disability #Leave <https://t.co/X70VRRzNIA> #FMLA #Insurance #Healthcare
16. The latest The CrossBrum Daily! <https://t.co/FOqrje2ilu> Thanks to @businessdeskwm @Birmingham_UK @HABHousing #strongerin #euref
17. What You Eat Affects Your Sleep #takecontrol #sleepisgood #healthylifestyle <https://t.co/8hdT3yYUUK>
18. #RADIO #70s #juice Now Playing #Stayin' Alive #Bee Gees #APPS Apple <https://t.co/RLfj5bTKBk> Google <https://t.co/6PBywmglVI>
19. Bernadette and Mitzi never lets a guy leave without sex: #Bernadette and Mitzi #lets #never #without #sex #leave <https://t.co/je9U9HP4tV>
20. Volbeat - The Hangman's Body Count <https://t.co/6gIXnAueLv> #TakeControl

Appendix C: Tweet Coding Scheme

Coding Procedure

1. Copy and paste the tweet URL (column C) into your browser and examine the *full context* of the tweet:
 - Scrutinise any images and animated GIFs that are included.
 - Watch videos in full if they are embedded in the tweet. If there is a link to an external video (e.g., on YouTube), review the title and description and skim the content of the video—you do not have to watch it in full.
 - If the tweet contains a URL (e.g., to a news article, blog post, campaign website, etc.), carefully inspect the snippet of the website that is embedded in the tweet. Then click on the URL and skim the content to get a sense of the tweet's meaning and context—you do not have to read the entire article or blog post.
 - If the tweet cannot be viewed on Twitter because a) it has been removed, b) the user's account is protected, or c) the user's account has been suspended, take a look at the tweet text (column D). Supplement this by trying to view any URLs (column AA) and/or media URLs (column AC) that might be included, and review the username (column Q), screenname (column R), and user description (column S), if included. The more context, the better!
2. Code for the variables described below, adding your label directly to the spreadsheet in the relevant column (E-O). Please write the label exactly as it appears in the scheme below (capitalise the first letter and use lowercase for the rest: e.g., 'Remain' as opposed to 'remain'; 'Other' as opposed to 'other').
3. If you feel the need to justify any of your coding decisions, or if you had difficulty making them, use the Notes column (P) to add any comments that you feel are necessary. You can, e.g., describe the content/perspective of a news article that the tweet links to, or explain why you felt one emotion was more appropriate than another.
 - If you label issue as 'Other' please add details in the Notes column (i.e., what you think the category might be).
 - If you label position as 'Ambiguous', please explain why in the Notes column.
4. **Do not overthink—take tweets at face value!** If it's not explicitly stated or strongly implied (after taking the full context into consideration), do not assume it is there, and do not code for it. This applies especially to playful tweets, and will save you both time and mental energy. Imagine that you are reading this tweet on Twitter back when it was posted, and go with your gut instinct.

Note: example tweets are provided for all of the coding categories. Specific components of the tweet that were helpful in determining which category it belongs to are **bolded**.

1. Position [Column E]

Which side of the referendum debate does the tweet support?

- Leave
- Remain
- Ambiguous
- N/A (not related to the EU referendum)

Guidance Notes

- In general, the position expressed by the tweet is the position espoused by its author. User information (username, screenname, description) can and should be used to verify the position that seems to be expressed by the tweet, if it is unclear.
 - The user description (column S) is especially useful. Sometimes users explicitly indicate the side they support in the referendum campaign in their description.
 - Note that as it has been almost a year since the referendum, many users have probably changed their descriptions (and/or even their usernames). If this is the case, please look at the description and username *in the CSV file*.
 - If you are able to view the tweet on Twitter, you can also examine the user's profile picture and background picture.
 - If the author's position seems to be completely at odds with the position expressed by the tweet, code according to the *position expressed by the tweet*.
- The usage of a hashtag that clearly indicates a side (e.g., #VoteLeave) does not in and of itself indicate that the tweet supports that side, as it might have been appropriated by campaigners from the other side. The full tweet text and any media contained need to be scrutinized.
 - Generally speaking, however, if there are many hashtags in a tweet that relate to a specific side, even if the content itself is neutral the tweet can be deemed to support that side.
- Try to avoid coding 'Ambiguous' if there is even the slightest hint of partisanship. If you do use this code, please explain why in the Notes column. Users who are undecided, unsure about which side they support, or upset with both sides definitely fall into this category.

EXAMPLES		
Position	Tweet URL	Tweet Text
Leave	https://twitter.com/brexitmarch/status/743403615985991680	European? British? These 'Brexit' Voters Identify as English https://t.co/SDPggnFc7w #BREXIT #StrongerIn #No2EU #EUref #LeaveEU #VoteLeave
Leave	https://twitter.com/M0nk33y/status/743714998732500993	Sweden prepares for Brexit as EU minister slams David Cameron's debate performances http://shr.gs/Kl2z46X #DavidCameron #VoteLeave #EUref
Remain	https://www.twitter.com/GrahamLawler/status/740161292573544448	If you don't have your national insurance number, you can still register to vote. Don't miss out: http://www.theguardian.com/politics/2015/apr/16/do-i-need-my-national-insurance-number-to-register-to-vote ... #StrongerIn

Remain	https://www.twitter.com/nprarobinson1/status/745341614617407493	Wish the crowd on #VoteRemain would shout and clap a little louder in there It's making me nervous how loud #VoteLeave are 😞 #BBCDebate
Ambiguous	https://twitter.com/AngryBritain/status/746116957355778048	Orkney Islands vote to #Remain, total votes cast 6, all related #EUref
Ambiguous	https://twitter.com/marcuschown/status/740245629574762497	Our most untruthful Prime Minister warning us about the untruths of the #Leave campaign! Bloody hell. Talk about muddying the waters.

2. Organisation [Column F]

Is the tweet from an organisation (e.g., BBC News, Vote Leave, Virgin) or from an individual representing an organisation (e.g., David Cameron, Will Straw, Richard Branson)?

- 0
- 1

Guidance Notes

- 0 represents non-organisation; 1 represents organisation.
- The word ‘organisation’ encompasses the media, politicians, campaign groups, political parties, and companies/corporations/brands (as opposed to private individuals). All of these entities exert a significant degree of influence upon the public; they have authority and reputation.
 - Note that any individual who represents an organisation should also be considered as an organisation: e.g., the CEO of a public company, Members of Parliament (MPs), Members of the European Parliament (MEPs). The only exception to this is if the individual adds a phrase along the lines of ‘all views my own’ in his/her description—in this case, the individual should *not* be coded as an organisation.
 - Politicians who are below MP level (e.g., Councillors) are *not* organisations.
 - Most, but not all, organisations are verified on Twitter (column T should be ‘true’; on Twitter there is a blue badge with a white check mark next to the username). An account does not have to be verified to be labelled as an organisation, but verified status strongly suggests that an account is an organisation.
 - A high number of followers (anything above 10K) also indicates that an account might be an organisation. Accounts with less followers might still be organisations, however.
- If you are unsure about whether an account is an organisation, try searching for the username on Google. Review any websites and/or press coverage that relate to the account.
 - A detailed Wikipedia page is a sign that the account is an organisation.
- As it has been almost a year since the referendum, many users have probably changed their descriptions and/or their usernames, as noted above. If an account has been renamed or rebranded, code with the user information available in the CSV file (as that is relevant to the tweet at the time it was sent), and make a note in the Notes column that the account has since been rebranded.

- Here is a list of organisational Twitter handles that relate to the referendum (it is not exhaustive):
 - Arron_banks, BeLeaveBritain, BeLeaveUK, BetterOffOut, BorisJohnson, bowgroup, BrexitE, BrexitTheMovie, BrugesGroup, Business4GO, CAEFtweet, ConnollyAssoc, consforbritain, DanHannanMEP, DanielJHannan, darrengrimes_, DemocracyMovemt, DouglasCarswell, duponline, EconForBrexit, eureferendum, ExittheEU, fishingforleave, forbritain, FreeBritain, georgegalloway, GetBritainOut, GlobalBritain, Grassroots_Out, greenleaves2016, jdwtweet, JimAllister, Labour4Go, labourleave, lawyers4britain, LeaveEUOfficial, LeaveHQ, Left4GO, LGBT4GO, matthew_elliott, michaelgove, Muslims4Britain, newcastlefirst, NewsUK, Nigel_Farage, Nireland4GO, PeoplesPledgeEU, RichardAENorth, RMTunion, ScienceBritain, Scotland4GO, Steel4GO, Stop_The_EU, SuzanneEvans1, tfa4freedom, TUSCoalition, UKIP, ukleave_eu, Uni4GO, utopiexperiment, Vapers4Britain, VeteransBritain, vote_leave, Vote_LeaveMedia, WAGTV_, WofEUUnion, WomenforBritain, Y_Eurosceptics, Adecco_UK, Another_Europe, Avaaz, BNE_Group, Brand_EU, BritInds4IN, britinfluence, CamForEurope, CarolineLucas, cityoflondon, ClientEarth, CommunityUnion, congroup4europe, ConservativesIN, CRAEDO, CWUnews, David_Cameron, E3G, euomove, euomovescot, fbunational, FedTrust, George_Osborne, GlobalJusticeUK, GMBPressOffice, HealthierIn, jeremycorbyn, jon_bartley, labour4europe, LawyersIn4B, lebara, LibDems, London_First, London4Europe, lsebrexitvote, NewEuropeans, NFUtweets, NicolaSturgeon, Number10gov, Plaid_Cymru, pleasedontgouk, ProEuropa_GB, ProgressOnline, RCorbettMEP, RemaininEU, RepSocialist, res_publica, Scientists4EU, scotgp, SDLPlive, SMFthinktank, stayinEU, StrongerIn, strongerinpress, Students4Europe, The_TUC, TheGreenParty, theSNP, UCATTUnion, UKLabour, UKLabourIN, unisontweets, unitetheunion, UsdawUnion, Virgin, wdjstraw, weare_europe, weareeuropeuk, wwwfoecouk, BrexitWatch, foreignoffice, PollstationUK, wef, BBCBreaking, BBCNews, BBCRealityCheck, BloombergTV, business, CityAM, Daily_Express, daily_politics, Daily_Star, DailyMailUK, DailyMirror, EveningStandard, FT, guardian, Independent, indy100, MailOnline, MetroUK, NewStatesman, POLITICOEurope, SkyNews, spectator, standardnews, Sunday_Mail, Telegraph, TheEconomist, thepaper, TheSun, thesundaypeople, thesundaytimes, thetimes, WSJ

EXAMPLES		
Organisation	Username & Profile URL	Profile Description
1	Ben Macpherson MSP https://twitter.com/BenMacpherson	Local MSP for Edinburgh Northern & Leith (@theSNP). Passionate about justice & progress. Scot & internationalist. Email: ben.macpherson.msp@parliament.scot.
1	Woking Stronger IN https://www.twitter.com/Woking_IN	#Woking is stronger, safer and better off in Europe. Don't sit on the sidelines - be part of the campaign to secure our future in Europe.
0	Steve Unwin https://www.twitter.com/SteveUnwin01	North Dorset UKIP Parliamentary Candidate 2015. Former Conservative Party Agent @LeaveEUOfficial @Grassroots_Out #Brexit #No2EU #LeaveEU RT not an endorsement
0	RAJAT DATTA https://www.twitter.com/Anderson_d_king	Resident Doctor, Dermatology; Nationalist, Right Wing Conservative, Veryfond of Analysing Political Trends

3. Rhetoric (Message Type) [Column G]

Which rhetorical strategy is evidenced in the tweet to persuade the audience?

- Identification
- Division
- Other

Guidance Notes

- To code tweets for this variable, please keep in mind the position of the user who is tweeting. If the position is ambiguous, the rhetorical strategy must be ‘Other’.
- *Identification*: appeal to shared beliefs, values, attitudes, interests, or memberships (e.g., ‘We stand for democracy’). Attempt to establish common ground with audience.
 - E.g., any positive statements made about the EU by a Remain campaigner, or about sovereignty / the UK by a Leave campaigner.
 - Inclusive pronouns are often indicators of this category (we, our).
 - Tweets about pro-Remain campaigners from pro-Remain accounts and pro-Leave campaigners from pro-Leave accounts are obvious examples of this category (many contain images of people holding flyers or banners).
- *Division*: critique of the opposition; creates a ‘we-they dichotomy’. In the context of the EU referendum, this can be expanded to encompass *anything negative said about anyone or anything related to the opposite side*. This includes—and is not limited to—people, institutions, campaigns, and policies.
 - E.g., a migrant slur or anti-EU statement from a pro-Leave user; an anti-UKIP or anti-nationalist statement from a pro-Remain user.
 - Any tweet about the risks or repercussions of leaving the EU from a pro-Remain user or staying in the EU from a pro-Leave user.
 - Name-calling is a prominent example of this category.
 - Includes tweets that *compare* both sides (implicitly critiquing the opposition). I.e., if the tweet contains both identification and division, label it as division.
- *Other*: tweets that contain neither identification nor division, as defined above. Do not be afraid to use this category if you think it is the most appropriate—do not try to force a tweet into being identification or division. Many tweets use other rhetorical devices.
 - If you think you can identify what the other rhetorical device is, please make a note of this in the Notes column.

EXAMPLES		
Message Type	Tweet URL	Tweet Text
Identification	https://www.twitter.com/dianessere/status/744678769361813505	RT @wdjstraw: Wonderful to be part of the grassroots #BigIN today in Hyde Park #StrongerIn #VoteRemain https://t.co/q4jJfm5nd
Identification	https://twitter.com/GrassrootsOX28/status/728978916019425280/photo/1	#Brexit #eureferendum #euref #leaveeu #voteleave #goactionday ANOTHER ONE FOR #WITNEY 🙌🙌🙌
Identification	https://twitter.com/jk_rowling/status/742718884696993793	I don't want to leave a union that has brought more peace and stability to the region than it has ever known #Remain
Division	https://twitter.com/owensheers/status/745595514985545728	Why I'll be voting #Remain tomorrow For @WalesStrongerIn. Let's put our shoulder to the wheel, not walk away .

Division	https://twitter.com/aveen_geologist/status/745139594308378625	#Mutts4Remain Cantona is worried he will get deported back to France #VoteRemain
Division	https://twitter.com/MissMossey/status/739188246526201857	Don't let the right-wing swivel-eyed loonies erode our rights. #EUreferendum #Remain #StrongerIn
Division	https://www.twitter.com/MercianSaxon/status/731082043745030144	RT @LeaveEUOfficial: '@GiselaStuart: When I talk about immigration , I talk about the fact it needs to be managed . #LeaveEU #Brexit
Other	https://twitter.com/wdjstraw/status/746024134602293248	Lost your polling card? No problem, use this tool to find your polling station : http://www.wheredoivote.co.uk #StrongerIn
Other	https://twitter.com/pavilionopinion/status/726172850005118976	History of Europe: War War War War War War War Arguments about bananas. To be honest, I'll probably go with banana arguments. #remain
Other	https://twitter.com/rowbar/status/745872885533843456	Germans will reserve sunbeds for Brits w their towels AND admit 1966 goal was legit say @BILD, if UK votes #remain!

4. Message Format [Column H]

Does the tweet contain any syntactic devices?

- 0
- 1

Guidance Notes

- 0 represents the absence of the syntactic devices defined below; 1 represents the presence of one or both of the devices.
- *Contrast*: the tweet contains an explicit contrast in words, meaning, or both. I.e., two words or ideas that are opposites or strikingly unlike each other are juxtaposed, usually in parallel form.
 - This device is also known as antithesis. It is a very general category.
 - Can be the juxtaposition of *media* (image, video, GIF) and text, or of two media items. E.g., an image of an interracial couple with the text 'We're voting for #Brexit because we're clearly racist', or an image of the UK before Brexit next to an image of the UK after Brexit. It is very important to examine the full context of the tweet.
 - Many images contain text that contains contrasts. Please inspect the textual content of images embedded in tweets.
 - As Grady and Potter (1985) assert, 'Political messages conveyed with the use of contrasts are naturally emphasized because, in effect, the core assertion is normally made twice—in a "positive" and a "negative" form' (p. 173).
- *List*: must contain at least three elements (can be of any length provided this condition is met). The elements can be the same (e.g., 'war, war, war').
- *List-Contrast*: both devices are present in the tweet.
- Do not be alarmed if most tweets do not contain these syntactic devices—they are adapted from Atkinson's (1984) study of transcripts and recordings of large-scale

political party rallies in the UK. He discovered that the two most dominant formats in political speeches are lists and contrastive pairs. They might not be as applicable to tweets, given their extreme brevity compared with speeches.

EXAMPLES		
Message Format	Tweet URL	Tweet Text
1 [Contrast]	https://twitter.com/owensheers/status/745595514985545728	Why I'll be voting #Remain tomorrow For @WalesStrongerIn. Let's put our shoulder to the wheel, not walk away.
1 [Contrast]	https://twitter.com/billybragg/status/743408035167633408	This image sums up why I voted #Remain: Not every Leave voter is a racist, but every racist will vote Leave.
1 [Contrast]	https://twitter.com/BrexitFTW/status/746029557208260609	#iVoted for the UK to leave EU, Britain never stands alone, Britain stands proud. #VoteLeave #Brexit #UK
1 [List]	https://twitter.com/BBCBreacking/status/746177461117943808	"A victory for real people, for ordinary people, for decent people " - Nigel Farage, #Leave http://bbc.in/264XBZu
1 [List]	https://www.twitter.com/TishWillsy/status/745739051559428096	RT @BradenDavy: Imagine how you'll feel if you wake up on Friday to see remain. Gutted, disappointed, sad and depressed. Don't regret your...
1 [List-Contrast]	https://twitter.com/pavilionopinion/status/726172850005118976	History of Europe: War War War War War War War Arguments about bananas. To be honest, I'll probably go with banana arguments. #remain
0	https://twitter.com/chrissharts/status/745909346123587584	Leaving the EU would be like Man Utd leaving the Premier League and going for a kick about down the park #Remain
0	https://twitter.com/DonaldTrumpWall/status/746011578722553856	#VoteLeave before #England turns into a 3rd world hellhole. #Muslims are infiltrating the #UK from within. #Brexit
0	https://twitter.com/TerryGDavis7/status/739379310851723264	When the EU army replaces National Armies which EU state will they send them into in the greater interest of the EU if they vote to #leave?

5. Question [Column I]

Does the tweet contain a question?

- 0
- 1

Guidance Notes

- 0 represents the absence of a question; 1 represents the presence of a question.
- The question does not have to be rhetorical—and indeed, in the context of Twitter it is often extremely difficult to judge this.
- This variable is extremely straightforward, but be careful: *the presence of a question mark in and of itself does not indicate that there is a question in the tweet.* It could represent an emoji, which cannot be displayed in the CSV due to Unicode translation errors. Always check

the tweet on Twitter if possible; if the tweet is no longer available on Twitter, use context to judge whether or not there is a question.

EXAMPLES		
Question	Tweet URL	Tweet Text
1	https://twitter.com/KKarkagiannis/status/740424203233570817	A solid opinion article in favour of #Remain by C Patten (but why this animosity towards referendums? Aristocracy?) https://t.co/PsIJkr5hOH
1	https://twitter.com/DAZZA3201/status/745690738047000576	My lads home from Uni so I asked how he was voting? He didn't know but his Uni mates all voting In. Today he said they're all voting #Leave
0	https://twitter.com/Fight4UK/status/724320985080553472	#Brexit ~ #GO ~ #LeaveEU ?? #TTIP < It's what the working class must avoid. If you are a Labour voter who's voting for it, shame on you.
0	https://twitter.com/MikeHPR/status/740528852687032320	In the space marked 'vision', left vacant by @vote_leave, there's now this ? ? ? #VoteStay #StrongerIn https://t.co/CPK0ztFFHe

6. Emotion [Columns J-L]

Does the tweet express enthusiasm? [Column J]

- 0
- 1

Does the tweet express anxiety? [Column K]

- 0
- 1

Does the tweet express aversion? [Column L]

- 0
- 1

Guidance Notes

- 0 represents absence of emotion; 1 represents presence of emotion.
- *Enthusiasm*: happiness, hope, pride, determination, excitement, joy; the tweet is optimistic and conveys general positive sentiment.
- *Anxiety*: fear, unease, worry; tweet might frame a dangerous or troublesome problem or threat.
- *Aversion*: anger, outrage, hatred, contempt, disgust; tweet might frame 'others' as agents of a problem (the 'other' can be an individual, a group, or an institution such as the European Union).
- These emotions are *not* mutually exclusive: a tweet can contain up to all three, or any combination of the three.

EXAMPLES		
Emotion	Tweet URL	Tweet Text
1 [Enthusiasm]	https://www.twitter.com/Woking/status/745902257774931968	Raining? Don't let the weather get in the way of you voting in the most important vote in a generation. #VoteRemain
1 [Enthusiasm]	https://twitter.com/BBCBreaKing/status/746177461117943808	"A victory for real people, for ordinary people, for decent people" - Nigel Farage, #Leave http://bbc.in/264XBZu
1 [Enthusiasm]	https://twitter.com/MissLauraMarcus/status/743110036491276288	"Remain? Hell yes! Because There Is No Alternative " by @MissLauraMarcus on @LinkedIn https://www.linkedin.com/pulse/remain-hell-yes-because-alternative-laura-marcus ... #EURef #Remain #VoteRemain
1 [Anxiety]	https://twitter.com/SJPowell/status/742829639089823744	Don't risk the NHS! #VoteLeave to avoid #TTIP privatisation
1 [Anxiety]	https://twitter.com/MrB47351012/status/724901989684023297	# Turkey Threatens EU Over Armenian Genocide Art Project http://bit.ly/1UefDBh #LeaveEU #Brexit #GO
1 [Anxiety]	https://twitter.com/Os5i3/status/743372785779679232	# Brexit fears spark investor sell off! #BigMoney buy cheap & make a killing, who's running the #Leave campaign?
1 [Aversion]	https://twitter.com/grahnort/status/745619814094348288	Complicated debate obvs but surely safe to say that if Farage Johnson and Gove think something is a good idea we can assume it isn't #remain
1 [Aversion]	https://www.twitter.com/lehain/status/738830745708400640	If you're complaining you're losing. #VoteLeave https://t.co/JtgOut5Zuv
1 [Aversion]	https://www.twitter.com/davidwebb27/status/735369549642649601	RT @Patriotic_Brit: If anyone from #remain would like to see how George & Dave are fooling them & why tweet me. We're not #strongerin #voteleave
0	https://twitter.com/barrymc2010/status/745727745095835648	#Leave want Brits to choose between different rules, EU or not, but won't specify what post-#Brexit rules would be.
0	https://twitter.com/wdistray/status/746024134602293248	Lost your polling card? No problem, use this tool to find your polling station: http://www.wheredoivote.co.uk #StrongerIn
0	https://twitter.com/StrongerInTheEU/status/743870164706398208	Have recent events changed your #euref voting intention? #StrongerIn #VoteLeave #Brexit #LeaveEU #Labour #UKIP

7. Resolution (action-orientation) [Column M]

How does the tweet relate to action? Does it tell the audience how to respond by calling for a specific action (e.g., read a blog post, watch a video, sign a petition, attend a rally, retweet this message)?

- 0
- 1

Guidance Notes

- 0 represents absence of resolution; 1 represents presence of resolution.
- Resolution can be broadly conceived as a call for action of any kind (from watching a video to voting). It can be anything from a suggestion to an outright demand.
- Examine the *verbs* in the tweet: are there imperatives? To what extent does the tweet command the audience to do something?
 - Resolution must be *explicit*: a statement of a voting intention or a picture of campaigners does not in and of itself contain any resolution.
 - Words and phrases such as ‘must’, ‘please’, ‘ASAP’, and ‘now’ are clear indications of resolution.
 - Any tweet with ‘please RT’ should be coded as ‘1’.
 - Note that sometimes hashtags representing campaigns that contain verbs in their names are used as verbs in the tweet—in which case they increase its resolution. #VoteLeave is a good example. If they are being used as topic tags, they do not increase resolution.
- Tweets that are informational or descriptive should be coded as ‘0’.

EXAMPLES		
Resolution	Tweet URL	Tweet Text
1	https://twitter.com/SJ_Powell/status/742829639089823744	Don't risk the NHS! #VoteLeave to avoid #TTIP privatisation
1	https://www.twitter.com/only1colenzoit/status/742695455201611777	#VoteRemain and the wipe the smile of Rupert Murdoch's face says @jonskewes
1	https://www.twitter.com/Sec_Socialism/status/739190501870567428	RT @MissMossley: Don't let the right-wing swivel-eyed loonies erode our rights. #EUreferendum #Remain #StrongerIn https://t.co/exuJXhdcWU
1	https://twitter.com/obviousreally/status/734363652745203712	Must watch #brexitthemovie https://www.brexitthemovie.com/ #voteremain #voteleave #brexit #takecontrol #beleave RT a.s.a.p. please!
1	https://www.twitter.com/TishWillisy/status/745739051559428096	RT @BradenDavy: Imagine how you'll feel if you wake up on Friday to see remain. Gutted, disappointed, sad and depressed. Don't regret your...
1	https://twitter.com/Hunter7Taylor/status/744459194053627904	BetterOffOut: Follow your heart, vote for democracy and hope. Vote Leave: http://www.telegraph.co.uk/news/2016/06/18/michael-gove-interview-vote-for-democracy-and-hope-vote-leave/ ... by Michael Gove via Telegraph #VoteLeave...

0	https://www.twitter.com/maggie_ms/status/745541964976644097	RT @Sue_Cowley: I asked the kid last night how she would vote if she could. She said "Duh! I'd vote to stay in of course, mummy." #Remain
0	https://twitter.com/M_Pope1/status/741699561312419840	The wife and I are both voting to leave. Yes we are racist and yes we are xenophobic..... of course. #VoteLeave
0	https://twitter.com/SE_StrongerIn/status/744536580849827842	Huge team out in #Oxford today knocking doors! #StrongerIn #InEverywhere

8. Playfulness [Column N]

Does the tweet play with language? Is there a discrepancy between literal and intended meanings?

- 0
- 1

Guidance Notes

- 0 represents absence of playfulness; 1 represents presence of playfulness.
- Playfulness encompasses broad categories of rhetorical figures: irony, sarcasm, metaphor, hyperbole, pun, personification, and paradox are some of the most common in political language. In general, there is a deviation from expectation (conventions of expression). The general effect is often humour or wit.
- The scrutiny of media items (images, videos, GIFs) is critical for the assessment of this variable—often playfulness is the result of the juxtaposition of text and media.
- Clichés should not be coded as playful, as they do not violate expectations.

EXAMPLES		
Playfulness	Tweet URL	Tweet Text
1	https://twitter.com/chrischarts/status/745909346123587584	Leaving the EU would be like Man Utd leaving the Premier League and going for a kick about down the park #Remain
1	https://twitter.com/MissMosley/status/739188246526201857	Don't let the right-wing swivel-eyed loonies erode our rights . #EUreferendum #Remain #StrongerIn
1	https://twitter.com/aveen_geologist/status/745139594308378625	#Mutts4Remain Cantona is worried he will get deported back to France #VoteRemain
0	https://twitter.com/ThomasEvansGB/status/738491123774132225	#BBCQT No viable alternative outside the #EU says Liz Truss MP The #Commonwealth is projected to overtake the EU economy by 2019. #VoteLeave
0	https://twitter.com/alexandroach1/status/745340195017474048	Wish the crowd on #VoteRemain would shout and clap a little louder in there It's making me nervous how loud #VoteLeave are 😊 #BBCDebate
0	https://twitter.com/davenelli/status/745858986684092416	#EU is Thatcherism on a continental scale #Leave + build a new internationalism based on socialist cooperation #TUSC http://bit.ly/28WNF9V

9. Issue [Column O]

What issue/ theme does the tweet contain, if any?

- Economy and labour
- Healthcare
- Immigration
- Security
- Sovereignty
- Other
- None

Guidance Notes

- *Economy and labour*: references to markets, trade, business, money, wealth, cost, pay, trade unions, TTIP, GDP, recession, treasury, value of currency; anything fiscal, financial, or work-related
- *Healthcare*: anything NHS-related; references to nurses, doctors, £350 million per week
- *Immigration*: references to migrants from the EU and elsewhere, free movement, racism, xenophobia
- *Security*: anything related to defence, terrorism, policing, war, safety, peace; references to the military or armed forces
- *Sovereignty*: the UK's ability to make its own laws and govern itself; references to democracy, independence, laws, justice, influence, EU bureaucracy and dictatorship
- *Other*: avoid choosing this category unless you are convinced that the tweet contains an issue not captured by the categories above. If you do choose it, please indicate what you think the issue/theme is in the Notes column.
 - This category should be chosen when two or more themes are *equally* dominant in the tweet, or when many themes are present but described at a very vague level.
- *None*: do not be alarmed if many of the tweets contain no issues/themes. Do not try to force the tweet to fit into an issue category—remember to take it at face value.
 - This category should be chosen whenever the tweet is more concerned with people (e.g., lying politicians) than with substantive issues.
- These issues are adapted (with modifications) from the [BBC EU Referendum issues guide](#) and analysis on Brexit tweets from the [GATE team at the University of Sheffield](#).
- If the tweet appears to contain more than one issue, choose the one that is more salient.
- Do not be alarmed if some of the categories are far less prominent than others.

EXAMPLES		
Issue	Tweet URL	Tweet Text
Economy and labour	https://www.twitter.com/JLNASH2006/status/730099259794726913	RT @RedHotSquirrel: Even if the EU were to stay intact after Brexit its GDP would be smaller than India's some time mid-century.
Economy and labour	https://twitter.com/davenelli/status/745858986684092416	#EU is Thatcherism on a continental scale #Leave + build a new internationalism based on socialist cooperation #TUSC http://bit.ly/28WNF9V
Healthcare	https://twitter.com/SJ_Powell/status/742829639089823744	Don't risk the NHS ! #VoteLeave to avoid #TTIP privatisation

Healthcare	https://twitter.com/DAILYSQLIB/status/742084507436748801	Healthcare Workers Say We Should Vote Leave to Save the #NHS http://bit.ly/1RQMJ8G #brexit #voteleave #euref #remain #go #no2eu
Immigration	https://www.twitter.com/Tiamat2020/status/739801114346946560	RT @2tweetaboutit: 'Cover-up' over migrants entering UK as data on illegals is kept hidden https://t.co/cAu3J7rcol Vote #LeaveEU the only w...
Immigration	https://twitter.com/M_Pope1/status/741699561312419840	The wife and I are both voting to leave. Yes we are racist and yes we are xenophobic of course. #VoteLeave
Security	https://twitter.com/TerryGDavis7/status/739379310851723264	When the EU army replaces National Armies which EU state will they send them into in the greater interest of the EU if they vote to #leave?
Security	https://twitter.com/jk_rowling/status/742718884696993793	I don't want to leave a union that has brought more peace and stability to the region than it has ever known #Remain
Sovereignty	https://twitter.com/BrexitFTW/status/746029557208260609	#iVoted for the UK to leave EU, Britain never stands alone, Britain stands proud . #VoteLeave #Brexit #UK
Sovereignty	https://twitter.com/Hunter7Taylor/status/744459194053627904	BetterOffOut: Follow your heart, vote for democracy and hope. Vote Leave: http://www.telegraph.co.uk/news/2016/06/18/michael-gove-interview-vote-for-democracy-and-hope-vote-leave/ ... by Michael Gove via Telegraph #VoteLeave...
Other	https://www.twitter.com/3fish/status/738460592307109889	RT @racheljoyce: Me too. Two kids. One abt start uni. I want them to be able to have a future here. #VoteLeave https://t.co/P5qcQV61RY
Other	https://twitter.com/CraigWhittington/status/735058819982708736	Anyone wearing a poppy but voting Remain needs to take a long look in the mirror. #Leave #EUReferendum
None	https://twitter.com/karinsmyth/status/745266493282213888	Excellent session #EUref campaigning in Bedminster with Ollie (our leader), Richard & Elaine #StrongerIn
None	https://twitter.com/wdjstraw/status/746024134602293248	Lost your polling card? No problem, use this tool to find your polling station: http://www.wheredoivote.co.uk #StrongerIn

Appendix D: Interview Questions Template

I. Twitter Motives, Practises, and Significance [Individuals]

What motivated you to create a Twitter account in [insert month and year here]?

What motivated you to start tweeting (as opposed to just reading tweets)?

[If the account is verified] I noticed that your account is verified. When and how did this happen?

What motivated you to tweet about the EU referendum?

How often did you tweet during the EU referendum period (15 April-23 June 2016), compared to how often you normally tweet?

Did you follow or engage with Leave campaigners? Who were they?

Did you follow specific hashtags? What were they?

I. Twitter Motives, Practises, and Significance [Organisations]

What motivated [organisation name] to establish a Twitter account?

What did [organisation name] hope to accomplish with this account during the referendum period?

Did you tweet from the Twitter account?

What did your role as a co-manager of the account entail?

Did you follow or engage with Remain/Leave campaigners? Who were they?

Did you follow specific hashtags? What were they?

Did you create any hashtags?

How much time did you spend managing the [organisation name] Twitter account during the referendum period?

How many other people co-managed the [organisation name] Twitter account during the referendum period?

How much time did they spend managing the account?

Were you assigned to the management of the Twitter account, or did you volunteer to do it?

Did managing the Twitter account provide you with any personal or professional benefits?

Could you tell me about these benefits?

[If respondent also has a personal account] How did your Twitter strategy for your personal account differ from your strategy for the [organisation name] account, if at all?

II. Tweet Strategies and Success

Now that we've talked about your tweeting strategies, let's discuss what you hope to achieve through these strategies. What does success on Twitter mean to you?

[If the respondent indicates that retweets and likes are measures of success] Is getting just one retweet or one like success? My analysis reveals that over half of tweets related to the EU referendum are not retweeted or liked at all.

Here are the tweets that you sent that appeared in my dataset during the referendum period. Can you point to any examples of successful tweets? [The average tweet from your account in my dataset had X retweets and X likes. This one had many more.]

For each successful tweet:

- *Could you talk me through the process of how you decided to publish this tweet? [Open the tweet in Twitter and watch any videos that are embedded]*
- *How much time did you take to think about how to phrase this tweet and what to include in it?*
- *Which factors did you consider when you composed this tweet that contributed to its success?*

Are there any more recent tweets you have sent that were surprisingly successful?

If you are comfortable doing so, can you point to any examples of less successful tweets that you sent? For each unsuccessful tweet:

- *What made this tweet less successful?*
- *How much time did you take to think about how to phrase this tweet and what to include in it?*

[If not already addressed] Which strategy do you think is more effective on Twitter: positive or negative campaigning?

[If not already addressed] X% of your tweets in my dataset do not have an image, video, or GIF. When do you think it is a good idea to include them in a tweet? Do you think media items have an impact on tweet success?

[If not already addressed] X% of your tweets have URLs. What types of links do you share? Do you think this had an impact on their success?

[If not already addressed] X% of your tweets mention at least one user. How do you select which users to mention? Would you say this was a successful strategy?

[If not already addressed] X% of your tweets include more than one hashtag. How do you select which hashtags to include? Do you think this had an impact on their success?

Let's now talk about temporal factors. Here is a bar chart of your tweet frequency per hour of day. Most of your tweets were published between [insert time intervals here]. Was this deliberate (did you consider time of day when tweeting)? Do you think this had an impact on their success?

Here is a bar chart of your tweet frequency per day of week. Most of your tweets were published on [insert day of week here]. Was this deliberate (did you consider day of the week when tweeting)? Do you think this had an impact on their success?

Are there any other strategies that you use to boost the success of your tweets that we have not yet discussed?

Here are some examples of tweets that you retweeted during the referendum period. What about these tweets motivated you to share them?

Here are the tweets that you have recently liked. What about these tweets motivated you to like them?

How often do you retweet content you disagree with?

How often do you like content you disagree with?

Can you explain the difference between a retweet and a like?

III. Social Media and Political Campaigns [Organisations]

What were the other social media platforms that [organisation name] used during the referendum period?

How important was Twitter compared with these other platforms?

How much money was spent on Twitter vs. the other platforms?

Can you rank these platforms in order of importance?

Apart from social media, what were the communication channels that [organisation name] used during the referendum period?

How important was social media for [organisation name] compared with these other channels?

What was the most important channel?

Do you think [organisation name]'s tweets had any influence on the referendum outcome? Please provide specific examples to support this conclusion.

IV. Conclusion

Is there anything else you would like to add? Is there anything important that I have omitted?

Who else do you think I should talk to?

Could you provide me with their contact information?

Would it be okay for me to use your name as a reference when contacting them?

Appendix E: Participant Information Sheet and Informed Consent Form



Participant Information Sheet

#VOTELEAVE OR #STRONGERIN: RESONANCE AND RHETORIC IN THE EU REFERENDUM

You are an important participant in my doctoral research because you actively tweeted about the UK EU membership referendum between 15 April and 23 June 2016, in your capacity as the [position] of [organisation name]. I am investigating rhetorical factors that influence the resonance of political messages on social media: *how can messages be structured and delivered most effectively?* To better understand the 26 million referendum-related tweets that I have collected and the rhetorical strategies of their authors, I am conducting a series of interviews.

To take part, you will have to answer questions about your tweeting motives and practises in the context of politics, and how you define success on Twitter. I will show you a selection of tweets that you published during the campaigning period, along with statistics on their retweet and like counts, and we will discuss why some were more resonant than others. You will emerge from the interview with a better understanding of how to publish engaging content on social media.

This project has been approved by the University of Oxford Central University Research Ethics Committee. Participation is voluntary, and you may withdraw from the interview without penalty at any time. If you are happy to take part, you will be asked to sign a consent form. The interview should take no more than an hour to complete, and can be conducted in person or via phone, Skype, or email depending on your preferences. The interview will be recorded, so that it can be transcribed and analysed. I am the only person with access to the recording, and it will be stored using the highest security standards (on a passcode-protected computer).

The outcomes of the research will be written up as a doctoral thesis, and may include verbatim quotes based on the interviews conducted. I will confirm any direct quotes from our interview that I would like to use, and you will be quoted as the [position] of [organisation name]. Upon successful submission, results may be published in academic journals or included in presentations, and the thesis will be deposited both in print and online in the University of Oxford archives to facilitate its use in future research. At this point, all interview recordings will be deleted.

If you have a question or concern about any aspect of this project, please contact me at yin.lu@oii.ox.ac.uk or one of my supervisors: Grant Blank at grant.blank@oii.ox.ac.uk or Jonathan Bright at jonathan.bright@oii.ox.ac.uk. Your concern should be acknowledged within ten working days, and an indication will be provided of how it will be dealt with. If you remain unhappy or wish to make a formal complaint, please contact the chair of the Social Sciences & Humanities Inter-Divisional Research Ethics Committee at the University of Oxford, who will seek to resolve the matter in a reasonably expeditious manner (ethics@socsci.ox.ac.uk; Research Services, University of Oxford, Wellington Square, Oxford OX1 2JD).

Thank you; I look forward to speaking with you if you are happy to be interviewed.

Best wishes,

Yin Yin Lu

oxford internet institute
university of oxford
one st giles oxford ox1 3js
united kingdom
tel +44(0)1865 287210 fax +44(0)1865 287211
enquiries@oii.ox.ac.uk
www.oii.ox.ac.uk





Informed Consent Form

#VOTELAVE OR #STRONGERIN: RESONANCE AND RHETORIC IN THE EU REFERENDUM

I consent to my participation in an interview about EU referendum tweets as part of Yin Yin Lu's doctoral thesis research project. I hereby declare that:

- ☐ I have read the participant information sheet;
- ☐ I have had the opportunity to ask questions about the study and have received satisfactory answers to questions, and any additional details requested;
- ☐ I understand that this project has been reviewed by, and received ethics clearance through, the University of Oxford Central University Research Ethics Committee;
- ☐ I understand that I may withdraw from the study at any time without penalty by advising the researcher of this decision, and that I do not have to provide a reason;
- ☐ I understand who will have access to personal data provided;
- ☐ I understand how the data will be stored, and what will happen to the data at the end of the project;
- ☐ I know that the research will be written up as a doctoral thesis, and understand how personal data included in that thesis will be published;
- ☐ I consent to being audio recorded;
- ☐ I understand how to raise a concern or make a complaint.

Name of Participant: _____

Signature: _____ Date: _____

Name of Researcher: _____


Signature: _____ Date: _____


oxford internet institute
university of oxford
one st giles oxford ox1 3js
united kingdom
tel +44(0)1865 287210 fax +44(0)1865 287211
enquiries@oii.ox.ac.uk
www.oii.ox.ac.uk














Appendix F: Bivariate Crosstabulations and Stacked Bar Charts









I. Sample Level















II. Population Level








Appendix G: Regression Diagnostics (Sample Level)



Pearson's Correlation Coefficients for Continuous Control Variables

	Followers count (log ₁₀)	Friends count (log ₁₀)	Statuses count (log ₁₀)	Account age	Days before referendum
Followers count (log ₁₀)	1.000***	0.804***	0.668***	0.365***	0.057**
Friends count (log ₁₀)	0.804***	1.000***	0.511***	0.376***	0.072***
Statuses count (log ₁₀)	0.668***	0.511***	1.000***	0.398***	0.025
Account age	0.365***	0.376***	0.398***	1.000***	-0.049**
Days before referendum	0.057**	0.072***	0.025	-0.049**	1.000***

* p < .05, ** p < .01, *** p < .001



Odds Ratios from Logistic Regression Models with Discrete Emotions

	WAS RETWEETED	WAS LIKED
CONTROL	Leave	0.743 *
	Remain	0.847
	Followers count ($\log_{10}$)	1.641 ***
	Account age	0.910 ***
	Small hours	0.541 *
	Morning	0.754 **
	Afternoon	0.946
	Monday	0.320 ***
	Tuesday	0.275 ***
	Wednesday	0.301 ***
	Thursday	0.292 ***
	Friday	0.368 ***
	Sunday	0.305 ***
	Days before referendum	1.011 ***
	Hashtag count	0.911 ***
DISCOVERABILITY	Mention count	1.050
	URL count (binary)	0.633 ***
	Media type count (binary)	2.182 ***
	Organisation	7.555 ***
	Enthusiasm	1.071
CONTENT	Anxiety	1.134
	Aversion	1.111
	Economy and labour	1.044
	Healthcare	0.639
	Immigration	1.005
	Other issue	1.028
	Security	1.121
	Sovereignty	0.770
	Resolution	0.852
	Message format	0.898
STYLE	Question	0.919
	Playfulness	0.838
Number of Observations		2,976
McFadden's R^2		0.12
AIC		3,592.41
		3,569.96

* $p < .05$, ** $p < .01$, *** $p < .001$

The omitted reference categories are ambiguous (position), night (hour of day), Saturday (day of week), and none (issue).

Odds Ratios from Logistic Regression Models With and Without High Leverage

	RT MODEL	RT MODEL (NO HIGH LEVERAGE)	LIKE MODEL	LIKE MODEL (NO HIGH LEVERAGE)
CONTROL	Leave	0.637 **	0.639 **	0.771
	Remain	0.725	0.730	1.057
	Followers count ($\log_{10}$)	1.647 ***	1.643 ***	1.579 ***
	Account age	0.911 ***	0.912 ***	0.976
	Small hours	0.547 *	0.572 *	0.590 *
	Morning	0.757 **	0.760 **	0.787 *
	Afternoon	0.949	0.951	0.711 ***
	Monday	0.323 ***	0.319 ***	0.299 ***
	Tuesday	0.278 ***	0.277 ***	0.338 ***
	Wednesday	0.306 ***	0.304 ***	0.353 ***
	Thursday	0.296 ***	0.295 ***	0.383 ***
	Friday	0.372 ***	0.370 ***	0.328 ***
	Sunday	0.308 ***	0.307 ***	0.321 ***
	Days before referendum	1.011 ***	1.011 ***	0.996
	Hashtag count	0.912 ***	0.913 ***	0.905 ***
	Mention count	1.043	1.019	0.888
DISCOVERABILITY	URL count (binary)	0.634 ***	0.635 ***	0.468 ***
	Media type count (binary)	2.166 ***	2.159 ***	1.797 ***
	Organisation	7.414 ***	7.354 ***	5.674 ***
	Identification	1.290	1.306	1.051
	Division	1.247	1.255	1.041
CONTENT	Emotion	1.054	1.056	1.223
	Economy and labour	1.020	1.018	0.901
	Healthcare	0.629	0.656	0.514 *
	Immigration	0.999	0.999	0.892
	Other issue	0.995	0.996	0.900
	Security	1.097	1.091	0.952
	Sovereignty	0.746	0.743	0.649 *
	Resolution	0.842	0.829	0.923
	Message format	0.885	0.883	1.015
STYLE	Question	0.924	0.922	0.743 *
	Playfulness	0.844	0.844	1.024
Number of Observations		2,976	2,969	2,427
McFadden's R^2		0.12	0.11	0.11
AIC		3,590.02	3,585.09	3,572.39

* $p < .05$, ** $p < .01$, *** $p < .001$

The omitted reference categories are ambiguous (position), night (hour of day), other (message framing), Saturday (day of week), and none (issue).

Appendix H: Linear Regression (Population Level)

OLS Regression for Retweet Count ($\log_{10}$)

	COEFFICIENT	STANDARD ERROR
CONTROL	Leave	0.0061 **
	Remain	-0.0142 ***
	Followers count ($\log_{10}$)	0.1509 ***
	Account age	-0.0223 ***
	Small hours	-0.0562 ***
	Morning	-0.0103 ***
	Afternoon	-0.0229 ***
	Monday	-0.0677 ***
	Tuesday	-0.0922 ***
	Wednesday	-0.0895 ***
	Thursday	-0.0791 ***
	Friday	-0.0612 ***
	Sunday	-0.0457 ***
	Days before referendum	0.0011 ***
	Hashtag count	-0.0129 ***
DISCOVERABILITY	Mention count	-0.0042 ***
	URL count (binary)	-0.0554 ***
	Photo	0.1537 ***
	Video	0.3430 ***
	GIF	0.0918 ***
	Organisation	0.6218 ***
	Division	-0.0105 ***
CONTENT & STYLE	Emotion	0.0152 ***
	Issue	0.0145 ***
	Resolution	-0.0441 ***
	Question	-0.0088 ***
Number of Observations		741,411
R ²		0.460
Adjusted R ²		0.460

* $p < .05$, ** $p < .01$, *** $p < .001$

The omitted reference categories are ambiguous (position), night (hour of day), other (message framing), and Saturday (day of week).

9. Bibliography

9.1 Books, Book Chapters, and Journal Articles

- Aaker, J. (1997). Dimensions of brand personality. *Journal of Marketing Research*, 34(3), 347-456.
- Ahluwalia, R., & Burnkrant, R. E. (2004). Answering questions about questions: A persuasion knowledge perspective for understanding the effects of rhetorical questions. *Journal of Consumer Research*, 31(1), 26-42.
- Alter, A. (2017). *Irresistible: Why we can't stop checking, scrolling, clicking, and watching*. London, UK: The Bodley Head.
- Ansolabehere, S., & Iyengar, S. (1995). *Going negative: How attack ads shrink and polarise the electorate*. London, UK: The Free Press.
- Anstead, N., & O'Loughlin, B. (2010). The emerging viewertariat: explaining Twitter responses to Nick Griffin's appearance on BBC Question Time. *University of East Anglia School of Political, Social and International Studies Working Paper Series 1*.
- Aristotle. (2007). *On rhetoric: A theory of civic discourse* (2nd ed.; G. A. Kennedy, Trans.). Oxford, UK: Oxford University Press.
- Atkinson, C. (2010). *The backchannel: How audiences are using Twitter and social media and changing presentations forever*. Berkeley, CA: New Riders.
- Atkinson, M. (1984). *Our masters' voices: The language and body language of politics*. London, UK: Methuen.
- Bamman, D., Eisenstein, J., & Schnoebelen, T. (2014). Gender identity and lexical variation in social media. *Journal of Sociolinguistics*, 18(2), 135-160.
- Bansal, P., Bansal, R., & Varma, V. (2015). Towards deep semantic analysis of hashtags. *arXiv:1501.03210 [cs.IR]*.
- Barash, V., & Kelly, J. (2012). Salience vs. commitment: Dynamics of political hashtags in Russian Twitter. *Berkman Center Research Publication* 2012.9.
- Baron, N. (2008). *Always on*. Oxford, UK: Oxford University Press.
- Barthes, R. (1990). *The pleasure of the text*. Oxford, UK: Basil Blackwell.
- Barthes, R. (2014). The old rhetoric: An aide-mémoire. In M. Bauer, A. Süerdem, & A. Biquelet (Eds.), *Textual analysis: Vol. 2* (pp. 197-256). London, UK: SAGE Publications.
- Bartlett, J., Miller, C., Reffin, J., Weir, D., & Wibberley, S. (2014). *Vox digitas*. London, UK: Demos Centre for the Analysis of Social Media. Retrieved from https://www.demos.co.uk/files/Vox_Digitas_-_web.pdf?1408832211

- Barton, D., & Lee, C. (2013). *Language online: Investigating digital texts and practices*. London, UK: Routledge.
- Bauer, M. (2000). Classical content analysis: A review. In M. Bauer & G. Gaskell (Eds.), *Qualitative researching with text, image and sound* (pp. 132-152). London, UK: SAGE Publications.
- Baym, N. K. (1995). The performance of humor in computer-mediated communication. *Journal of Computer-Mediated Communication*, 1(2). doi:10.1111/j.1083-6101.1995.tb00327.x
- Becker, H. S. (1958). Problems of inference and proof in participant observation. *American Sociological Review*, 23(6), 652-660.
- Bennett, W. L., & Iyengar, S. (2008). A new era of minimal effects? The changing foundations of political communication. *Journal of Communication*, 58(4), 707-731.
- Berelson, B. (1952). *Content analysis in communication research*. New York, NY: Free Press.
- Berezin, M. (1994). Cultural form and political meaning: State-subsidized theater, ideology, and the language of style in fascist Italy. *American Journal of Sociology*, 99(5), 1237-1286.
- Berezin, M. (2002). Secure states: towards a political sociology of emotion. *The Sociological Review*, 50(S2), 33-52.
- Berger, J. (2013). *Contagious: Why things catch on*. London, UK: Simon & Schuster.
- Berger, J., & Milkman, K. (2012). What Makes Online Content Viral? *Journal of Marketing Research*, (49)2, 192-205.
- Berger, J. (2011). Arousal increases social transmission of information. *Psychological Science*, 22(7), 891-893.
- Berlyne, D. E. (1971). *Aesthetics and psychobiology*. New York, NY: Appleton.
- Berners-Lee, T. (1999). *Weaving the Web*. London, UK: Orion Business Books.
- Bimber, B., Flanagin, A. J., & Stohl, C. (2005). Reconceptualising collective action in the contemporary media environment. *Communication Theory*, 15(4), 365-388.
- Bitzer, L. (2014). The rhetorical situation. In M. Bauer, A. Süerdem, & A. Biquelet (Eds.), *Textual analysis: Vol. 2* (pp. 257-267). London, UK: SAGE Publications.
- Black, E. (1992). *Rhetorical questions: Studies of public discourse*. London, UK: The University of Chicago Press.
- Blank, G. (2017). The digital divide among Twitter users and its implications for social research. *Social Science Computer Review*, 35(6), 679-697.
- Blank, G., & Reisdorf, B. (2012). The participatory Web. *Information, Communication & Society*, 15(4), 537-554.

- Blaszka, M., Burch, L. M., Frederick, E. L., Clavio, G., & Walsh, P. (2012). #WorldSeries: An empirical examination of a Twitter hashtag during a major sporting event. *International Journal of Sport Communication*, 5(4), 435-453.
- Bode, L., Hanna, A., Yang, J., & Shah, D. (2014). Candidate networks, citizen clusters, and political expression: Strategic hashtag use in the 2010 midterms. *The ANNALS of the American Academy of Political and Social Science*, 659(1), 149-165.
- Borchers, T. (2013). *Persuasion in the media age* (3rd ed.). Long Grove, IL: Waveland Press.
- Bowen, G. A. (2008). Naturalistic inquiry and the saturation concept: A research note. *Qualitative Research*, 8(1), 137-152.
- boyd, d., & Crawford, K. (2012). Critical questions for big data: Provocations for a cultural, technological, and scholarly phenomenon. *Information, Communication & Society*, 15(5), 662-679.
- boyd, d., Golder, S., & Lotan, G. (2010). Tweet, tweet, retweet: Conversational aspects of retweeting on Twitter. In *Proceedings of the Forty-Third Hawai'i International Conference on System Sciences*. Washington, DC: IEEE Computer Society.
- Brader, T. (2005). Striking a responsive chord: How political ads motivate and persuade voters by appealing to emotions. *American Journal of Political Science*, 49(2), 388-405.
- Brader, T. (2006). *Campaigning for hearts and minds: How emotional appeals in political ads work*. London, UK: University of Chicago Press.
- Bright, J., Hale, S., Ganesh, B., Bulovsky, A., Margetts, H., & Howard, P. (2017). Does campaigning on social media make a difference? Evidence from candidate use of Twitter during the 2015 and 2017 UK elections. *arXiv:1710.07087 [cs.SI]*.
- Bringer, J., Johnston, L., & Brackenridge, C. (2006). Using computer-assisted qualitative data analysis software to develop a grounded theory project. *Field Methods*, 18(3), 245-266.
- Brinkmann, S. (2013). *Qualitative interviewing*. Oxford, UK: Oxford University Press.
- British Sociological Association. (2017a). *BSA statement of ethical practice*. Durham, UK: BSA Publications. Retrieved from https://www.britsoc.co.uk/media/24310/bsa_statement_of_ethical_practice.pdf
- British Sociological Association. (2017b). *Ethics guidelines and collated resources for digital research: Statement of ethical practice annexe*. Durham, UK: BSA Publications. Retrieved from https://www.britsoc.co.uk/media/24309/bsa_statement_of_ethical_practice_annexe.pdf
- Brooks, D. J., & Geer, J. G. (2007). Beyond negativity: The effects of incivility on the electorate. *American Journal of Political Science*, 51(1), 1-16.
- Bruns, A. (2012). How long is a tweet? Mapping dynamic conversation networks on Twitter using Gawk and Gephi. *Information, Communication, & Society*, 15(9), 1323-1351.

- Bruns, A., & Burgess, J. E. (2011). The use of Twitter hashtags in the formation of ad hoc publics. Presented at the *Sixth European Consortium for Political Research General Conference*. Reykjavik: University of Iceland.
- Bruns, A., Enli, G., Skogerbo, E., Larsson, A., & Christensen, C. (Eds.). (2017). *The Routledge companion to social media and politics*. Oxford, UK: Oxford University Press.
- Bryman, A. (2012). *Social research methods* (4th ed.). Oxford, UK: Oxford University Press.
- Burke, K. (1969). *A Rhetoric of Motives*. Berkeley, CA: University of California Press.
- Burke, K. (1984). The rhetoric of Hitler's 'battle'. In M. J. Shapiro (Ed.), *Language and politics* (pp. 61-80). Oxford, UK: Basil Blackwell.
- Burscher, B., Odijk, D., Vliegthart, R., de Rijke, M., & de Vreese, C. H. (2014). Teaching the computer to code frames in news: Comparing two supervised machine learning approaches to frame analysis. *Communication Methods and Measures*, 8(3), 190-206.
- Carraro, L., Gawronski, B., & Castelli, L. (2010). Losing on all fronts: The effects of negative versus positive person-based campaigns on implicit and explicit evaluations of political candidates. *British Journal of Social Psychology*, 49(3), 453-470.
- Cassidy, W. P. (2006). Gatekeeping similar for online, print journalists. *Newspaper Research Journal*, 27(2), 6-23.
- Castells, M. (2009). *Communication power*. Oxford, UK: Oxford University Press.
- Ceron, A., & d'Adda, G. (2016). E-campaigning on Twitter: The effectiveness of distributive promises and negative campaign in the 2013 Italian election. *New Media & Society*, 18(9), 1935-1955.
- Chadwick, A. (2013). *The hybrid media system: Politics and power*. Oxford, UK: Oxford University Press.
- Chaffee, S., & Metzger, M. (2001). The end of mass communication. *Mass Communication and Society*, 4(4), 365-379.
- Chang, H. (2010). A new perspective on Twitter hashtag use: Diffusion of innovation theory. *Proceedings of the American Society for Information Science and Technology*, 47(1), 1-4.
- Charteris-Black, J. (2005). *Politicians and rhetoric: The persuasive power of metaphor*. Basingstoke, UK: Palgrave Macmillan.
- Chilton, P. (2004). *Analysing political discourse: Theory and practice*. London, UK: Routledge.
- Cho, J., & Trent, A. (2006). Validity in qualitative research revisited. *Qualitative Research*, 6(3), 319-340.
- Cho, Y. (2008). Intercoder reliability. In P. J. Lavrakas (Ed.), *Encyclopedia of survey research methods* (pp. 345-346). Thousand Oaks, CA: SAGE Publications.

- Cialdini, R. B. (2007). *Influence: The psychology of persuasion* (revised ed.). New York, NY: Collins.
- Clark-Carter, D. (1997). *Doing quantitative psychological research: From design to report*. Hove, UK: Psychology Press.
- Coleman, J. S. (1990). *Foundations of social theory*. London, UK: Belknap Press.
- Corbett, E. P. J. (1999). *Classical rhetoric for the modern student* (4th ed). Oxford, UK: Oxford University Press.
- Crawford, K. (2009). Following you: Disciplines of listening in social media. *Continuum: Journal of Media & Cultural Studies*, 23(4), 525-535.
- Creswell, J. W., & Plano Clark, V. L. (2011). *Designing and conducting mixed methods research* (2nd ed.). London, UK: SAGE Publications.
- Curtice, J. (2016). *The two poles of the referendum debate: Immigration and the economy*. London, UK: NatCen Social Research. Retrieved from <http://whatukthinks.org/eu/wp-content/uploads/2016/01/Analysis-paper-4-The-two-poles-of-the-referendum-debate.pdf>
- Curtice, J., & Evans, G. (2015). Britain and Europe: Are we all Eurosceptics now? In J. Curtice & R. Ormston (Eds.), *British Social Attitudes: The 32nd Report*. London, UK: NatCen Social Research. Retrieved from <http://www.bsa.natcen.ac.uk/latest-report/british-social-attitudes-32/europe.aspx>
- Crystal, D. (2006). *Language and the internet*. Cambridge, UK: Cambridge University Press.
- Crystal, D. (2009). *The future of language*. Abingdon, UK: Routledge.
- Dahl, D., Frankenberger, K., & Manchanda, R. (2003). Does it pay to shock? Reactions to shocking and nonshocking advertising content among university students. *Journal of Advertising Research*, 43(3), 268-280.
- Dainton, M., & Zelle, E. D. (2004). *Applying communication theory for professional life: A practical introduction*. London, UK: SAGE Publications.
- Danet, B. (2001). *Cyberpl@y: Communicating online*. Oxford, UK: Berg.
- DeCuir-Gunby, J., Marshall, P., & McCulloch, A. (2011). Developing and using a codebook for the analysis of interview data: An example from a professional development research project. *Field Methods*, 23(2), 136-155.
- Demos. (2016). *The use of misogynistic terms on Twitter*. London, UK: Demos Centre for the Analysis of Social Media. Retrieved from <https://www.demos.co.uk/wp-content/uploads/2016/05/Misogyny-online.pdf>
- Demšar, J., & Zupan, B. (2013). Orange: Data mining fruitful and fun—a historical perspective. *Informatica: An International Journal of Computing and Informatics*, 37(1), 55-60.

- Demšar, J., Curk, T., Erjavec, A., Gorup, Č., Hočvar, T., Milutinovič, M., ... Zupan, B. (2013). Orange: Data mining toolbox in Python. *Journal of Machine Learning Research*, 14(8), 2349-2353.
- Dillard, J. P., & Shen, L. (Eds.). (2013). *The SAGE handbook of persuasion: Developments in theory and practice* (2nd ed.). London, UK: SAGE Publications.
- Driscoll, K., & Walker, S. (2014). Working within a black box: Transparency in the collection and production of big Twitter data. *International Journal of Communication*, 8, 1745-1764.
- Dubois, E., & Ford, H. (2015). Trace interviews: An actor-centered approach. *International Journal of Communication*, 9, 2067-2091.
- Efron, M. (2010). Hashtag retrieval in a microblogging environment. In *Proceedings of the 33rd International ACM SIGIR Conference on Research and Development in Information Retrieval*. New York, NY: ACM.
- Eisenhardt, K. M. (1989). Building theories from case study research. *Academy of Management Review*, 14(4), 532-550.
- Entman, R. (2014). Framing: Toward clarification of a fractured paradigm. In M. Bauer, A. Süerdem, & A. Bicquelet (Eds.), *Textual analysis: Vol. 2* (pp. 269-276). London, UK: SAGE Publications.
- Ferragina, P., Piccinno, F., & Santoro, R. (2015). On analysing hashtags in Twitter. In *Proceedings of the Ninth International AAAI Conference on Web and Social Media*. Palo Alto, CA: AAAI Press.
- Ferrara, E., Varol, O., Davis, C., Menczer, F., & Flammini, A. (2016). The rise of social bots. *Communications of the ACM*, 59(7), 96-104.
- Finkel, S. E., & Geer, J. G. (1998). A spot check: Casting doubt on the demobilizing effect of attack advertising. *American Journal of Political Science*, 42(2), 573-595.
- Flanagin, A. J., Stohl, C., & Bimber, B. (2006). Modelling the structure of collective action. *Communication Monographs*, 73(1), 29-54.
- Frank, J. (1990). You call that a rhetorical question? Forms and functions of rhetorical questions in conversation. *Journal of Pragmatics*, 14(5), 723-738.
- Fridkin, K. L., & Kenney, P. J. (2011). Variability in citizens' reactions to different types of negative campaigns. *American Journal of Political Science*, 55(2), 307-325.
- Fu, P., Lin, Z., Yuan, F., Wang, W., & Meng, D. (2018). Learning sentiment-specific word embedding via global sentiment representation. In *Proceedings of the 32nd AAAI Conference on Artificial Intelligence*. Palo Alto, CA: AAAI.
- Gainous, J., & Wagner, K. M. (2014). *Tweeting to power: The social media revolution in American politics*. New York, NY: Oxford University Press.
- Gans, H. J. (1979). *Deciding what's news*. New York, NY: Pantheon.

- Gayo-Avello, D. (2012). A balanced survey on election prediction using Twitter data. *arXiv:1204.6441v1 [cs.CY]*.
- Geertz, C. (1973). *The interpretation of cultures: Selected essays*. New York, NY: Basic Books.
- Gerbaudo, P. (2012). *Tweets and the streets: Social media and contemporary activism*. London, UK: Pluto Press.
- Gibson, J. J. (1979). *The ecological approach to visual perception*. New York, NY: Houghton Mifflin.
- Glaser, B., & Strauss, A. L. (1967). *The discovery of grounded theory: Strategies for qualitative research*. London, UK: Wiedenfeld and Nicholson.
- Goel, S., Anderson, A., Hofman, J., & Watts, D. (2016). The structural virality of online diffusion. *Management Science*, 62(1), 180-196.
- Gold, R. (1997). The ethnographic method in sociology. *Qualitative Inquiry*, 3(4), 388-402.
- Goodwin, J., Jasper, J. M., & Polletta, F. (Eds.). (2001). *Passionate politics: Emotion and social movements*. Chicago, IL: The University of Chicago Press.
- Gorodnichenko, Y., Pham, T., & Talavera, O. (2018). Social media, sentiment, and public opinions: Evidence from #Brexit and #USElection. *NBER Working Paper Number 24631*.
- Grady, K., & Potter, J. (1985). Speaking and clapping: a comparison of Foot and Thatcher's oratory. *Language & Communication*, 5(3), 173-183.
- Gray, D. E. (2014). *Doing research in the real world* (3rd ed.). Los Angeles, CA: SAGE Publications.
- Grice, H. P. (1991). *Studies in the way of words*. Cambridge, MA: Harvard University Press.
- Grimmer, J., & Stewart, B. M. (2013). Text as data: The promise and pitfalls of automatic content analysis methods for political texts. *Political Analysis*, 21(3), 267-297.
- Griswold, W. (1986). *Renaissance revivals: City comedy and revenge tragedy in the London theatre, 1576-1980*. London, UK: The University of Chicago Press.
- Griswold, W. (1987a). The fabrication of meaning: Literary interpretation in the United States, Great Britain, and the West Indies. *American Journal of Sociology*, 92(5), 1077-1117.
- Griswold, W. (1987b). A methodological framework for the sociology of culture. *Sociological Methodology*, 17, 1-35.
- Griswold, W. (2013). *Cultures and societies in a changing world* (4th ed.). London, UK: SAGE Publications.
- Gurak, L. (1997). *Persuasion and privacy in cyberspace: The online protests over Lotus MarketPlace and the Clipper Chip*. London, UK: Yale University Press.

- Hair, J. F., Black, W. C., Babin, B. J., & Anderson, R. E. (2010). *Multivariate data analysis* (7th ed.). London, UK: Pearson.
- Hall, Todd H. (2015). *Emotional diplomacy: Official emotion on the international stage*. Ithaca, NY: Cornell University Press.
- Hammersley, M. (2004). Case study. In M. S. Lewis-Beck, A. Bryman, & T. F. Liao (Eds.), *The SAGE encyclopedia of social science research methods* (pp. 93-95). Thousand Oaks, CA: SAGE Publications.
- Hannah, D., & Lautsch, B. (2011). Counting in qualitative research: Why to conduct it, when to avoid it, and when to closet it. *Journal of Management Inquiry*, 20(1), 14-22.
- Hart, R. P. (2000). *Campaign talk: Why elections are good for us*. Princeton, NJ: Princeton University Press.
- Hayes, A., & Krippendorff, K. Answering the call for a standard reliability measure for coding data. *Communication Methods and Measures* 1(1), 77-89.
- Hayes, R. A., Carr, C. T., & Wohn, D. Y. (2016). One click, many meanings: Interpreting paralinguistic digital affordances in social media. *Journal of Broadcasting & Electronic Media*, 60(1), 171-187.
- Heritage, J., & Greatbatch, D. (1986). Generating applause: a study of rhetoric and response at party political conferences. *American Journal of Sociology*, 92(1), 110-157.
- Hermida, A. (2010). Twittering the news: The emergence of ambient journalism. *Journalism Practice*, 4(3), 297-308.
- Herring, S. C. (1999). Interactional coherence in CMC. *Journal of Computer-Mediated Communication*, 4(4). doi:10.1111/j.1083-6101.1999.tb00106.x
- Herring, S. C. (2004). Computer-mediated discourse analysis: An approach to researching online behaviour. In S. A. Barab, R. Kling, & J. H. Gray (Eds.), *Designing for virtual communities in the service of learning* (pp. 338-376). New York, NY: Cambridge University Press.
- Herring, S. C. (2013). Discourse in Web 2.0: Familiar, reconfigured, and emergent. In D. Tannen & A. M. Trester (Eds.), *Discourse 2.0: Language and new media* (pp. 1-25). Washington, DC: Georgetown University Press.
- Heverin, T., & Zach, L. (2010). Microblogging for crisis communication: Examination of Twitter use in response to a 2009 violent crisis in the Seattle-Tacoma, Washington area. Presented at the *Seventh International ISCRAM Conference*. Seattle, WA: Crowne Plaza Hotel.
- Highfield, T., & Leaver, T. (2015). A methodology for mapping Instagram hashtags. *First Monday*, 20(1).
- Hodgson, Geoffrey M. (2006). What are institutions? *Journal of Economic Issues*, 40(1), 1-25.

- Hopkins, D. J., & King, G. (2010). A method of automated nonparametric content analysis for social science. *American Journal of Political Science*, 54(1), 229-247.
- Honeycutt, C., & Herring, S. (2009). Beyond microblogging: Conversation and collaboration via Twitter. In *Proceedings of the Forty-second Hawai'i International Conference on System Sciences*. Los Alamitos, CA: IEEE Press.
- Housley, W., Procter, R., Edwards, A., Burnap, P., Williams, M., Sloan, L., Rana, O., Morgan, J., Voss, A., & Greenhill, A. (2014). Big and broad social data and the sociological imagination: A collaborative response. *Big Data & Society*, 1(2), 1-15.
- Howard, P. N., & Kollanyi, B. (2016). Bots, #StrongerIn, and #Brexit: Computational Propaganda during the UK-EU referendum. Comprop Research Note 2016.1. Oxford, UK: Project on Computational Propaganda.
- Huang, J., Thornton, K. M., & Efthimiadis, E. N. (2010). Conversational tagging in Twitter. In *Proceedings of the Twenty-first ACM Conference on Hypertext and Hypermedia*. New York, NY: ACM.
- Huberman, B. A., Romero, D. M., & Wu, F. (2008). Social networks that matter: Twitter under the microscope. *First Monday*, 14(1).
- Iyengar, S., & Reeves, R. (1997). *Do the media govern? Politicians, voters, and reporters in America*. London, UK: SAGE Publications.
- Jain, N., Agarwal, P., & Pruthi, J. (2015). HashJacker—Detection and analysis of hashtag hijacking on Twitter. *International Journal of Computer Applications*, 114(19), 17-20.
- Jamieson, K. H. (1988). *Eloquence in an electronic age: The transformation of political speechmaking*. Oxford, UK: Oxford University Press.
- Jansen, B., Zhang, M., Sobel, K., & Chowdury, A. (2009). Twitter power: Tweets as electronic word of mouth. *Journal of the American Society for Information Science and Technology*, 60(11), 2169-2188.
- Java, A., Finin, T., Song, X., & Tseng, B. (2007). Why we Twitter: understanding microblogging usage and communities. In *Proceedings of the Joint Ninth WEBKDD and First SNA-KDD Workshop*. New York, NY: ACM. 56-65.
- Jawahar, G., Gupta, M., & Varma, V. (2017). Interpretation of semantic tweet representations. In *Proceedings of the 2017 IEEE/ACM International Conference on Advances in Social Networks Analysis and Mining*. New York, NY: ACM. 95-102.
- Jeffares, S. (2014). *Interpreting hashtag politics: Policy ideas in an era of social media*. London, UK: Palgrave Macmillan.
- Jungherr, A. (2016). Twitter use in election campaigns: A systematic literature review. *Journal of Information Technology & Politics*, 13(1), 72-91.
- Jungherr, A., Jürgens, P., & Schoen, H. (2012). Why the Pirate Party won the German election of 2009 or The trouble with predictions: A response to Tumasjan, A., Sprenger, T. O.,

- Sander, P. G., & Welp, I. M. 'Predicting elections with Twitter: What 140 characters reveal about political sentiment'. *Social Science Computer Review*, 30(2), 229-234.
- Kahn, K. F., & Kenney, P. J. (1999). Do negative campaigns mobilize or suppress turnout? Clarifying the relationship between negativity and participation. *The American Political Science Review*, 93(4), 877-889.
- Kanouse, D. E., & Hanson, L. R. (1971). Negativity in evaluation. In E. E. Jones, D. E. Kanouse, H. H. Kelley, R. E. Nisbett, S. Valins, & B. Weiner (Eds.), *Attribution: Perceiving the causes of behaviour* (pp. 47-62). Morristown, NJ: General Learning Press.
- Kastely, J. (2004) Pathos: Rhetoric and emotion. In W. Jost & W. Olmsted (Eds.), *A companion to rhetoric and rhetorical criticism* (pp. 221-237). Oxford, UK: Blackwell Publishing.
- Keesing, R. M. (1974). Theories of culture. *Annual Review of Anthropology*, 3, 73-97.
- Kehoe, A., & Gee, M. (2011). Social tagging: A new perspective on textual 'aboutness'. *Studies in Variation, Contacts and Change in English*, 6(5).
- Kernell, S. (1977). Presidential popularity and negative voting: An alternative explanation of the midterm congressional decline of the President's party. *The American Political Science Review*, 71(1), 44-66.
- Kim, Y., Kim, Y., Lee, J. S., Oh, J., & Lee, N. Y. (2015). Tweeting the public: Journalists' Twitter use, attitudes toward the public's tweets, and the relationship with the public. *Information, Communication, & Society*, 18(4), 443-458.
- Koshik, I. (2005). *Beyond rhetorical questions: Assertive questions in everyday interaction*. Amsterdam, NL: John Benjamins Publishing Company.
- Krippendorff, K. (2004). *Content analysis: An introduction to its methodology*. Thousand Oaks, CA: Sage.
- Krippendorff, K. (2011). Agreement and information in the reliability of coding. *Communication Methods and Measures*, 5(2), 93-112.
- Krishnamurthy, B., Gill, P., & Arlitt, M. (2008). A few chirps about Twitter. In *Proceedings of the First Workshop on Online Social Networks*. New York, NY: ACM.
- Kwak, H., Lee, C., Park, H., & Moon, S. (2010). What is Twitter, a social network or a news media? In *Proceedings of the 19th International Conference on World Wide Web*. New York, NY: ACM.
- Lau, R. R. (1982). Negativity in political perception. *Political Behaviour*, 4(4), 353-377.
- Lau, R. R. (1985). Two explanations for negativity effects in political behaviour. *American Journal of Political Science*, 29(1), 119-138.
- Lau, R. R., & Rovner, I. B. (2009). Negative campaigning. *Annual Review of Political Science*, 12, 285-306.

- Lau, R. R., Sigelman, L., & Rovner, I. B. (2007). The effects of negative political campaigns: A meta-analytic reassessment. *The Journal of Politics*, 69(4), 1176-1209.
- Lazarsfeld, P. F., Berelson, B., & Gaudet, H. (1968). *The people's choice: How the voter makes up his mind in a presidential campaign* (3rd ed.). London, UK: Columbia University Press.
- Lazer, D. M. J., Baum, M. A., Benkler, Y., Berinsky, A. J., Greenhill, K. M., Menczer, F., ... Zittrain, J. L. (2018). The science of fake news. *Science*, 359(6380), 1094-1096.
- Leavitt, A. (2013). From #FollowFriday to YOLO: Exploring the cultural salience of Twitter memes. In K. Weller, A. Bruns, J. Burgess, M. Mahrt, & C. Puschmann (Eds.), *Twitter and society* (pp. 137-154). New York, NY: Peter Lang.
- Lee, C. K. M. (2011). Micro-blogging and status updates on Facebook: Texts and practices. In C. Thurlow & K. Mroczek (Eds.), *Digital discourse: Language in the new media* (pp. 110-128). Oxford, UK: Oxford University Press.
- Leech, G. N. (1973). *A linguistic guide to English poetry*. Harlow, UK: Longman.
- Lehmann, J., Gonçalves, B., Ramasco, J. J., & Cattuto, C. (2012). Dynamical classes of collective attention in Twitter. In *International World Wide Web Conference Committee*. New York, NY: ACM.
- Leith, S. (2011). *You talkin' to me?: Rhetoric from Aristotle to Obama*. London, UK: Profile Books.
- Lewin, K. (1947). Frontiers in group dynamics: Concept, method and reality in science; social equilibria and social change. *Human Relations*, 1(1), 5-40.
- Lewis, S., Zamith, R., & Hermida, A. (2013). Content analysis in an era of big data: A hybrid approach to computational and manual methods. *Journal of Broadcasting & Electronic Media*, 57(1), 34-52.
- Lieberman, J. N. (1977). *Playfulness: Its relationship to imagination and creativity*. London, UK: Academic Press.
- Lieberson, S. (2000). *A matter of taste: How names, fashions, and culture change*. London, UK: Yale University Press.
- Lilleker, D. G. (2006). *Key concepts in political communication*. London, UK: Sage Publications.
- Lin, Y., Margolin, D., Keegan, B., Baronchelli, A., & Lazer, D. (2013). #Bigbirds never die: Understanding social dynamics of emergent hashtags. In *Proceedings of the Seventh International Conference on Weblogs and Social Media*. Palo Alto, CA: The AAAI Press. 370-379.
- Loader, B. D., & Mercea, D. (Eds.). (2012). *Social media and democracy: Innovations in participatory politics*. London, UK: Routledge.
- Lombard, M., Snyder-Duch, J., & Bracken, C. (2002). Content analysis in mass communication: Assessment and reporting of intercoder reliability. *Human Communication Research*, 28(4), 587-604.

- Ma, Z., Sun, A., & Cong, G. (2012). Will this #hashtag be popular tomorrow? In *Proceedings of the Thirty-fifth International ACM SIGIR Conference on Research and Development in Information Retrieval*. New York, NY: ACM. 1173-1174.
- Maity, S., Ghuku, B., Upmanyu, A., & Mukherjee, A. (2015). Out of vocabulary words decrease, running texts prevail and hashtags coalesce: Twitter as an evolving sociolinguistic system. *arXiv:1509.05096 [cs.SI]*.
- Malhotra, A., Malhotra, C. K., & See, A. (2012). How to get your messages retweeted. *MIT Sloan Management Review*, 53(2), 61-66.
- Manovich, L. (2001). *The language of new media*. London, UK: The MIT Press.
- Marcus, G. E., Neuman, W. R., & MacKuen, M. (2000). *Affective intelligence and political judgement*. London, UK: The University of Chicago Press.
- Margetts, H., John, P., Hale, S., & Yasseri, T. (2015). *Political turbulence: How social media shape collective action*. Princeton, NJ: Princeton University Press.
- Marlin, R. (2002). *Propaganda and the ethics of persuasion*. Ormskirk, UK: Broadview Press.
- Martin, P. S. (2004). Inside the black box of negative campaign effects: three reasons why negative campaigns mobilize. *Political Psychology*, 25(4), 545-562.
- Marwick, A. E. (2013). *Status update: Celebrity, publicity, and branding in the social media age*. New Haven, CT: Yale University Press.
- Marwick, A. E., & boyd, d. (2011). I tweet honestly, I tweet passionately: Twitter users, context collapse, and the imagined audience. *New Media & Society*, 13(1), 114-133.
- Maybaum, R. (2013). 'Good morning tweethearts!' The diffusion of a lexical innovation in Twitter. *Newcastle University Selected Papers from Sociolinguistics Summer School 4*, 19(1), 128-137.
- Maynard, D., Bontcheva, K., & Rout, D. (2012). Challenges in developing opinion mining tools for social media. In *Proceedings of @NLP can u tag #usergeneratedcontent?! Language Resources and Evaluation Workshop*. Istanbul, Turkey: LREC.
- Maynard, D., & Funk, A. (2012). Automatic detection of political opinions in tweets. In R. Garcia-Castro, D. Fensel, & G. Antoniou (Eds.), *The semantic web: ESWC 2011 workshops* (pp. 88-99). London, UK: Springer.
- McLuhan, M. (1964). *Understanding media: The extensions of man*. New York, NY: McGraw-Hill.
- McMillan, S. J. (2000). The microscope and the moving target: The challenge of applying content analysis to the World Wide Web. *Journalism & Mass Communication Quarterly*, 77(1), 80-98.
- McQuarrie, E. F., & Mick, D. G. (1996). Figures of rhetoric in advertising language. *Journal of Consumer Research*, 22(4), 424-438.

- Medina, E. (2011). *Cybernetic revolutionaries: Technology and politics in Allende's Chile*. London, UK: The MIT Press.
- Messina, E., Fersini, E., & Zammit-Lucia, J. (2017). *All atwitter about Brexit: Lessons for the election campaigns*. London, UK: Radix Group.
- Metaxas, P. T., Mustafaraj, E., Wong, K., Zeng, L., O'Keefe, M., & Finn, S. (2015). What do retweets indicate? Results from user survey and meta-review of research. In *Proceedings of the Ninth International AAAI Conference on Web and Social Media*. Palo Alto, CA: AAAI Press.
- Miles, M. B., & Huberman, M. (1994). *Qualitative data analysis: An expanded sourcebook* (2nd ed.). London, UK: SAGE Publications.
- Miller, C., Arcostanzo, F., Smith, J., Krasodonski-Jones, A., Wiedlitzka, S., Jamali, R., & Dale, J. (2016). *From Brussels to Brexit: Islamophobia, xenophobia, racism, and reports of hateful incidents on Twitter*. London, UK: Demos Centre for the Analysis of Social Media. Retrieved from <https://www.demos.co.uk/wp-content/uploads/2016/07/From-Brussels-to-Brexit-Islamophobia-Xenophobia-Racism-and-Reports-of-Hateful-Incidents-on-Twitter-Research-Prepared-for-Channel-4-Dispatches-%E2%80%98Racist-Britain%E2%80%99-.pdf>
- Min, Y. (2004). News coverage of negative political campaigns: An experiment of negative campaign effects on turnout and candidate preference. *Harvard International Journal of Press/Politics*, 9(4), 95-111.
- Mishler, E. G. (1991). *Research interviewing: Context and narrative*. London, UK: Harvard University Press.
- Mislove, A., Lehmann, S., Ahn, Y., Onnela, J., & Rosenquist, J. N. (2011). Understanding the demographics of Twitter users. In *Proceedings of the Fifth International Conference on Weblogs and Social Media*. Barcelona, Spain.
- Mohammad, S., & Kiritchenko, S. (2014). Using hashtags to capture fine emotion categories from tweets. *Computational Intelligence*, 0(0).
- Morstatter, F., Pfeffer, J., & Liu, H. (2014). When is it biased? Assessing the representativeness of Twitter's Streaming API. *arXiv:1401.7909 [cs.SI]*.
- Mothersbaugh, D. L., Huhmann, B. A., & Franke, G. R. (2002). Combinatory and separative effects of rhetorical figures on consumers' effort and focus in ad processing. *Journal of Consumer Research*, 28(4), 589-602.
- Mulkay, M. (1988). *On humour: Its nature and its place in modern society*. Cambridge, UK: Polity Press.
- Murthy, D. (2013). *Twitter: Social communication in the Twitter age*. Cambridge, UK: Polity Press.
- Myers, G. (2010). *The discourse of blogs and wikis*. London, UK: Continuum International Publishing Group.

- Naaman, M., Becker, H., & Gravano, L. (2011). Hip and trendy: Characterizing emerging trends on Twitter. *Journal of the American Society for Information Science and Technology*, 62(5), 1-26.
- Nahon, K., & Hemsley, J. (2013). *Going viral*. Cambridge, UK: Polity Press.
- Napoli, P. M. (2003). *Audience economics: Media institutions and the audience marketplace*. New York, NY: Columbia University Press.
- Neuendorf, K. A. (2002). *The content analysis guidebook*. Thousand Oaks, CA: Sage Publications.
- Neuman, W. R., Marcus, G. E., MacKuen, M., & Crigler, A. (Eds.). (2007). *The affect effect: Dynamics of emotion in political thinking and behaviour*. Chicago, IL: The University of Chicago Press.
- Noor Al-Deen, H. S., & Hendricks, J. A. (Eds.). (2012). *Social media: Usage and impact*. Lanham, MD: Lexington Books.
- Oktaý, J. S. (2012). *Grounded theory*. Oxford, UK: Oxford University Press.
- Olsen, W. (2012). Interviews. In *Data collection: Key debates and methods in social research* (pp. 33-39). London: SAGE Publications.
- O'Reilly, D. 2005. Cultural brands/branding cultures. *Journal of Marketing Management*, 21(5), 573-588.
- Page, B. I. (1996). *Who deliberates? Mass media in modern democracy*. Chicago, IL: The University of Chicago Press.
- Page, R. (2012). The linguistics of self-branding and micro-celebrity in Twitter: The role of hashtags. *Discourse and Communication*, 6(2), 181-201.
- Palmer, J. (1994). *Taking humour seriously*. London, UK: Routledge.
- Papacharissi, Z. (2012). Without you, I'm nothing: Performances of the self on Twitter. *International Journal of Communication*, 6, 1989-2006.
- Papacharissi, Z. (2015). *Affective publics: sentiment, technology, and politics*. Oxford, UK: Oxford University Press.
- Papacharissi, Z., & Oliveira, M. (2012). Affective news and networked publics: The rhythms of news storytelling on #Egypt. *Journal of Communication*, 62(2), 266-282.
- Pascoe, C. J. (2007). *Dude, you're a fag: Masculinity and sexuality in high school*. London, UK: University of California Press.
- Payne, G., & Williams, M. (2005). Generalization in qualitative research. *Sociology*, 39(2), 295-314.

- Pedregosa, F., Varoquaux, G., Gramfort, A., Michel, V., Thirion, B., Grisel, O., ... Duchesnay, E. (2011). Scikit-learn: Machine learning in Python. *Journal of Machine Learning Research*, 12(10), 2825-2830.
- Penney, J. (2017). *The citizen marketer: Promoting political opinion in the social media age*. Oxford, UK: Oxford University Press.
- Petty, R. E., Cacioppo, J. T., & Heesacker, M. (1981). Effects of rhetorical questions on persuasion: A cognitive response analysis. *Journal of Personality and Social Psychology* 40(3), 432-440.
- Popping, R. (2000). *Computer-assisted text analysis*. London, UK: SAGE Publications.
- Portes, J. (2016). Immigration, free movement, and the EU referendum. *National Institute Economic Review*, 236, 14-22.
- Posch, L., Wagner, C., Singer, P., & Strohmaier, M. (2013). Meaning as collective use: Predicting semantic hashtag categories on Twitter. In *Proceedings of the Twenty-second International Conference on World Wide Web Companion*. Geneva, Switzerland: International World Wide Web Conferences Steering Committee. 621-628.
- Pöschko, J. (2011). Exploring Twitter hashtags. *arXiv:1111.6553 [cs.CL]*.
- Potts, L., Seitzinger, J., Jones, D., & Harrison, A. (2011). Tweeting disaster: Hashtag constructions and collisions. In *Proceedings of the Twenty-ninth ACM International Conference on Design of Communication*. New York, NY: ACM. 235-240.
- Poynter, R. (2010). *The handbook of online and social media research: Tools and techniques for market researchers*. New York, NY: Wiley.
- Pratto, F., & John, O. P. (1991). Automatic vigilance: The attention-grabbing power of negative social information. *Journal of Personality and Social Psychology*, 61(3), 380-391.
- Protess, D., & McCombs, M. E. (Eds.). (1991). *Agenda setting: Readings on media, public opinion, and policymaking*. London, UK: Taylor & Francis.
- Quinlan, L., Shephard, M., & Paterson, L. (2015). Online discussion and the 2014 Scottish independence referendum: Flaming keyboards or forums for deliberation? *Electoral Studies*, 38, 192-205.
- Raschka, S. (2014). Naïve Bayes and text classification I—Introduction and theory. *arXiv:1410.5329v4 [cs.LG]*.
- Rice, R., & Atkin, C. (2011). Communication Campaigns. *Oxford Bibliographies in Communication*. doi: 10.1093/obo/9780199756841-0055
- Riffe, D., Lacy, S. R., & Fico, F. G. (2005). *Analysing media messages: Using quantitative content analysis in research* (2nd ed.). Mahwah, NJ: Lawrence Erlbaum Associates.
- Robson, C. (2011). *Real world research* (3rd ed.). Chichester, UK: Wiley.

- Romero, D., Meeder, B., & Kleinberg, J. (2011). Differences in the mechanics of information diffusion across topics: Idioms, political hashtags, and complex contagion on Twitter. In *Proceedings of the 20th International Conference on World Wide Web*. New York, NY: ACM. 695-704.
- Roseman, I. (1984). Cognitive determinants of emotions: A structural theory. In P. Shaver (Ed.), *Review of personality and social psychology* (Vol. 5, pp. 11-36). Beverly Hill, CA: Sage.
- Roseman, I., Abelson, R., & Ewing, M. (1986). Emotions and political cognition: Emotional appeals in political communication. In R. Lau & D. Sears (Eds.), *Political cognition* (pp. 279-294). London, UK: Lawrence Erlbaum Associates.
- Rudkowsky, E., Haselmayer, M., Wastian, M., Jenny, M., Emrich, Š., & Sedlmair, M. (2018). More than bags of words: Sentiment analysis with word embeddings. *Communication Methods and Measures*, 1-18.
- Rupke, H., & Blank, G. (2009). 'Country roads' to globalization: Sociological Models for Understanding American Popular Music in China. *The Journal of Popular Culture*, 42(1), 126-146.
- Russ, C. (2007). Online crowds—Extraordinary mass behaviour on the internet. In *Proceedings of I-MEDIA '07 and I-SEMANTICS '07*. Graz, Austria: Know-Center.
- Salganik, M. J., Dodds, P. S., & Watts, D. J. (2006). Experimental study of inequality and unpredictability in an artificial cultural market. *Science*, 311(5762), 854-856.
- Sanli, C., & Lambiotte, R. (2015). Local variation of hashtag spike trains and popularity in Twitter. *PLoS ONE*, 10(7).
- Sayre, B., Bode L., Shah, D., Wilcox, D., & Shah, C. (2010). Agenda setting in a digital age: Tracking attention to California Proposition 8. *Policy & Internet*, 2(2). doi: 10.2202/1944-2866.1040
- Schmidt, J. (2013). Twitter and the rise of personal publics. In K. Weller, A. Bruns, J. Burgess, M. Mahrt, & C. Puschmann (Eds.), *Twitter and society* (pp. 3-14). New York, NY: Peter Lang.
- Schnoebelen, T. (2012). Gender, styles, and social networks in Twitter. Presented at *New Ways of Analysing Variation 41*. Bloomington, IN: Indiana University.
- Schroeder, R. (2014). Big data: Towards a more scientific social science and humanities? In M. Graham & W. H. Dutton (Eds.), *Society and the internet: How networks of information and communication are changing our lives* (pp. 164-176). Oxford, UK: Oxford University Press.
- Schudson, M. (1989). How culture works: Perspectives from media studies on the efficacy of symbols. *Theory and Society*, 18(2), 153-180.
- Segerberg, A., & Bennett, W. L. (2011). Social media and the organisation of collective action: Using Twitter to explore the ecologies of two climate change protests. *The Communication Review*, 14(3), 197-215.

- Settles, B. (2011). Closing the loop: Fast, interactive semi-supervised annotation with queries on features and instances. In *Proceedings of the 2011 Conference on Empirical Methods in Natural Language Processing*. Stroudsburg, PA: ACL. 1467-1478.
- Shao, G. (2009). Understanding the appeal of user-generated media: A uses and gratification perspective. *Internet Research*, 19(1), 7-25.
- Shoemaker, P. J., Vos, T. P., & Reese, S. D. (2009). Journalists as gatekeepers. In K. Wahl-Jorgensen & T. Hanitzsch (Eds.), *The handbook of journalism studies* (pp. 73-87). Abingdon, UK: Routledge.
- Simmel, G. (1902). The number of persons as determining the sociological form of the group. *American Journal of Sociology*, 8, 1-46.
- Simon, H. A. (1971). Designing organisations for an information-rich world. In M. Greenberger (Ed.), *Computers, communication, and the public interest* (pp. 32-72). Baltimore, Maryland: Johns Hopkins Press.
- Simons, H. W. (1970). Requirements, problems, and strategies: A theory of persuasion for social movements. *Quarterly Journal of Speech*, 56(1), 1-11.
- Simons, H. W. (1976). *Persuasion: Understanding, practice, and analysis*. Boston, MA: Addison-Wesley.
- Simons, H. W. (2004). The rhetorical legacy of Kenneth Burke. In W. Jost & W. Olmsted (Eds.), *A companion to rhetoric and rhetorical criticism* (pp. 152-167). Oxford, UK: Blackwell Publishing.
- Sinclair, J. (2005). Corpus and text: Basic principles. In M. Wynne (Ed.), *Developing linguistic corpora: A guide to good practice* (pp. 1-16). Oxford, UK: Oxbow Books.
- Singer, J. B. (1997). Still guarding the gate?: The newspaper journalist's role in an online world. *Convergence: The International Journal of Research into New Media Technologies*, 3(1), 72-89.
- Sjøvaag, H., & Stavelin, E. (2012). Web media and the quantitative content analysis: Methodological challenges in measuring online news content. *Convergence: The International Journal of Research into New Media Technologies*, 18(2), 215-229.
- Smith, A., Stavros, C., Westberg, K., Wilson, B., & Boyle, C. (2014). Alcohol-related player behavioural transgressions: Incidences, fan media responses, and a harm-reduction alternative. *International Review for the Sociology of Sport*, 49(3/4), 400-416.
- Sterne, J. (2012). *MP3: The meaning of a format*. Durham, NC: Duke University Press.
- Stewart, C. J., Smith, C.A., & Denton, R. E. (2007). *Persuasion and social movements*. Long Grove, IL: Waveland Press.
- Su, H. (2007). The multilingual and multiorthographic Taiwan-based internet: Creative uses of writing systems on college-affiliated BBSs. In B. Danet & S. C. Herring (Eds.), *The multilingual internet: Language, culture, and communication online* (pp. 64-86). Oxford, UK: Oxford University Press.

- Sundar, S. S. (2008). Self as source: Agency and customization in interactive media. In E. A. Konijn, S. Utz, M. Tanis, & S. B. Barnes (Eds.), *Mediated interpersonal communication* (pp. 58-74). Abingdon, UK: Routledge.
- Sundar, S. S., Oh, J., Kang, H., & Sreenivasan, A. (2013). How does technology persuade? Theoretical mechanisms for persuasive technologies. In J. P. Dillard & L. Shen (Eds.), *The SAGE handbook of persuasion: Developments in theory and practice* (2nd ed., pp. 388-404). London, UK: SAGE Publications.
- Swidler, A. (1986). Culture in action: Symbols and strategies. *American Sociological Review*, 51(2), 273-286.
- Symes, A. (2013). *A linguistic introduction to the discursive conventions of Twitter* (Doctoral dissertation). Retrieved from the Gothenburg University Publications Electronic Archive. <http://hdl.handle.net/2077/32982>
- Tang, D., Wei, F., Yang, N., Zhou, M., Liu, T., & Qin, B. (2014). Learning sentiment-specific word embedding for Twitter sentiment classification. *ACL*, 1, 1555-1565.
- Thelwall, M., Goriunova, O., Vis, F., Faulkner, S., Burns, A., Aulich, J., Mas-Bleda, A., Stuart, E., & D'Orazio, F. (2015). Chatting through pictures? A classification of images tweeted in one week in the UK and USA. *Journal of the Association for Information Science and Technology*, 66(1).
- Tinkham, S. F., & Weaver-Lariscy, R. A. (1990). Advertising message strategy in U.S. congressional campaigns: Its impact on election outcome. *Current Issues & Research in Advertising*, 13(1), 207-226.
- Tsur, O., & Rappoport, A. (2012). What's in a hashtag? Content-based prediction of the spread of ideas in microblogging communities. In *Proceedings of the Fifth ACM International Conference on Web Search and Data Mining*. New York, NY: ACM. 643-652.
- Tumasjan, A., Sprenger, T. O., Sandner, P. G., & Welpe, I. M. (2010). Predicting elections with Twitter: What 140 characters reveal about political sentiment. In *Proceedings of the Fourth International AAAI Conference on Weblogs and Social Media*. Menlo Park, CA: The AAAI Press.
- Tumasjan, A., Sprenger, T. O., Sandner, P. G., & Welpe, I. M. (2012). Where there is a sea there are pirates: Response to Jungherr, Jürgens, and Schoen. *Social Science Computer Review*, 30(2), 235-239.
- Usherwood, S., & Wright, K. (2017). Twitter campaigning strategies in the European Union referendum. *The British Journal of Politics and International Relations*, 19(2), 371-388.
- Van Dijck, J. (2013). *The culture of connectivity: A critical history of social media*. Oxford, UK: Oxford University Press.
- van Zoonen, W., & van der Meer, T. (2016). Social media research: The application of supervised machine learning in organisational communication research. *Computers in Human Behaviour*, 63, 132-141.

- Vis, F. (2013). Twitter as a reporting tool for breaking news: Journalists tweeting the 2011 UK riots. *Digital Journalism*, 1(1), 27-47.
- Vis, F., Faulkner, S., Parry, K., Manyukhina, Y., & Evans, L. (2013). Twitpic-ing the riots: analysing images shared on Twitter during the 2011 UK riots. In K. Weller, A. Bruns, J. Burgess, M. Mahrt, & C. Puschmann (Eds.), *Twitter and society* (pp. 385-398). New York, NY: Peter Lang.
- Vitter, J. S. (1985). Random sampling with a reservoir. *ACM Transactions on Mathematical Software*, 11(1), 37-57.
- Vosoughi, S., Roy, D., & Aral, S. (2018). The spread of true and false news online. *Science*, 359(6380), 1146-1151.
- Wang, X., Wei, F., Liu, X., Zhou, M., & Zhang, M. (2011). Topic sentiment analysis in Twitter: A graph-based hashtag sentiment classification approach. In *Proceedings of the Twentieth ACM International Conference on Information and Knowledge Management*. New York, NY: ACM. 1031-1040.
- Warnick, B., & Heineman, D. (2012). *Rhetoric online: The politics of new media* (2nd ed.). Oxford, UK: Peter Lang.
- Watson, C. (2006). Unreliable narrators? 'Inconsistency' (and some inconstancy) in interviews. *Qualitative Research*, 6(3), 367-384.
- Weaver-Lariscy, R. A., & Tinkham, S. F. (1996). Advertising message strategies in U.S. congressional campaigns: 1982, 1990. *Journal of Current Issues & Research in Advertising*, 18(1), 53-66.
- Webb, H., Jirotko, M., Procter, R., Stahl, B., Rana, O., Burnap, P., Housley, W., Edwards, A., & Williams, M. (2017). The ethical challenges of publishing Twitter data for research dissemination. In *Proceedings of the 2017 ACM on Web Science Conference*. New York, NY: ACM.
- Weber, I., Garimella, V., & Teka, A. (2013). Political hashtag trends. In P. Serdyukov, P. Braslavski, S. O. Kuznetsov, J. Kamps, S. Rüger, E. Agichtein, I. Segalovich, & E. Yilmaz (Eds.), *Advances in information retrieval* (pp. 857-860). Berlin, Germany: Springer.
- Weber, R. P. (1990). *Basic content analysis*. Newbury Park, CA: Sage University Paper Series on Quantitative Applications in the Social Sciences.
- Weimann, G., Weiss-Blatt, N., Mengistu, G., Tregerman, M. M., & Oren, R. (2014). Reevaluating 'The End of Mass Communication?', *Mass Communication and Society*, 17(6), 803-829.
- Weller, K., Bruns, A., Burgess, J., Mahrt, M., & Puschmann, C. (Eds.). (2013). *Twitter and society*. Oxford, UK: Peter Lang.
- Weng, L., Flammini, A., Vespignani, A., & Menczer, F. (2012). Competition among memes in a world with limited attention. *Nature Scientific Reports*, 2(335).

- Weng, L., Menczer, F., & Ahn, Y. (2013). Virality prediction and community structure in social networks. *Nature Scientific Reports*, 3(2522).
- West, D. M. (2005). *Air Wars: Television advertising in election campaigns, 1952-2004* (4th ed.). Washington, DC: CQ Press.
- White, D. M. (1950). The 'gate keeper': A case study in the selection of news. *Journalism Quarterly*, 27(4), 383-390.
- Wilkerson, C. (1976). A rhetorical definition of movements. *Central States Speech Journal*, 27, 88-94.
- Williams, B. A., & Carpin, M. X. D. (2004). Monica and Bill all the time and everywhere: The collapse of gatekeeping and agenda setting in the new media environment. *American Behavioral Scientist*, 47(9), 1208-1230.
- Williams, M., Webb, H., Housley, W., Jirotk, M., Burnap, P., Williams, M., Edwards, A., Rana, O., Procter, R., & Stahl, B. (2016). 'We have much more fighting to do': Analysis of Twitter-based online campaigns. Manuscript in preparation, Department of Computer Science, University of Oxford, Oxford, United Kingdom.
- Woolley, S., & Howard, P. (Eds.). (2018). *Computational propaganda: Political parties, politicians, and political manipulation on social media*. Oxford, UK: Oxford University Press.
- Wu, T. (2016a). Attention brokers. Presented at the Innovation Policy Colloquium, NYU. Retrieved from http://www.law.nyu.edu/sites/default/files/upload_documents/Tim%20Wu%20-%20Attention%20Brokers.pdf
- Wu, T. (2016b). *The attention merchants: From the daily newspaper to social media, how our time and attention is harvested and sold*. London, UK: Atlantic Books.
- Yang, L., Sun, T., Zhang, M., & Mei, Q. (2012). We know what @you #tag: Does the dual role affect hashtag adoption? In *Proceedings of the 21st International Conference on World Wide Web*. New York, NY: ACM. 261-270.
- Yardi, S., Romero, D., Schoenebeck, G., & boyd, d. (2010). Detecting spam in a Twitter network. *First Monday*, 15(1).
- Yin, R. (1981). The case study crisis: Some answers. *Administrative Science Quarterly*, 26(1), 58-65.
- Yin, R. (1994). Case study research (2nd ed.). Thousand Oaks, CA: SAGE Publications.
- Yuce, S., Agarwal, N., Wigand, R. T., Lim, M., & Robinson, R. S. (2014). Studying the evolution of online collective action: Saudi Arabian women's 'Oct26Driving' Twitter campaign. In *Proceedings of the SBP 2014 Conference on Social Computing, Behavioural Modelling, and Prediction*, Lecture Notes in Computer Science 8393. Cham, Switzerland: Springer. 413-420.
- Zappavigna, M. (2011). Ambient affiliation: A linguistic perspective on Twitter. *New Media & Society*, 13(5), 788-806.

- Zappavigna, M. (2012). *Discourse of Twitter and social media*. London, UK: Continuum.
- Zappavigna, M. (2015). Searchable talk: the linguistic functions of hashtags. *Social Semiotics*, 25(3), 274-291.
- Zhao, D., & Rosson, M. B. (2009). How and why people Twitter: The role that micro-blogging plays in informal communication at work. In *Proceedings of the ACM 2009 International Conference on Supporting Group Work*. New York, NY: ACM. 243-252.
- Zillmann, D. (1972). Rhetorical elicitation of agreement in persuasion. *Journal of Personality and Social Psychology*, 21(2), 159-165.

9.2 Website Articles

- American Dialect Society. (2013, January 4). 'Hashtag' is the 2012 word of the year. Retrieved from <http://www.american-dialect.org/hashtag-2012>
- American Dialect Society. (2015, January 9). 2014 word of the year is '#blacklivesmatter'. Retrieved from <http://www.american-dialect.org/2014-word-of-the-year-is-blacklivesmatter>
- Bar-Joseph, U. (2014, February 10). 5 brilliant ways to use #hashtags in social media marketing. *Search Engine Watch*. Retrieved from <http://searchenginewatch.com/sew/how-to/2327748/5-brilliant-ways-to-use-hashtags-in-social-media-marketing>
- BBC News. (2013, 23 January). David Cameron promises in/out referendum on EU. Retrieved from <https://www.bbc.com/news/uk-politics-21148282>
- BBC News. (2015, 21 May). Timeline: Campaigns for a European Union referendum. Retrieved from <https://www.bbc.co.uk/news/uk-politics-15390884>
- BBC News. (2016a, 12 January). EU referendum timeline: What will happen when? Retrieved from <http://www.bbc.co.uk/news/uk-politics-33141819>
- BBC News. (2016b, 8 February). EU referendum: How will official lead campaigns be chosen? Retrieved from <http://www.bbc.co.uk/news/uk-politics-eu-referendum-35522695>
- BBC News. (2016c, 20 February). EU referendum timeline: Countdown to the vote. Retrieved from <https://www.bbc.co.uk/news/uk-politics-33141819>
- BBC News. (2016d, April 13). Lead EU referendum campaigns named. Retrieved from <https://www.bbc.co.uk/news/uk-politics-36038672>
- Beer, J. (2017, May 9). Carter got his Wendy's nuggs and a Twitter world record. *Fast Company*. Retrieved from <https://www.fastcompany.com/40419538/carter-got-his-wendys-nuggs-and-a-twitter-world-record>

- Beese, J. (2013, February 11). Twitter promoted trends will now cost you \$200,000 daily. *Sprout Social*. Retrieved from <http://sproutsocial.com/insights/twitter-promoted-trends-cost/>
- Boag, P. (2017, February 21). Call to action: The 10 most effective techniques. Boagworks. Retrieved from <https://boagworld.com/design/call-to-action>
- Borthwick, J. (2015, March 7). Media hacking. *Medium*. Retrieved from <https://render.betaworks.com/media-hacking-3b1e350d619c>
- Bosker, B. (2014, May 10). How #hashtags became the raised fist of punctuation. *HuffPost Tech*. Retrieved from http://www.huffingtonpost.com/2014/05/10/bringbackourgirls_n_5298223.html
- Bosker, B. (2016, November). The binge breaker. *The Atlantic*. Retrieved from <https://www.theatlantic.com/magazine/archive/2016/11/the-binge-breaker/501122/>
- Botometer. (2018, November). FAQ. Retrieved from <https://botometer.iuni.iu.edu/#!/faq#how-does-it-works>
- Botsman, R. (2015, October 20). The changing rules of trust in the digital age. *Harvard Business Review*. Retrieved from <https://hbr.org/2015/10/the-changing-rules-of-trust-in-the-digital-age>
- Boyd, S. (2007, August 28). Hash tags = Twitter groupings. Retrieved from <http://stoweboyd.com/post/39877198249/hash-tags-twitter-groupings>
- Brooks, R. (2010, December). What is Paper.li? A newspaper of tweets? Retrieved from <https://www.takeflyte.com/flyte/2010/12/what-is-paper-li-a-newspaper-of-tweets.html>
- Byford, S. (2019, January 6). SpaceX moon tourist scores most retweeted tweet ever with promise of free money. *The Verge*. Retrieved from <https://www.theverge.com/2019/1/6/18171386/most-retweeted-tweet-ever-yusaku-maczawa-spacex>
- Chacon, B. (2018, December 9). How to calculate your Instagram engagement rate. Retrieved from <https://later.com/blog/instagram-engagement-rate/>
- Chakelian, A. (2015). The EU referendum campaign groups: who's In and who's Out? *New Statesman*. Retrieved from <http://www.newstatesman.com/world/2015/10/eu-referendum-campaign-groups-who-s-and-who-s-out>
- Communication Arts. (2013). Hashtag killer. Retrieved from <http://www.commartarts.com/interactive/cai13/hashtagkiller.html>
- Confessore, N., & Yourish, K. (2016, March 15). \$2 billion worth of free media for Donald Trump. *The New York Times*. Retrieved from <https://www.nytimes.com/2016/03/16/upshot/measuring-donald-trumps-mammoth-advantage-in-free-media.html>

- Cullum, B. (2010, December 17). What makes a Twitter hashtag successful? Retrieved from <http://www.movements.org/blog/entry/what-makes-a-twitter-hashtag-successful>
- Custom Fit Online. (2014, September 5). Too many hashtags can be a bad practice. Retrieved from <https://www.customfitonline.com/news/2014/9/5/too-many-hashtags-bad-practice>
- Dumenco, S. (2010, May 21). The real reason Twitter radically reworked its trending topics algorithm. *AdAge*. Retrieved from <https://adage.com/article/trending-topics/twitter-radically-reworked-trending-topics-algorithm/144023>
- The Electoral Commission. (2017). EU referendum results. Retrieved from <https://www.electoralcommission.org.uk/find-information-by-subject/elections-and-referendums/past-elections-and-referendums/eu-referendum/electorate-and-count-information>
- Elgot, J. (2016, June 9). EU referendum debates: when and where to watch them. *The Guardian*. Retrieved from <https://www.theguardian.com/politics/2016/jun/02/eu-referendum-tv-debates-when-where-watch-them>
- Ellering, N. (2017, October 16). What 20 studies say about the best times to post on social media. *CoSchedule Blog*. Retrieved from <https://coschedule.com/blog/best-times-to-post-on-social-media>
- Empson, R. (2013, September 19). Chirpify drops its 'hashtag commerce' consumer marketplace to focus on big brand marketing and 'actiontags'. *TechCrunch*. Retrieved from <http://techcrunch.com/2013/09/19/chirpify-drops-its-hashtag-commerce-consumer-marketplace-to-focus-on-big-brand-marketing-and-actiontags>
- European Union. (2017). The history of the European Union. *Europa.eu*. Retrieved from https://europa.eu/european-union/about-eu/history_en
- Grammarly. (2017). Imperative verbs: Definition and examples. *Grammarly Blog*. Retrieved from <https://www.grammarly.com/blog/imperative-verbs>
- Greenfield, R. (2013, July 25). A complete guide to the art of Twitter favourites. *The Atlantic*. Retrieved from <https://www.theatlantic.com/technology/archive/2013/07/complete-guide-art-twitter-favorites/312884>
- Greenwald, T. (2014, December 23). How Facebook and Candy Crush got you hooked. *Wired*. Retrieved from <https://www.wired.com/2014/12/how-to-build-habit-forming-products>
- Greenwood, S., Perrin, A., & Duggan, M. (2016, November 11). Social media update 2016. *Pew Research Center*. Retrieved from <http://www.pewinternet.org/2016/11/11/social-media-update-2016>
- Grimes, T. (2013, June 13). How Facebook hashtags can impact brands. *Social Media Today*. Retrieved from <http://socialmediatoday.com/timgrimes/1532906/how-facebook-hashtags-can-impact-brands>

- Harris, T. (2016, May 19). How technology hijacks people's minds—from a magician and Google's design ethicist. Retrieved from <http://www.tristanharris.com/essays/>
- Hendy, D. (2018). Smethwick. *BBC*. Retrieved from <https://www.bbc.co.uk/historyofthebbc/memories/elections/smethwick>
- Hodson, T. (2018, September 21). Peston launches new Wednesday show. *itv*. Retrieved from <https://www.itv.com/presscentre/press-releases/peston-launches-new-wednesday-show>
- Hunt, E. (2016, December 17). What is fake news? How to spot it and what you can do to stop it. *The Guardian*. Retrieved from <https://www.theguardian.com/media/2016/dec/18/what-is-fake-news-pizzagate>
- Hutchinson, A. (2016, October 21). New report looks at optimal hashtag use on Twitter, Instagram and Facebook. *Social Media Today*. Retrieved from <http://www.socialmediatoday.com/social-networks/new-report-looks-optimal-hashtag-use-twitter-instagram-and-facebook>
- Internet Live Stats. (2018, 12 December). Twitter usage statistics. Retrieved from <http://www.internetlivestats.com/twitter-statistics>
- Jones, A. A. (2016, November 1). Challenging 'slacktivism': Activism on social media is not enough. *Huffington Post*. Retrieved from http://www.huffingtonpost.com/entry/challenging-slacktivism-activism-on-social-media_us_5817c2dbe4b09b190529c8ae
- Karr, D. (2013, January 28). 13 tools for hashtag research and management. Retrieved from <http://www.marketingtechblog.com/hashtag-research-tools>
- Kaufman, A. (2015, March 21). Jack Dorsey credits journalists with Twitter's Rise. *Huffington Post*. Retrieved from http://www.huffingtonpost.com/2015/03/21/jack-dorsey-journalists_n_6916508.html
- Knibbs, K. (2015). The FBI says retweets are endorsements. *Gizmodo*. Retrieved from <https://gizmodo.com/the-fbi-says-retweets-are-endorsements-1731526051>
- Koman, T. (2018, July 24). Checking in with Carter Wilkerson, the kid who got free Wendy's nuggets for a year. *Delish*. Retrieved from <https://www.delish.com/food-news/a22529419/where-are-they-now-wendys-nuggets-tweet-carter-wilkerson>
- Kumar, A. (2015, November 3). Hearts on Twitter. Retrieved from https://blog.twitter.com/official/en_us/a/2015/hearts-on-twitter.html
- Lenhart, A., & Fox, S. (2009, February 12). Twitter and status updating. *Pew Research Centre*. Retrieved from www.pewinternet.org/Reports/2009/Twitter-and-status-updating.aspx
- Liu, E. (2015, November 17). TF-IDF, Term frequency-inverse document frequency. Retrieved from https://ethen8181.github.io/machine-learning/clustering_old/tf_idf/tf_idf.html

- Lohr, S. (2018, March 8). It's true: False news spreads faster and wider. And humans are to blame. *The New York Times*. Retrieved from <https://www.nytimes.com/2018/03/08/technology/twitter-fake-news-research.html>
- Lunden, I. (2018, October 29). No, Twitter is not removing its like button 'soon', but it is considering it. *TechCrunch*. Retrieved from <https://techcrunch.com/2018/10/29/twitter-heart-break>
- Madrigal, A. C. (2013). How to actually get a job on Twitter. *The Atlantic*. Retrieved from <https://www.theatlantic.com/technology/archive/2013/07/how-to-actually-get-a-job-on-twitter/278246>
- Madu, Z. (2017). The important questions: Why the quote-tweet is the absolute worst Twitter function. *National Post*. Retrieved from <https://nationalpost.com/life/the-important-questions-why-the-quote-tweet-is-the-absolute-worst-twitter-function>
- Mann, R. (2016). How the 'Daisy' ad changed everything about political advertising. *Smithsonian*. Retrieved from <https://www.smithsonianmag.com/history/how-daisy-ad-changed-everything-about-political-advertising-180958741>
- Manson, M. (2014, December 4). In the future, our attention will be sold. Retrieved from <https://markmanson.net/attention>
- Marketing Terms. (2017). Call to action. Retrieved from https://www.marketingterms.com/dictionary/call_to_action
- Martínez, A. (2018a, February 23). How Trump conquered Facebook—without Russian ads. *Wired*. Retrieved from <https://www.wired.com/story/how-trump-conquered-facebookwithout-russian-ads>
- Martínez, A. (2018b, March 2). Facebook's equal time quandary. *Wired*. Retrieved from <https://www.wired.com/story/facebook-equal-time-advertising>
- Messina, C. (2007, August 25). Groups for Twitter; or a proposal for Twitter tag channels. Retrieved from <http://factoryjoe.com/blog/2007/08/25/groups-for-twitter-or-a-proposal-for-twitter-tag-channels>
- Messina, C. (2007, October 22). Twitter hashtags for emergency coordination and disaster relief. Retrieved from <http://factoryjoe.com/blog/2007/10/22/twitter-hashtags-for-emergency-coordination-and-disaster-relief>
- Messina, C. (2007, December 23). Making the most of hashtags. Retrieved from <http://factoryjoe.com/blog/2007/12/23/making-the-most-of-hashtags>
- Meyer, R. (2018, March 8). The grim conclusions of the largest-ever study of fake news. *The Atlantic*. Retrieved from <https://www.theatlantic.com/technology/archive/2018/03/largest-study-ever-fake-news-mit-twitter/555104>

- Micallef, J. (2016). Brexit: Consequences and significance. *The World Post*. Retrieved from http://www.huffingtonpost.com/joseph-v-micallef/brexit-consequences-and-s_b_10659356.html
- Moore, R. J. (2009, October 5). Twitter data analysis: An investor's perspective. *TechCrunch*. Retrieved from <https://techcrunch.com/2009/10/05/twitter-data-analysis-an-investors-perspective-2>
- Moreau, E. (2017, May 30). 12 very funny Twitter parody accounts. *Lifewire*. Retrieved from <https://www.lifewire.com/very-funny-twitter-parody-accounts-3486602?i10c>
- Morgans, J. (2017, May 21). The secret ways social media is built for addiction. *Vice*. Retrieved from https://www.vice.com/en_uk/article/the-secret-ways-social-media-is-built-for-addiction
- Murphy, M. (2016, July 6). Twitter Live is a new thing from Twitter that actually seems useful. *Quartz*. Retrieved from <https://qz.com/724892/twitter-live-is-a-new-thing-from-twitter-that-actually-seems-useful>
- National Union of Journalists. (2015, July 1). BBC facts and stats. Retrieved from <https://www.nuj.org.uk/news/bbc-facts-and-stats>
- Newton, C. (2015, November 3). Twitter officially kills off favorites and replaces them with likes. *The Verge*. Retrieved from <https://www.theverge.com/2015/11/3/9661180/twitter-vine-favorite-fav-likes-hearts>
- Orlean, S. (2010, June 29). Hash. *The New Yorker*. Retrieved from <http://www.newyorker.com/online/blogs/susanorlean/2010/06/hash.html>
- Oxford Dictionaries. (2014, June 13). From flexitarian to evil genius: New words in the OED. *OxfordWords Blog*. Retrieved from <http://blog.oxforddictionaries.com/2014/06/oed-update-june-2014>
- Oxford Dictionaries. (2016). Word of the Year 2016 is... Retrieved from <https://en.oxforddictionaries.com/word-of-the-year/word-of-the-year-2016>
- Parker, A. (2011, June 10). Twitter's secret handshake. *The New York Times*. Retrieved from <http://www.nytimes.com/2011/06/12/fashion/hashtags-a-new-way-for-tweets-cultural-studies.html?smid=tw-nytimes>
- Parsons, J. (2017, April 1). What is the ideal Twitter follower to following ratio? Retrieved from <https://follows.com/blog/2017/04/ideal-follower-following-ratio>
- Perez, S. (2017, November 7). Twitter officially expands its character count to 280 starting today. *TechCrunch*. Retrieved from <https://techcrunch.com/2017/11/07/twitter-officially-expands-its-character-count-to-280-starting-today>
- Perez, S. (2018, February 28). Twitter launches Bookmarks, a private way to save tweets. *TechCrunch*. Retrieved from <https://techcrunch.com/2018/02/28/twitter-launches-bookmarks-a-private-way-to-save-tweets>

- Public Media Alliance. (2017). British Broadcasting Corporation. Retrieved from <https://publicmediaalliance.org/member-page/british-broadcasting-corporation>
- Rambukkana, N. (2013, June 29). CFP—Hashtag Publics: The Power and Politics of Networked Discourse Communities. Retrieved from <http://complexsingularities.net/call-for-papers-hashtag-publics>
- Rusbridger, A. (2010, November 19). Why Twitter matters for media organisations. *The Guardian*. Retrieved from <http://www.theguardian.com/media/2010/nov/19/alan-rusbridger-twitter>
- Schofield, J. (2018, September 18). Twitter brings back chronological timelines. *ZDNet*. Retrieved from <https://www.zdnet.com/article/twitter-brings-back-chronological-timelines>
- Seward, Z. (2014, August 11). Twitter admits that as many as 23 million of its active users are automated. *Quartz*. Retrieved from <http://qz.com/248063/twitter-admits-that-as-many-as-23-million-of-its-active-users-are-actually-bots>
- Smith, K. (2017, December 17). 45 incredible and interesting Twitter statistics. *Brandwatch*. Retrieved from <https://www.brandwatch.com/blog/44-twitter-stats>
- Sorgatz, R. (2008, June 17). The microfame game. *New York Magazine*. Retrieved from <http://nymag.com/news/media/47958>
- Smith, N. (2015, April 27). Why is the NHS so important? BBC. Retrieved from <http://www.bbc.co.uk/programmes/p02q2kvn>
- Statista. (2018). Number of monthly active Twitter users worldwide from 1st quarter 2010 to 3rd quarter 2018 (in millions). Retrieved from <https://www.statista.com/statistics/282087/number-of-monthly-active-twitter-users/>
- Stewart, H., & Mason, R. (2016, June 16). Nigel Farage's anti-migrant poster reported to police. *The Guardian*. Retrieved from <https://www.theguardian.com/politics/2016/jun/16/nigel-farage-defends-ukip-breaking-point-poster-queue-of-migrants>
- Stone, B. (2009, November 19). What's happening? *Twitter Blogs*. Retrieved from <https://blog.twitter.com/2009/whats-happening>
- Tiwari, P., & Venkataramakrishnan, S. (2018, February 5). Brexit 2018: Insights Analysis—Jan Edition. *Cronycle*. Retrieved from <https://cronycle.com/insights/brexit-2018-jan-edition>
- Twitter. (2018). About. Retrieved from <https://about.twitter.com>
- Twitter Developers. (2014). Frequently asked questions. Retrieved from <https://dev.twitter.com/docs/faq>

- Twitter Developers. (2018a). Optimise Tweets with Cards. Retrieved from <https://developer.twitter.com/en/docs/tweets/optimize-with-cards/overview/abouts-cards>
- Twitter Developers. (2018b). Search Tweets. Retrieved from <https://developer.twitter.com/en/docs/tweets/search/overview>
- Twitter Help Centre. (2017a). About replies and mentions. Retrieved from <https://support.twitter.com/articles/14023>
- Twitter Help Centre. (2017b). Liking a tweet or Moment. Retrieved from <https://support.twitter.com/articles/20169874>
- Twitter Help Centre. (2017c). Request to verify an account. Retrieved from <https://support.twitter.com/articles/20174631>
- Twitter Help Centre. (2018a). About your activity dashboard. Retrieved from <https://help.twitter.com/en/managing-your-account/using-the-tweet-activity-dashboard>
- Twitter Help Centre. (2018b). How to create a thread on Twitter. Retrieved from <https://help.twitter.com/en/using-twitter/create-a-thread>
- Twitter Terms of Service. (2018). Twitter Terms of Service. Retrieved from <https://twitter.com/en/tos>
- Vander Wal, T. (2007, February 2). Folksonomy. Retrieved from <http://vanderwal.net/folksonomy.html>
- Vote Leave. (2016). Let's give our NHS the £350 million the EU takes every week. LSE Digital Library. Retrieved from <https://digital.library.lse.ac.uk/objects/lse:paf829reh>
- Waldman, K. (2014, May 6). Hashtags are the new scare quotes. *Slate*. Retrieved from http://www.slate.com/blogs/future_tense/2014/05/06/_blessed_hashtags_are_the_new_scare_quotes.html
- Warzel, C. (2014, April 15). Meet the man behind Twitter's most infamous phrase. *BuzzFeed News*. Retrieved from <https://www.buzzfeednews.com/article/charliewarzel/meet-the-man-behind-twitters-most-infamous-phrase>
- Weiss, S. (2014, May 26). The power of #YesAllWomen. *The New Yorker*. Retrieved from <http://www.newyorker.com/online/blogs/culture/2014/05/the-power-of-yesallwomen.html>
- Williams, J. (2016, June 23). EU referendum: What are the main arguments for and against leaving the European Union? *Manchester Evening News*. Retrieved from <http://www.manchestereveningnews.co.uk/news/greater-manchester-news/eu-referendum-issues-for-against-11020375>
- Wintour, P., & Watt, N. (2014, May 26). UKIP wins European elections with ease to set off political earthquake. *The Guardian*. Retrieved from

<https://www.theguardian.com/politics/2014/may/26/ukip-european-elections-political-earthquake>

Wishpond. (2017). The 25 best words to use in your call-to-action buttons. Retrieved from <http://blog.wishpond.com/post/103290853633/the-25-best-words-to-use-in-your-call-to-action>