CRANFIELD UNIVERSITY

Lorna Walker

What factors influence whether politicians' tweets are

retweeted?

Using CHAID to build an explanatory model of the retweeting of

politicians' tweets during the 2015 UK General Election campaign

SCHOOL OF MANAGEMENT

PhD THESIS

Academic Year 2016-17

Supervisor: Professor Paul Baines

September 2016

This thesis is submitted in partial fulfilment of the requirements
for the degree of Doctor of Philosophy

Abstract

Twitter is ever-present in British political life and many politicians use it as part of their campaign strategies. However, little is known about whether their tweets engage people, for example by being retweeted. This research addresses that gap, examining tweets sent by MPs during the 2015 UK General Election campaign to identify which were retweeted and why.

A conceptual model proposes three factors which are most likely to influence retweets: the characteristics of (1) the tweet's sender, (2) the tweet and (3) its recipients. This research focuses on the first two of these. Content and sentiment analysis are used to develop a typology of the politicians' tweets, followed by CHAID analysis to identify the factors that best predict which tweets are retweeted.

The research shows that the characteristics of tweet and its sender do influence whether the tweet is retweeted. Of the sender's characteristics, number of followers is the most important – more followers leads to more retweets. Of the tweet characteristics, the tweet's sentiment is the most influential. Negative tweets are retweeted more than positive or neutral tweets. Tweets attacking opponents or using fear appeals are also highly likely to be retweeted.

The research makes a methodological contribution by demonstrating how CHAID models can be used to accurately predict retweets. This method has not been used to predict retweets before and has broad application to other contexts. The research also contributes to our understanding of how politicians and the public interact on Twitter, an area little studied to date, and proposes some practical recommendations regarding how MPs can improve the effectiveness of their Twitter campaigning. The finding that negative tweets are more likely to be retweeted also contributes to the ongoing debate regarding whether people are more likely to pass on positive or negative information online.

Acknowledgements

When I started work on this PhD I had no idea what a huge undertaking it would be, and how reliant I would be on the support of friends, family and colleagues for both practical help and emotional support.

Thanks must go to all at Cranfield School of Management who have supported me on this journey. In particular, to my panel members Radu Dimitriou and Emma MacDonald and to my supervisor Paul Baines. It's taken many twists and turns and a few dead ends to get to this point and all three of them have been supportive, encouraging and generous with their time and expertise throughout.

This thesis has stretched my brain in ways that I didn't know it could be stretched. Thanks are due to my amazingly knowledgeable friends and colleagues who helped me develop this project in all kinds of practical ways. In particular, I'm grateful to Jarlath Quinn for his sage advice about using SPSS Modeler, Anabel Gutierrez and Andy Neal for allowing me to bounce my statistics questions off them, and Steve Hanlon for help accessing the Twitter API and allowing me to mither him with my technical questions. Additional thanks are also due to Andy, Steve and my mother, Valerie Walker, for their proof reading help.

My friends and family have had a lot to put up with over the last six years. There have been many cancelled engagements and many a well-meaning enquiry about how the PhD is going has been met with grumpy mumbling (thanks are due here to Terry and Jenny Burridge their efforts to save everyone from a lot of awkwardness by getting "Don't ask me about my PhD!" emblazoned on a mug for me). I really couldn't have done it without the support and love of Valerie Walker, Neil Walker and, most of all, Matthias Feist. Thank you.

Table of contents

List of figures

List of tables

Chapter 1 Overview of the research topic

1.1.   Chapter introduction

Over the last decade social media, particularly Twitter, has had a huge impact on political campaigning. Politicians are no longer confined to official, top-down campaigns run by their parties. Social media frees them to speak directly to voters and campaign for themselves, sometimes with mixed results. A successful social media strategy is frequently credited as a key part of Barak Obama's victory in 2008. However, for every Barak Obama there is also an Emily Thornberry, sacked from the Labour Party Shadow Cabinet for an injudicious tweet, an Anthony Weiner, the US Congressman forced to resign after accidentally tweeting a sexually explicit picture of himself, or a Jack Dromey, an MP with a number of Twitter gaffes to his name including publicly favouriting a tweet linking to a porn site.

Twitter is ubiquitous in politics now. At the time of writing, 563 of the 650 British MPs were on Twitter along with all American senators and 430 out of 435 members of the US House of Representatives. However, having a presence on Twitter and using it effectively to engage with voters are two different things. There is a growing body of research investigating how politicians use Twitter, and an equivalent body examining how voters use it. However, there is little literature considering how the two overlap, how voters respond to politicians' tweets, or what works and what does not when it comes to politicians stimulating voter engagement on Twitter. The research presented here addresses that gap. Specifically, it focuses on retweets as a measure of the extent to which politicians effectively engage with their Twitter followers, and uses CHAID decision trees to build a series of predictive models that determine how likely particular tweets are to be retweeted and identify the factors that most influence whether they are retweeted. CHAID modelling has not previously been used in this way and so this research makes a methodological contribution as well as addressing a gap in the political marketing literature.

This chapter puts forward the rationale for the research, along with an overview of the topic, before discussing the overarching research question to be addressed. The methods and findings of the research are briefly outlined and the research contributions pinpointed. The chapter concludes by outlining the structure of the thesis and summarising the key points of its argument.

1.2.   Rationale for the topic

This thesis examines how British politicians used Twitter during the 2015 General Election. Specifically, it identifies the factors which influenced whether or not politicians' tweets were retweeted. Retweets have been chosen as the variable upon which to focus because they are a simple measure of the extent to which a tweet has stimulated engagement amongst its recipients.

The research is based on analysis of all original tweets sent by incumbent MPs during the official campaign period of the 2015 General Election, from the Dissolution of Parliament on 30 March to the election itself on 7 May[1]. A combination of descriptive statistical analysis along with computerised and manual content and sentiment analysis is used to identify the factors that best predict whether tweets are retweeted or not. These findings form the basis of a set of practical recommendations for politicians who wish to make more effective use of Twitter as a campaigning tool.

The rationale for this research is threefold. Firstly, it has a methodological rationale, based on developing a new approach to predicting retweets. Secondly, it has a political rationale, based on addressing a gap in our current understanding of Twitter as a political communication tool. Finally, it has a marketing rationale, based on the wider relevance of this research to our understanding of how social media can be used as a marketing tool. Each is now discussed in more detail.

---

[1] Strictly speaking MPs resign once parliament is dissolved and the election campaign starts and so for the period of the campaign they are not MPs but are candidates like any other. However, as this research only considers the tweets of those politicians who were MPs and not of any other candidates, the term 'MP' is used throughout.

1.2.1. Methodological rationale

This research contributes to the growing body of research aimed at predicting either whether tweets will get retweeted or how many retweets they will generate (e.g. Suh *et al*., 2010; Yang *et al.,* 2010; Hong *et al.,* 2011; Kupavskii *et al.,* 2012). In general, extant literature treats this question as an abstract computing problem rather than as a marketing problem so focuses on generating accurate predictions against a random sample of tweets rather than on identifying the factors that influence the chances of retweeting or on discussing the practical implications of findings for tweeters. This thesis moves retweeting research forward by applying a modelling technique that has not previously been used in this way, as well as by addressing a specific tweeting context rather than a random sample, and by focusing not only on whether tweets get retweeted but also on explaining why.

Numerous different approaches are taken to predicting retweets in the extant literature. However, as chapter four discusses, analytical approaches like logistic regression or neural networks can produce a score estimating the likelihood of a tweet being retweeted fairly accurately, but they do not make the factors that influence the retweet transparent, so offer limited practical benefit to anyone wanting to understand how the different elements of their tweets influence the chances of them being retweeted.

This research uses CHAID decision tree algorithms to predict which tweets will get retweeted and to identify the factors that most influence retweeting. CHAID is commonly used in predictive analytics, but does not appear to have yet been used to predict retweets. CHAID offers many benefits over other prediction methods (discussed in more depth in chapter four), but one particular advantage is that its output is easy to understand, particularly when compared to the output of alternatives such as logistic regression or neural networks. If the aim of the research is purely to generate an accurate score that predicts the chances of a tweet being retweeted (as tends to be the case in extant literature), then one does not necessarily need to understand how that score is derived. However, if the aim of the research is to shed light on the factors that influence retweeting so that people can use that

knowledge to craft more effective tweets, then having easily understandable output matters greatly.

Extant research predicting retweets almost exclusively uses as predictors the structural variables which are easily accessible from the Twitter API (application programming interface), rather than variables relating to the content of the tweet. Common sense suggests that the content of the tweet probably influences whether it gets retweeted. However, when content analysis is used in extant retweeting research, it is purely based on machine-based methods such as automated counts of the number of times that particular words appear in a tweet (e.g. Petrovic *et al.,* 2011; Yang *et al.,* 2010). Likewise, if sentiment analysis is used it is based on an automated sentiment score (e.g. Kupavskii *et al.,* 2012) rather than one derived from manual sentiment analysis. However, computer-based sentiment analysis is not as accurate as manual sentiment analysis (Canhoto and Padmanabhanb, 2015), and computerised content analysis techniques are not yet sophisticated enough to be able to identify accurately either the topic of a particular tweet or the purpose which the author had in mind when writing it (Conway, 2006; Krippendorff, 2013; Lewis *et al.,* 2013). The lack of manually-derived content and sentiment variables in existing retweet prediction models is a significant limitation, and one which the research presented in this thesis addresses. A secondary contribution is that this research provides a comparison of manual and machine-based content and sentiment analysis as applied to political tweets.

An additional methodological limitation of existing retweet prediction research is it tends to be based on random samples of all tweets rather than on samples of particular kinds of tweets. Using random samples from across the whole Twitter firehose[2] means that no understanding can be gained of which factors are important in determining whether certain kinds of tweets get retweeted. The factors that

---

[2] The Twitter firehose and the Twitter Streaming API are the two ways researchers can access Twitter data. The Twitter Streaming API is free but only enables access to a sample of tweets matching the searcher's criteria. The Twitter firehose guarantees that a search will return every tweet that matches the criteria set however it is prohibitively expensive to access and Twitter restricts access to a handful of companies.

influence whether a politician's tweet gets retweeted may differ from those that influence whether a brand tweet gets retweeted, or a tweet from a personal friend. Thus it is worthwhile to extend the literature on predicting retweets by moving on from methods using general samples to those which focus on specific groups of tweets, as the research presented here does.

In summary, this research extends existing methodological approaches to predicting retweets in three ways:

1. Using CHAID decision tree algorithms to identify the factors which most influence the chances of a tweet being retweeted in a transparent and intuitive way. CHAID does not appear to have been used for this purpose before.

2. Using variables generated via manual content and sentiment analysis of tweets as predictors in the models to further develop understanding of how content and sentiment of tweets affect retweeting. Extant literature tends either to not use content-related variables or to use computer-generated sentiment and content variables only.

3. Focusing on identifying the factors which influence the chances of particular types of tweets – those sent by politicians during an election campaign – of being retweeted rather than simply using a large sample of random tweets.

1.2.2. Political rationale

Social media now plays an important role in politics, and Twitter particularly offers new ways in which politicians can relate to the public and vice versa (Newman, 2010). By 2010 Twitter had already established itself as a core communication tool for politicians and an essential source of real-time information for journalists, shown by the fact over 600 candidates were active on Twitter during the 2010 general election campaign, along with hundreds of journalists, party workers and advisors (Newman, 2010). By the 2015 election the number of candidates active on Twitter had grown to over 2,000 and now the vast majority of UK MPs and virtually all American members

of Congress are on Twitter. Jeremy Corbyn (leader of the Labour Party at the time of writing) argues that his party is now forced to use social media as its primary method of communication with voters, as it can no longer get a voice through traditional media (Mason, 2016). During the political turmoil that followed the EU Referendum in June 2016, Twitter was the dominant mode by which MPs made announcements and disseminated their views. For example, the majority of the members of the Labour Party shadow cabinet who resigned in the days following the vote announced their resignations on Twitter first (*Labour Shadow Cabinet and ministers' resignations - the letters in full*, 2016).

Twitter is a powerful communication tool which has much to offer MPs, but it can also be dangerous as many politicians have found to their cost. Twitter can shine a light on the political process and provide the public with a hitherto impossible way to engage with politicians personally, but its use can also interfere with the democratic process and hinder clear communication (as will be discussed further in chapter two). Thus, a better understanding of how politicians use Twitter is of benefit both to politicians themselves, who wish to make more effective use of Twitter personally, and to those in wider society concerned about the growing role of social media as a political communication tool.

There is a growing body of research examining how politicians use Twitter and, in parallel, literature which considers how citizens use Twitter for political communication. However, to date there is little literature that considers how the two overlap. It is commonly reported, for example, that politicians tend to use Twitter as a one way broadcast mechanism rather than for truly engaging with citizens (e.g. Hemphill *et al.,* 2013; Momoc, 2012; Larsson and Hallvard, 2011 amongst others). However, given that many politicians' tweets are retweeted or replied to, some engagement clearly does take place, but little is known about the factors that stimulate this. There is also a third body of literature focusing on understanding and predicting retweets in general terms, but nothing that looks specifically at retweets of politicians' tweets. Figure 1 shows how the research presented in this thesis

contributes by bringing together areas of literature which have, to date, been considered separately.

Figure 1 - Gap in extant literature which this research addresses
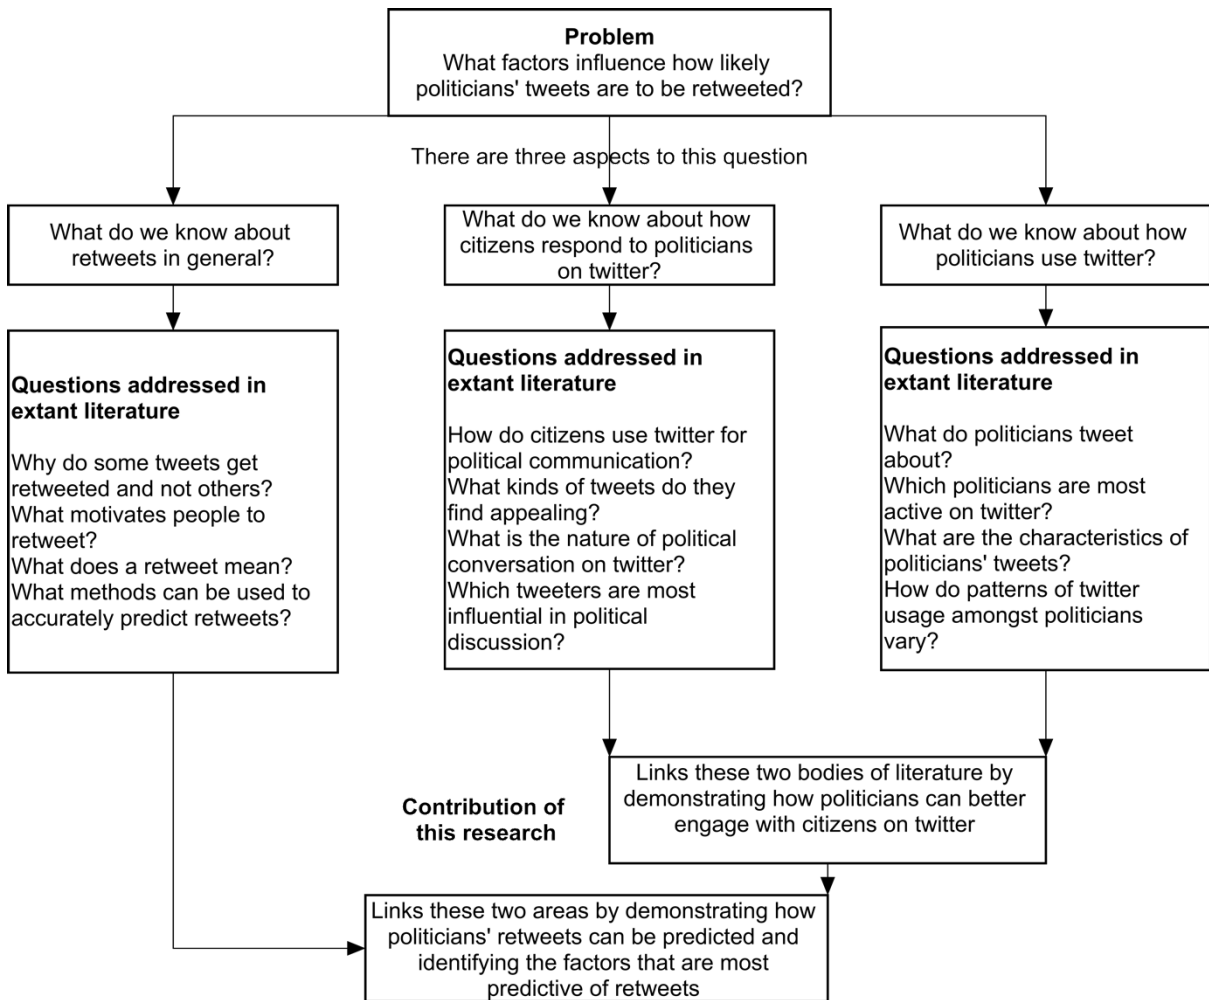
The power of Twitter to break political careers is clear when one considers examples such as Emily Thornberry[3], however politicians are still largely left to their own devices by their parties when it comes to deciding how to use Twitter. The potential of Twitter to enable better engagement with citizens is clearly recognised by politicians such as John Prescott and Stella Creasy, who each have thousands of Twitter followers and

---

[3] The Labour MP was forced to resign from the shadow cabinet in November 2014. On the day of the Rochester and Strood by-election she tweeted a photograph of a house in the constituency which was displaying several St George's cross flags and had a white van parked outside, with the caption 'images from Rochester'. She was accused of snobbery and disrespect for voters and was criticised by many fellow Labour MPs. By the end of the day she had resigned her shadow cabinet position.

use it to communicate with voters on a daily basis. However, many politicians make limited use of Twitter, have few followers, tweet rarely and tend only to use Twitter in broadcast mode to tell people what they are doing, achieving little of significance this way. Most politicians would benefit from better understanding how Twitter works as a communication tool and how it can be used to effectively engage with voters, particularly during campaign times.

Citizens would also benefit from better understanding the role of Twitter in political life. Social networking sites are now used for political communication on a significant scale. In the US 20% of registered voters say they use social media to encourage others to vote (Cook, 2013) and in 2010 an experiment by Facebook suggested that a single election day message was responsible for encouraging an additional 340,000 people to vote (Bond *et al.*, 2012). This may seem like a small number when considering the scale of US elections (it represents a 0.4% uplift in voter turnout) however such small numbers can be significant. George Bush's win in 2000 hinged on the outcome in a single state – Florida – where the margin of victory was only 537 votes. Thus, even if the impact of social media on voters is small, it still has the potential to have a significant effect on the democratic process.

The use of social media for political marketing is largely unregulated, and this is reflected in the changing ways that politicians have used the internet and social media over time (discussed in more depth in appendix one), moving from treating websites simply as online brochures to using the internet for interaction and as a way of circumventing the rules that govern traditional marketing media (Towner and Dulio, 2012). Parties do not have to declare responsibility for any images which they post on Twitter, leading to a growing number of 'tweetgraphics' containing campaign messages without any direct party branding (such as the example in Figure 2).

Figure 2 - Attack on Ukip tweeted by the official Liberal Democrat account

Such communications are not subject to offline campaigning rules (Mason, 2014), so parties are not obliged to take responsibility for the messages contained within them or to ensure that they are accurate. Hence, tweetgraphics are generally used negatively as attack ads. That said, traditional negative print and TV ads can also be designed to mislead voters and present them with incorrect information (Hall Jamieson, 1992) so this is not a completely new development but it is fair to say that online communication is subject to much less regulation may have more influence than traditional political ads. The 2016 EU referendum campaign was heavily criticised for the use of lies and misinformation, particularly by the 'leave' camp. Katherine Viner, Guardian editor, argues that social media enables such mistruths to spread unchecked in a way that would have been impossible when we relied on newspapers and TV for our news, suggesting that social media has led us into an era of 'post-truth politics' (Viner, 2016). This lack of regulation combined with growing usage and potential for reach makes politicians' use of social media an important area of research.

1.2.3. Marketing rationale

Twitter also plays a growing role in marketing more broadly, in the commercial as well as the political sphere. Large organisations virtually all have a Twitter presence. In turn, social media savvy consumers understand the power of Twitter to force companies to listen to their concerns. Consumers use Twitter as a powerful form of word-of-mouth, to let many people know about their experience of dealing with companies, be that positive or negative. Improving our understanding of the factors that motivate retweeting feeds into more general research examining electronic word-of-mouth (eWOM). Even though the context here is political marketing, the hope is that the findings will be of interest to commercial marketers who want to better understand retweeting. Additionally, although the findings may be specific to the context of political marketing, the method used is not. As explained above, this research demonstrates a new method of identifying the factors that influence retweets and that method is highly transferable between contexts. One of the biggest benefits of CHAID – that it produces highly intuitive output which is easy to operationalise – offers a clear benefit to commercial marketers who want to better understand which of their tweets are getting picked up and retweeted and, crucially, how they can increase that number in the future.

1.3. Research objectives

The primary research question to be addressed in this thesis is: **what factors influence whether politicians' tweets are retweeted?**

In order to answer this question, the research has six objectives.

1. To identify the factors that extant literature suggests will determine whether politicians' tweets are retweeted.
2. To identify other factors which might also play a role in influencing whether politicians' tweets are retweeted.

3. To propose a typology of the tweets sent by UK politicians during the 2015 General Election campaign and identify which most effectively generate retweets.

4. To test the extent to which the factors identified in objectives one and two do indeed determine the chances of a tweet being retweeted by building predictive models using CHAID.

5. To demonstrate a new method of predicting retweets – CHAID analysis – which could be of use to social media researchers and marketers in other fields.

6. To provide practical advice for politicians regarding how to best to use Twitter to engage with citizens as part of a campaign communication strategy.

1.4.   Why focus on Twitter?

Twitter is one of many social media sites. Wikipedia lists over 200 currently active. In the UK the most visited sites are Facebook, Twitter, LinkedIn, Google+, Pinterest and Instagram (Rose, 2014). However, Twitter is by far best suited to this research as it is both highly relevant to political marketing and also easy to research from a practical point of view.

The primary reasons for focusing on Twitter are outlined below:

- Twitter is the network most frequently researched in a political context, so there is a critical mass of existing literature to give context to this research.

- Twitter is inherently outward-facing so it is possible to access all communications that meet particular criteria (in this case all tweets sent by particular people). This is not true of other social networks such as Facebook, which are not primarily designed for public communication but rather for communication between friends. Accessing information posted by users on

Facebook poses both ethical and logistical challenges (Ampofo, Anstead and O'Loughlin, 2012) which are not present if using Twitter.

- Twitter enables messages to be sent quickly and off the cuff. Thus it is the social network most commonly used for political communication and has played a central role in many high profile political events such as the Iranian election protests (2009-10), the Tunisian revolution (2010-11) and the Egyptian revolution of 2011 (Eltantawy and Wiest, 2011; Lotan *et al.*, 2011; Wilson and Dunn, 2011; Hermida, Lewis and Zamith, 2012).

Twitter has become the social network of choice for politicians. On Election Day in 2008 the Obama campaign tweeted once to an audience of around one million followers. In contrast, on Election Day in 2012 more than 300 tweets were sent to over 27 million followers (Helm, 2013). When politicians wish to respond to national events or tragedies, they do it first on Twitter. It is not unusual to see policy announced on Twitter before it is announced to journalists, or even to other politicians (Richards, 2015).

Voters too are increasingly turning to Twitter for political information. In January 2016 44% of American adults reported that they had learned something about the presidential election via social media in the previous week, and just under a quarter said that they looked to the candidates' social media posts for information, making social media a more popular news choice than campaign websites and emails combined (Mitchell, Holcomb and Weisel, 2016). Newspapers and other offline media have normalised the use of social media as source material, thus amplifying the impact of social networking sites, and many politicians use social media as a way of building relationships with the media (Newman, 2010). Political parties are also using social networks as an internal tool for management of and communication with activists, party workers, volunteers and journalists.

Based on the growing use of Twitter for political communication by politicians, citizens and the media, it seems clear that Twitter will continue to play a significant role in politics for the foreseeable future and hence is a worthy area of study.

1.5. Why focus on individual politicians?

The units of analysis in this research are the tweets sent by individual politicians rather than by party central offices or campaign accounts. This is because Twitter is at its heart a personal medium. A politician's Twitter account represents them as an individual rather than their party, and the appeal to followers is that they feel they can interact with that particular person rather than with a faceless organisation. Politicians can circumvent their official party campaigns and talk directly to citizens, offering a sense of transparency and engagement by revealing more about themselves than can be covered in a political interview, as the tweet from Gerry Adams in Figure 3 shows[4].

Figure 3 - Tweet sent by Gerry Adams 27/10/14

This ability to circumvent the official campaign communication strategy and speak one's mind on Twitter also has risks for politicians as well as benefits. As already discussed, there are many well-publicised examples of politicians sending misguided tweets, often with career-limiting effects. Indeed, in the run up to the 2015 General Election, Ukip warned its members to stay off Twitter, as they were committing so

---

[4] More of Gerry Adams' work on Twitter can be seen in his recently released book, My Little Book of Tweets, available from the Sinn Fein bookshop (http://www.sinnfeinbookshop.com/my-little-book-of-tweets-by-gerry-adams/)

many embarrassing gaffes (Boffey, 2014). The freedom that Twitter offers MPs also comes with risks for their parties as it may damage the party's ability to appear coherent and 'on message', but the flip side of this is that if politicians appear more accessible to voters then this may help their parties to build a sense of community and belonging, an important part of modern political branding strategy (Dean, Croft and Pich, 2015). Thus, although the focus of this research is on individual politicians, it should still be of interest to party strategists keen to better understand how Twitter can both help and hinder them.

In commercial marketing companies are generally advised to give their corporate Twitter feeds a personal face by letting people know who is actually sending the tweets. Whilst there does not appear to be any extant research examining whether people are more likely to follow individual politicians than party or campaign Twitter feeds, there is some evidence suggesting that people prefer to get political information from the Twitter feeds of individual journalists rather than from the feeds of media organisations (Lotan *et al.,* 2011) and it seems reasonable to assume that the same may be true of political parties and individual politicians. Anecdotal evidence supports this assumption. At the time of writing the Conservative Party's official Twitter feed had 198,000 followers and Labour's had 311,000. In comparison David Cameron had 1.4 million followers and Jeremy Corbyn 429,000. More people engage with the individual politicians than with the party machines.

Politicians are becoming more aware of the potential value of social media to their campaigns. A briefing document from the European Parliamentary Research Service (Davies, 2014) suggests a number of ways in which politicians might benefit from using social media, in particular drawing attention to the value of network effects – when someone likes or retweets something it can reach more people. The advice for politicians in this document is about how they can effectively use Twitter as part of their individual campaigns, rather than about how Twitter could be used as part of the general party campaign. For these reasons it was felt that focusing this research on how individual politicians use Twitter was likely to yield the most interesting results, fit best into the existing research in this field, and make the biggest contribution.

1.6.   Why focus on election campaigns?

Examining Twitter during election campaigns offers a richer analysis as there is more political activity and conversation on Twitter during these times (Dang-Xuan *et al.*, 2013). Twitter activity seems to be largely determined by offline events, with big spikes in activity around events such as televised debates and of course election day itself (Larsson and Hallvard, 2011). Politicians are more active on Twitter during campaigning periods (Golbeck, Grimes and Rogers, 2010). Clearly, tweets sent during campaigns are not representative of tweets sent at other times, so considering campaign tweets only limits the generalisability of the research. However, the research aims to contribute to understanding of how Twitter can be used as a political marketing tool and so campaign tweets are most relevant to this objective.

1.7.   Why focus on retweets?

The focus here is on retweets as a measure of engagement rather than on replies or favourites. Chapter three will argue that the retweet is a form of electronic word-of-mouth and that the act of retweeting is the most substantive way that someone can engage with a tweet. Factors that influence retweet volume have been identified and are discussed in chapter seven, giving some guidance to politicians who wish to increase the number of retweets that they attract, but the focus of the predictive element of the research is simply on identifying whether tweets are retweeted or not.

1.8.   Research method and design

A literature review was conducted (discussed in chapters two, three and four), from which a conceptual model was developed as show in Figure 4. This shows that there are three factors that influence retweets: the characteristics of the sender of the tweet, of the tweet itself and of the recipient of the tweet. As will be explained in more depth in chapter three, this research focuses on the first two of these only, because they are the only ones over which the politicians sending the tweets have any control, and because information about a tweet's recipients is not easily available.

Figure 4 - Conceptual model showing factors influencing retweets
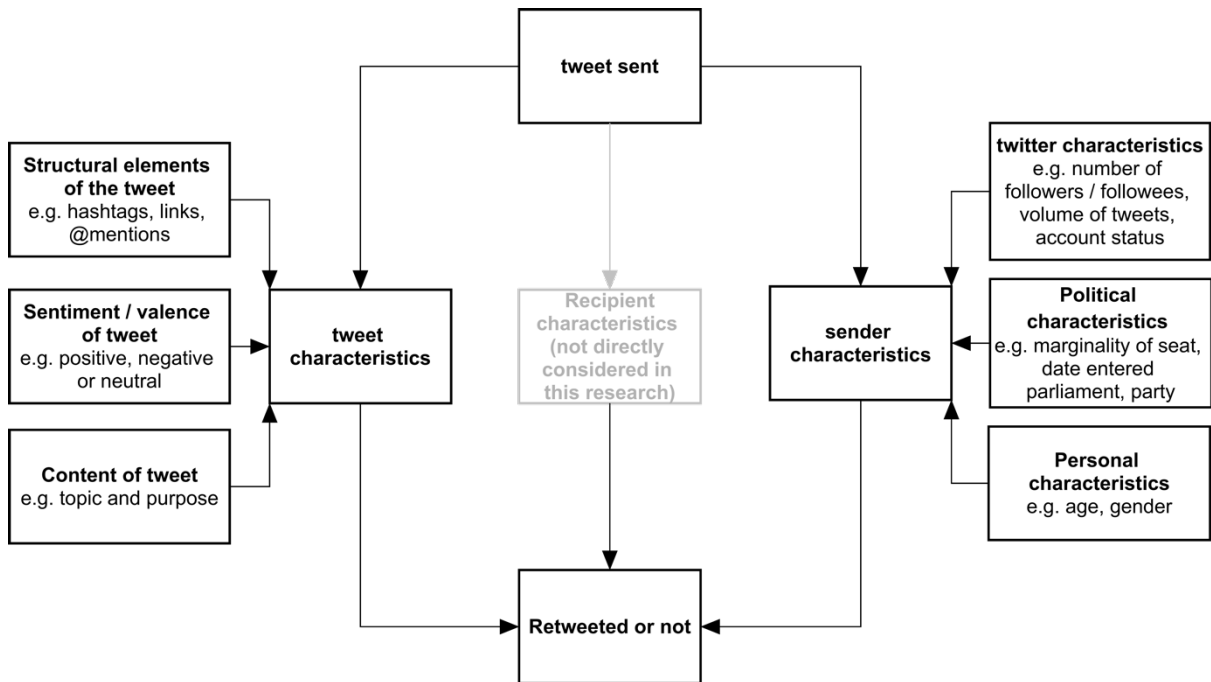
In order to test this conceptual model, all 154,565 tweets sent by sitting MPs during the 2015 UK General Election campaign were collected using Brandwatch, a commercial social media listening tool that enables access to the entire Twitter firehose. All replies to other people or retweets of other people's content were removed, leaving 42,444 original tweets to form the basis of the analysis. The analysis was performed using SPSS Statistics, SPSS Modeler and SPSS Text Analytics (discussed further in chapter five). The stages were as follows:

1. **Descriptive statistical analysis** of all 42,444 original tweets to identify patterns in the MPs' tweeting behaviour, enable comparison of the behaviour of this group of politicians with that of other groups reported in extant literature, and to begin to develop an understanding of the factors that might influence retweeting.

2. **A smaller sample of tweets was created**, matching the 6,510 tweets that were not retweeted with a random sample of the same number of tweets that

were retweeted, thus creating a sample of 13,020 tweets in which exactly half were retweeted and half were not.

3. **A series of predictive CHAID models were built** and run on the matched sample of 13,020 tweets. Different types of variables were used in each model in order to better understand how each type influenced retweeting, as follows:

    a) Models based on variables related to the structural elements of the tweets (e.g. whether they include hashtags, links, mentions of other people or other structural elements).

    b) Models based on variables related to the authors of the tweets (e.g. number of followers, gender, party affiliation).

4. **Machine-based content and sentiment analysis** was conducted, creating new content and sentiment-related variables which were used as the basis of a further set of CHAID models.

5. **Manual content and sentiment analysis** was performed on a smaller sample of 1,212 tweets (evenly split between retweeted and not) and the resulting variables used in a further set of CHAID models.

6. **A 'master model' was built** bringing all the predictive variables together into one model.

The dependent variable used throughout was a yes/no flag indicating whether or not each tweet had been retweeted. The independent variables came from a variety of different sources, as follows:

- Appended by Twitter to each tweet and included in the download from its API (e.g. number of retweets).

- Researched from a third party source and added to the dataset manually (e.g. party affiliation of sender).

- Calculated by the researcher from data provided by Twitter (e.g. ratio of followers to followees).

- Appended to the basic data available from Twitter by Brandwatch (e.g. the Kred[5] score for each tweet author).

- Variables generated by computer-based content and sentiment analysis of tweets.

- Variables generated by manual content and sentiment analysis of tweets.

Broadly speaking, these variables divide into three types:

- Those that relate to the author of the tweet.

- Those that relate to the content of the tweet and its sentiment.

- Those that relate to the structural features of the tweet.

A list of all the variables and explanation of what each means is given in appendix two.

1.9.   Summary of argument

This research shows that CHAID analysis is an effective way both of determining which tweets are most likely to get retweeted, and of identifying which factors most influence retweets. The models built show that the content of the tweet and, more specifically, its sentiment have the most powerful effect on whether the tweet will be retweeted. Negative tweets are both more likely to be retweeted in the first place than either positive or neutral tweets, and also attract a much higher volume of retweets. When tweets are coded by content rather than by sentiment a similar pattern is seen. Those tweets most likely to be retweeted are attacking tweets and, more specifically, tweets which employ fear appeals, retweeted over 80% of the time. This is relevant to the debate surrounding whether negative campaigning turns people off politics (e.g. Ansolabehere *et al.*, 1994) or encourages participation (e.g. Finkel and Geer, 1998; Ridout *et al.*, 2004). This research shows that negative tweets seem to engage people much more than positive tweets. However, we do not know what effect these tweets ultimately have on voting behaviour and cannot assume that

---

[5] Kred (http://home.kred/rules/) specialises in measuring influence online. It gives tweet authors an influence score out of 1,000 based on how likely they are to get retweeted, mentioned, replied to and followed on Twitter and other social networks.

retweets equate to votes. People may respond to the tweets but still feel less motivated to vote.

The finding that negative tweets are much more likely to get retweeted than positive tweets also contributes to an ongoing debate in the literature on online virality by providing clear evidence that, in this context at least, negative content is much more likely to get passed on than positive, in contrast to the findings of some of the main works in the field (e.g. Berger and Milkman, 2012).

As far as the debate about the role of social media in politics is concerned, most of the evidence presented in this thesis provides support for the techno-pessimist perspective, the view that social media simply represents a new medium through which the same political messages can be disseminated by the same actors who have power in the offline word. The number of followers someone has is a powerful determinant of both whether their tweets get retweeted and how many times. The MPs with the highest numbers of followers are those with the highest profiles. There is little evidence to suggest that MPs who do not already have a significant voice in the traditional media are using Twitter to build one. That said, some of the smaller, newer parties do significantly better than larger, more established parties in terms of activating their MPs to tweet and getting a high proportion of their tweets retweeted many times, so there is a glimmer of hope for techno-optimists who see social media as providing an opportunity to get heard for those who struggle to get a voice in the traditional media.

1.10.   Contribution to knowledge

This is an inductive piece of research in which the focus is more on developing theory than it is on testing it. Social media research is still a young field and retweeting research in particular lacks any agreed theoretical underpinnings that could have been used as the basis for this research – many papers are published with no theoretical element at all and those which do include theory come from many different perspectives. This diversity of approaches means that this is an interdisciplinary

project which brings together ideas from academic marketing theory together with ideas from political science, computing and linguistics. The relative newness of the field combined with the length of the academic publishing cycle also means that many of the sources cited are by necessity conference papers, working papers and practitioner-oriented publications combined with niche journals focusing on particular sub-disciplines, as well as more traditional journals leading in particular disciplines. The focus here is on consolidating what can be learned from extant research in relevant areas in order to build a conceptual model which is then be tested using a new method. Thus, this research contributes to knowledge in three different domains – methodology, theory and practice – discussed below.

1.10.1. Methodological contribution

This research demonstrates how CHAID decision tree algorithms can be used to identify the factors that influence retweets, a method which does not appear to have been used in this way before. This method has broad application and could also be used by marketers wishing to understand the factors that influence the chances of any other kind of tweet being retweeted. Additionally, the research contributes to literature on content and sentiment analysis by comparing the effectiveness of machine-based content and sentiment analysis with manual coding and demonstrating that machine-based analysis still has some substantial limitations when it comes to social media data.

1.10.2. Theoretical contribution

This research makes a theoretical contribution in three main areas:

- **Social media use in political marketing** – this research contributes to literature examining how social media is used in political campaigns. In particular, it contributes to our knowledge of the types of tweets that politicians send and the extent to which those tweets effectively engage people, as well as contributing more broadly to the literature on negative campaigning by demonstrating that negative tweets, attacking tweets and fear appeals are more engaging than positive tweets. The research also proposes a new

typology of politicians' campaign tweets. This represents a contribution to political marketing literature, offering a more refined categorisation of politicians' tweets combined with an analysis of the impact of the nature of the tweet on the chances of it being retweeted, something which does not appear to have been done before in the political marketing arena.

- **Influence on Twitter** – studying influence is of value to marketers as understanding how ideas spread and how people's behaviour and attitudes can be influenced is at the heart of the ability to run successful campaigns, both for politicians and commercial marketers. This research identifies the factors that determine which tweets are most likely to be influential in the specific context of election campaigns but the findings could shed some light on the factors that might be influential in other contexts as well.

- **Word-of-mouth marketing** – this research focuses on how political messages spread via social media, but also contributes to a greater understanding of how other kinds of messages might spread, as many of the factors identified as influential are not specific to political messages but are widely present in commercial marketing tweets as well. This has implications not only for politicians but also for brands that want to harness the power of word-of-mouth marketing through social media.

1.10.3. Practical contribution

This research makes a practical contribution by identifying the factors that influence the chances of politicians' tweets being retweeted. Using this knowledge, practical recommendations have been developed which politicians and campaign managers can use in order to make more effective use of Twitter as a campaigning tool. Although these recommendations are specific to a political context, some may also have broader relevance in other contexts.

1.11.   Outputs to date

Walker, L. and Baines, P. (2013) 'A content analysis of Northern Irish political marketing print advertising', in *AM2013, 9-11 July*. Cardiff. Available at: https://marketing.conference-services.net/resources/327/3554/pdf/AM2013_0134_paper.pdf.

Walker, L. and Baines, P. (2013) 'Fear appeal use in Northern Irish print election campaign material', in *7th International Political Marketing Conference*. Stockholm 19-20 September.

Antonetti, P., Baines, P. and Walker, L. (2015) 'From elicitation to consumption: assessing the longitudinal effectiveness of negative emotional appeals in social marketing', Journal of Marketing Management, 31(9-10), pp. 940–969. doi: 10.1080/0267257X.2015.1031266.

Walker, L., Baines, P., Dimitriu, R. and MacDonald, E. (2016) 'Antecedents of retweeting in a political marketing context', submitted to Psychology & Marketing.

1.12.   Structure of thesis

This thesis is structured as follows.

- Chapter one presents an overview of the research questions and the rationale for the research.
- Chapter two presents a review of literature relating to social media and politics, including a discussion of what is already known about how politicians use Twitter.
- Chapter three reviews literature relating to the key areas of theory most relevant to this research: diffusion of innovations and the spread of new ideas, electronic word-of-mouth and online virality. It concludes by presenting the conceptual model built.

- Chapter four presents a discussion of the extant literature on the prediction of retweets with a particular focus on the methods used in order to provide justification for the methods used here.

- Chapter five begins with an overview of the ontological and epistemological assumptions underpinning this research before outlining the research methods adopted and justifying their selection.

- Chapter six presents an analysis and discussion of the data collected during the 2015 British General Election campaign.

- Chapter seven presents findings of the research and shows how the research questions have been addressed.

- Chapter eight summarises the main conclusions of the thesis, suggests how it contributes to knowledge and presents some recommendations for consideration alongside a discussion of the limitations of the research.

1.13.   Chapter conclusion

Twitter is a force of growing importance in political life but our understanding of what works and does not when campaigning on Twitter campaigning remains limited. This introductory chapter has presented the research question examined in this thesis and explained why this question was considered to be worthy of study as well as outlining the ways in which the research presented in this thesis will help further develop politicians' understanding of how they can most effectively use Twitter to engage with citizens. The next stage of the research is to identify factors that the literature suggests might influence whether politicians' tweets get retweeted (research objective one). To this end, chapter two reviews the extant literature on social media in politics, with a particular focus on determining what is already known about how politicians use Twitter.

Chapter 2 How social media is changing politics

2.1. Chapter introduction

There are three ways in which extant literature has informed this research, as shown in Figure 5. Broadly, the literature on how politicians currently use Twitter provides contextual understanding and background for the research presented here (and is the focus of this chapter). Literature concerned with how ideas spread provides theoretical insight to inform the research (discussed in chapter three), and literature that predicts retweets in other contexts informs the methodological decisions made during this project (considered in chapter four).

Figure 5 - How extant literature informs this research
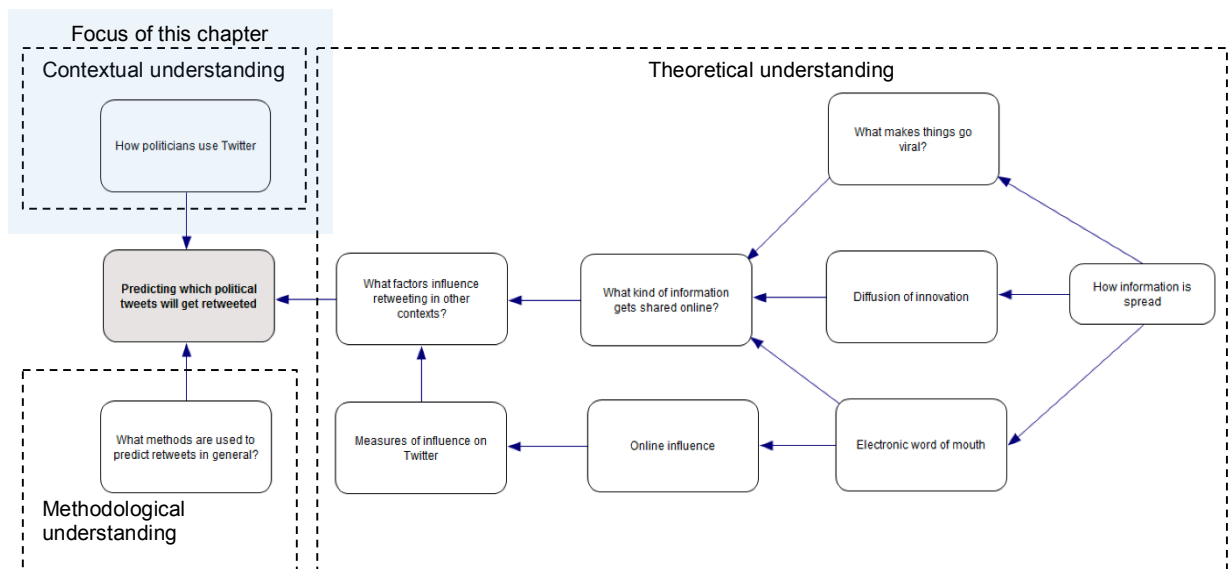
This chapter provides context for the research by showing how social media has fundamentally changed the nature of political campaigns, with a particular focus on how politicians use Twitter. It outlines the key debates relating to how the internet influences the political process before moving into a more focused consideration of the effect that social media has had on different aspects of the political process. The

chapter concludes with a presentation of what is currently known about how politicians use Twitter, feeding into the first research objective of identifying factors that might influence whether politicians' tweets get retweeted. The chapter argues that, whilst Twitter has the power to fundamentally change the nature of political campaigns, relatively little is known about how politicians use it, and virtually nothing is known about how citizens engage with politicians' tweets.

All parties recognise the importance of social media to their communication strategies. All have official Twitter feeds representing 'corporate' party communication and encourage individual candidates and serving politicians to tweet[6]. One could be forgiven for thinking that the widespread acceptance of social media by politicians as a method of communicating with citizens is a sign that social media is a force for good in politics, reducing the distance between politicians and citizens and facilitating greater transparency of the political process. However, the effect of the internet in general and social media in particular on politics is much debated and there is no consensus over whether its effects are positive or negative, as will be further discussed below.

2.2.   The effect of the internet on politics

There is an ongoing and unresolved debate in the literature about whether social media is a democratising force or not. Current literature examining the impact of the internet on politics divides into two schools of thought – techno-optimists and techno-ms (Turnsek and Jankowski, 2008) – a summary of the core arguments of these two groups is provided in Table 1.

---

[6] With the exception of Ukip which, as already mentioned, advised its candidates not to tweet during the 2015 election campaign, after a string of highly publicised social media mistakes by party members (Boffey, 2014).

Table 1 - Comparison of key arguments of techno-optimists and techno-pessimists

Techno-optimist view	Techno-pessimist view
• The internet is a positive democratic force	• The internet does not change the relationship between politicians and the public at all
• Makes political information more accessible to people	• It just facilitates more of the same
• Reduces reliance on journalists and other gatekeepers	• The same groups and people who have the loudest voices offline also have the loudest voices online
• Lowers the cost of political participation	• The internet perpetuates existing political structures
• Provides more opportunities for political participation	• People only engage with material that supports their pre-existing view thus reducing political debate
• Gives a voice to under-represented groups	• People are much less likely to be confronted with views with which they don't agree
• Exposes people to a wider range of political information that they would previously have accessed	

In essence, techno-optimists view the internet as a positive democratic force both for politicians and the electorate because it makes political information more accessible to people and encourages political participation by lowering its cost and increasing opportunities. This is also known as equalisation theory, on the basis that the internet removes some of the barrier which have previously led to some groups being favoured in the political process whilst others struggle to make their voices heard (Morris, 2011[7]).

By giving a voice to underrepresented groups and enabling them to compete on a more level playing field, techno-optimists argue that the internet moves us closer to Thomas Jefferson's vision of direct democracy. According to this view, the internet enables us to communicate with one another directly and make our views known so

---

[7] This reference is to the Kindle edition of Morris's book published in 2011 however the original work was published in 1999.

we are no longer reliant on journalists and other intermediaries to provide us with information and tell us what to think. Indeed, Morris (2011) imagines a future in which decisions can be put to a public vote and politicians are guided almost entirely by public opinion.

By contrast, techno-pessimists see the internet as just another information channel, one that does not fundamentally change the relationship between politicians and the public at all, instead facilitating 'more of the same'. This is also known as the normalisation theory and, according to this view, the internet is merely a new tool used to perpetuate existing political structures (Sunstein, 2001). Techno-pessimists argue that, whilst in theory the web enables people to access a much wider range of political information than they would normally be exposed to, in practice people only engage with material that supports their pre-existing views, resulting in much less political debate. Sunstein (2001) imagines a future where print newspaper readership has declined almost totally with people instead getting their news from a personalised 'Daily Me' of content they have selected for themselves from sources that confirm their biases.

Indeed, as people move more towards 'curating' their own news streams using tools such as RSS feeds and social media to control what they see, it seems like the Daily Me may be becoming a reality and, that being the case, perhaps people are exposed to a much narrower range of information and viewpoints than they would be if they still got most of their information from television or newspapers. This, techno-pessimists argue, is bad for democracy in two ways: firstly, it means that people are much less likely to be confronted with views or materials that they have not chosen in advance, important as a guard against extremism, and secondly, such information filtering leads to fragmentation of views between highly divergent groups who do not understand each other's positions at all.

The so-called 'filter bubble' (Pariser, 2011) whereby social media algorithms ensure that people only see that information which supports their existing worldview, has never been more apparent than during the EU referendum campaign. Many people

who voted remain in the referendum found the leave result particularly baffling because their social media feeds gave the impression that the remain campaign was clearly winning. In the referendum aftermath, internet activist Tom Steinberg addressed this issue on his Facebook page:

> "I am actively searching through Facebook for people celebrating the Brexit leave victory, but the filter bubble is SO strong, and extends SO far into things like Facebook's custom search that I can't find anyone who is happy *despite the fact that over half the country is clearly jubilant today* and despite the fact that I'm *actively* looking to hear what they are saying…We are getting countries where one half just doesn't know anything at all about the other half." (Steinberg, 2016)

If the internet tells you that everyone agrees with you and the campaign is won, perhaps that might mean that you are less likely to vote. Perhaps it might mean that you are less likely to give proper consideration to an issue before you do vote as you are only exposed to one side of it. Some theorists take this techno-pessimist view even further, suggesting that companies such as Google could use the information that they hold on their customers to try and influence the results of elections by altering search results based on individual political preferences and showing only information designed to persuade a person to vote in a particular way (Oboler, Welsh and Cruz, 2012).

The most popular social networks (Facebook, Twitter and most recently Instagram) have all moved to controlling what users see via algorithmic prediction of which posts are most likely to be of interest to them, rather than presenting posts in chronological order. "Algorithms such as the one that powers Facebook's newsfeed are designed to give us more of what they think we want – which means that the version of the world we encounter every day in our own personal stream has been invisibly curated to reinforce our pre-existing beliefs." argues Guardian editor, Katherine Viner (Viner, 2016).

Facebook is now one of the most important distributors of news online and there has recently been some controversy about whether its 'trending news' sidebar has a built-in bias against right wing content, due largely to the fact that it is curated by a group of young, liberal journalists (Nunez, 2016), suggesting that Oboler et al's vision (2012) may not be so far-fetched. Additionally, Facebook recently announced that it would be changing its algorithm to prioritise posts from friends and family rather than news posts, meaning people are even more likely to see only posts that they agree with.

The techno-optimist and techno-pessimist views represent extreme ends of a scale and both arguments have weaknesses. On the one hand, the techno-optimists overlook the fact that many of the most popular websites are those of traditional media outlets and so many people who now get their news from the internet are getting the same information that they used to get offline, just through a different medium. On the other hand, the techno-pessimists ignore the fact that people have always chosen to engage with news that supports their own point of view, through the act of selecting a particular daily newspaper or a preferred news channel on TV. Again, the role of the internet here is to provide further facilitation for something that was already happening rather than to change the game. Additionally, the central assumption of techno-pessimism, that people online are only exposed to information they agree with, may not hold completely true. The 'hijacking' of hashtags from the other side of the political debate happens regularly in politics (Conover *et al.*, 2011; Stieglitz and Dang-Xuan, 2013a; Raynauld and Greenberg, 2014) and is an example of how people can be exposed to views with which they do not agree on Twitter even if they are only following people with whom they do agree.

If the innovation hypothesis put forward by the techno-optimist camp is correct then we would expect to see newer, less established parties making more use of web 2.0[8] tools as they would have the most to gain from them. However, Vergeer *et al.* (2010)

---

[8] The term 'web 2.0' came into popular usage in the mid-2000s and refers to websites which encourage user generated content, interaction, collaboration and social activity rather than the passive consumption of content characteristic of so-called 'web 1.0' sites. The social web is a key part of web 2.0 technologies.

found that candidates from older, more established parties were more likely to be using Twitter. Similarly, Larsson and Hallvard (2011) found that the most active political Twitter users were politicians and journalists or well-established bloggers, suggesting that Twitter was serving more of a normalisation function, offering another outlet to established voices rather than enabling new voices to be heard.

Further evidence that politicians' online behaviour mirrors their offline behaviour comes from Druckman *et al.* (2010) who examined US congressional candidates' websites between 2002 and 2006, arguing that during this time candidates' online behaviour largely mirrored their offline behaviour. They were no more or less likely to go negative online than they were offline, and site designers aimed their sites squarely at the typical undecided voter – the target of most offline political marketing communication – in spite of a wealth of evidence that visitors to candidate websites are much more likely to be engaged supporters of that candidate. This, they argue, supports the normalisation hypothesis because it suggests that online campaigning is 'business as usual' with the internet simply being another communication medium rather than changing the nature of political communication in any meaningful way. Similarly Lawless (2012) analysed the Twitter and Facebook communications of members of Congress during summer 2009 and found that essentially what the social networks offered was "two new ways to send the same old messages" (Lawless, 2012 p209). However, much of this research is very dated. Druckman's research is from a pre-social media time and Lawless was looking at how members of Congress used social media when it was in its infancy. Times have changed. People have a better understanding now of the potential that social media has for subverting traditional methods of political communication.

Political campaigning has become increasing professionalised over the years, with the emphasis more towards party-based campaigns directed from the centre and focusing only on strategically relevant seats, thus taking the focus away from what happens on the doorstep (Lee, 2014). The internet, and particularly web 2.0 technologies such as social media, offer the potential for politicians to connect more closely with voters and campaign on their own behalves, something which would provide further support

for the techno-optimist argument. However, Lee's (2014) analysis of 204 campaigns in the North West of England during the 2010 General Election provided little evidence that this happened in practice. Interviews with campaign managers (Lee, 2014) confirmed that interactivity was very limited, suggesting that, whilst politicians were keen on Twitter, they saw it more as a way of keeping in touch with journalists than of communicating with voters. Again, this provides support for the normalisation hypothesis, with Twitter providing another medium through which existing networks of power can be reinforced rather than enabling new relationships to be established.

A substantial limitation of the bulk of research examining the internet and politics is the tendency of researchers to talk about 'the internet' as if it were a relatively static and unchanging thing (Karpf, 2012). In reality the internet is evolving and changing constantly and the internet of 2005 is not at all the same thing as the internet of 2016. Politicians may have been cautious in the early days of their internet and social media usage and become more sophisticated users over time. Thus research showing how politicians used Twitter in its early days may not tell us much about how they use it now.

2.3. The effect of social media on political campaigning

While the jury may be out on whether or to what extent the internet has changed the fundamental nature of politics, it is undeniable that web 2.0 has changed the nature of political campaigning, largely by freeing politicians from party campaign machines and enabling them to communicate with voters directly, as well as reducing the psychological distance between themselves and their constituents, presenting themselves as normal people with families and interests outside politics (Vergeer, Hermans and Sams, 2010; Jackson and Lilleker, 2011; Adi, Erickson and Lilleker, 2013), although there is some debate in the literature about the extent to which they do this in practice (Zamora Medina and Zurutuza Munoz, 2014). This freedom offers benefits but also comes with risks. The power of Twitter 'outrage' has forced numerous politicians to resign in a way that could not have happened without social media, so it is clear that the growing incorporation of Twitter into politics does not simply

represent business as usual for politicians but exposes them to both new opportunities and new risks.

As with the internet in general, so with social media there is debate regarding how it affects politics and whether its influence is positive or negative. Most commentators agree that social media adds to the personalisation of politics by removing many of the logistical and cost barriers to communication inherent within traditional media, meaning that political parties and journalists have less control over political communication than used to be the case (Newman, 2010; Stieglitz and Dang-Xuan, 2013b). However, this is not necessarily always a positive thing as politicians can also use social media to circumvent the legal restrictions that regulate their use of other media, and social media is not subject to journalistic standards regarding accuracy meaning that it is much easier to disseminate misinformation through social media to the extent that information which is palpably false becomes the de facto 'truth' once it has been repeated enough times on social media. For example, during the 2012 American presidential election campaign Republicans used Twitter to direct supporters to sites which contained misinformation about President Obama (Gainous and Wagner, 2014). The proliferation of misinformation means that the recipients of information have to evaluate the quality of the information themselves much more than is the case when information is filtered through the normal rules of journalism (Turnsek and Jankowski, 2008) so the fact that Twitter enables citizens to access a wider range of information without filter does not necessarily make it any easier for them to discern the 'truth'.

It is also generally agreed that social media moves the emphasis away from the political party towards individual politicians and hence facilitates more personalised campaigning than can readily be achieved through traditional campaigning methods (e.g. Enli and Skogerbø, 2013; Lilleker and Koc-Michalska, 2013). Politicians can communicate directly with citizens and potentially use that communication to build up relationships with voters as well as to gauge their support on key issues and solicit their opinions (Steiglitz and Dang-Xuan, 2013). However, there is little evidence that politicians are actually using social media like this in any meaningful way, even though

they may say that they want to (Grant, Moon and Busby Grant, 2010), and one could also argue that, even if they were to do so, it would not necessarily be a positive democratic development. Voters on social media are not representative of all voters. Politicians could use social media to gauge which issues are of concern to people and identify emerging concerns, but this could lead to a situation in which policy positions are changing constantly in response to short term changes in public opinion. Twitter does not represent all of public opinion – evidence suggests that a small number of users are responsible for virtually all political communication on Twitter (Effing, Hillegersberg and Huibers, 2011) – and so overreliance on Twitter by politicians could simply lead to the creation of a new political elite.

Techno-pessimists would argue that social networks like Twitter function as an 'echo chamber' with Twitter users generally choosing who to follow based on whose tweets most chime with their existing views and indeed extant research does provide some support for this view. Smith, Rainie *et al.,* (2014) show that political conversations on Twitter tend to develop quickly into two highly polarised groups – liberals and conservatives – and that the members of these two groups use different hashtags, reference different websites and rarely communicate with people on the other side, supporting the argument that partisan people on Twitter tend to select information which supports their pre-existing views.

However, just because people may use Twitter in a partisan way does not mean that they are not exposed to alternative view points, and there are several ways that this can happen. For example, Twitter now enables organisations to pay for promoted tweets, which they can show to people based on their interests. Thus political parties could pay for tweets to be shown to floating voters or their opponents' supporters. Additionally, as already discussed, hashtags can be 'hijacked' by opposition supporters as a way of forcing their content into the timelines of their opponents (Conover *et al.*, 2011; Raynauld and Greenberg, 2014). For example, the tweet in Figure 6 shows an anti-EU campaigner using the #remain hashtag.

Figure 6 - Example of opposition hashtag use
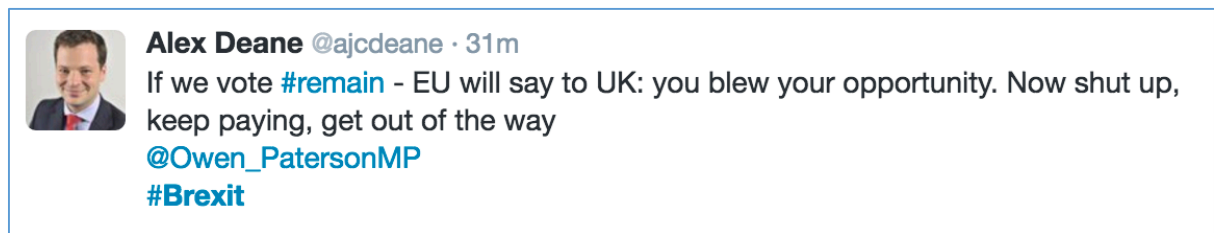
There is a substantial body of research examining the link between offline and online political participation but as yet no agreement on whether social media does or does not influence political participation. Some extant research in this area takes a broad view of online behaviour, for example Bachmann and de Zuniga (2013) show that people who prefer to access news online rather than offline are more likely to participate politically, both on and offline. Likewise Park (2013) shows that people who are (self-reported) opinion leaders on Twitter are more likely to engage in political discussion and to participate politically. There is also some evidence that social media use increases political participation as it enables people to discuss political ideas with each other and disseminate their own views (Steiglitz and Dang-Xuan, 2013), thus developing the online social capital they need in order to be motivated to participate in politics (Gainous and Wagner, 2014). However, there is debate about whether those who are politically active on Twitter are the same people who are politically active in other spheres or whether Twitter genuinely facilitates political participation for new groups of people who were previously politically unengaged.

Just as offline political science literature suggests some groups of voters are highly politically engaged whilst others are not, so the same is true on Twitter. McKelvey *et al.,* (2014) argue that some Twitter users are intensive discussers of political issues and that these users tend to have a sophisticated understanding of Twitter as a medium so make extensive use of Twitter conventions such as hashtags and @mentions, whilst other Twitter users have a poorer understanding of how to use the conventions of Twitter and their political tweets tend to be much simpler and more

basic in both form and content. They find that the online messages of those with a simple tweeting style better track election results than the messages of elite, 'super-users' of Twitter. They suggest that this may be because advanced Twitter users who engage in sophisticated political discussions via the medium are more likely to be political devotees and so their tweets will be less influenced by the day-to-day ups and downs of a campaign than are those of less sophisticated, less engaged tweeters.

McKelvey et al's (2014) research examines citizens who tweet and assumes that there is a relationship between the citizen's level of political engagement and their online political behaviour, essentially supporting the techno-pessimist view that online behaviour mirrors offline rather than extending it. However, it says nothing about the role that politicians' own tweets and other political communications play in stimulating these discussions. It may also be the case that some politicians are sophisticated Twitter users, whilst others have only a basic grasp of Twitter conventions. Just as it is possible to segment citizens according to their relative sophistication on Twitter, so it may be possible to do the same with politicians. The research discussed in this thesis addresses this gap, categorising political tweets according to a range of factors, and then examining the relative effectiveness of different kinds of politicians' tweets in stimulating citizen engagement.

2.4.   The influence of Twitter on voting

We know that voters are influenced by other people when deciding for whom to vote (Lazarsfeld, Berelson and Gaudet, 1968). We also know that people are influenced by others when it comes to deciding what products to buy and that this influence can be exercised both face-to-face through a network of close contacts (Lazarsfeld and Katz, 1955) and also online via various forms of electronic word-of-mouth representing a network of people with whom one has weak ties (Chevalier and Mayzlin, 2006). Thus, if purchasing decisions can be influenced by both strong and weak ties, both online and offline, then the same could be true for voting decisions. Indeed, in 2010 Facebook undertook an experiment on 61 million of its users to see if their intention to vote could be manipulated by displaying personalised messages on election day,

letting people know which of their Facebook friends had already voted (Bond *et al.*, 2012). Users who saw these personalised messages were 2.08% more likely to click a button indicating that they themselves had voted than were users who received a non-personalised message about how to vote. Comparison with actual voting records showed that the group who saw the social message were also 0.39% more likely to actually vote, representing about 282,000 additional votes (Bond *et al.*, 2012). This is a small effect but one that could still have the potential to influence the outcome of an election. Thus the question of whether social networks in general can influence voting behaviour is an important one to consider.

We know that the more actively people use digital media, the more likely they are to also be politically active (e.g. Bimber and Copeland, 2013; Gainous and Wagner, 2014) however this does not address the 'chicken or egg?' question. Does being active on digital media lead to people being more active politically, or are people who are politically active more likely to also make extensive use of digital media? Extant research tends to assume the relationship between digital participation and political participation is linear and getting stronger over time, with digital media playing a more substantial role in elections as time goes on. This puts the technology in the foreground, suggesting that it is the technology itself which is the thing that drives participation.

Instead, Bimber and Copeland (2013) argue that more attention should be paid to the content of the messages and to examining how political and digital participation vary between election cycles. They examine data from the American Election Studies over five election cycles. Whilst their findings do broadly support the contention that there is a trend towards a stronger relationship between internet use and political participation over time, they also find a substantial amount of variance between elections and show that there are no two elections in which internet use predicts the same political participation acts, meaning that the characteristics of the elections themselves also play an important role in determining levels of participation. For example, perhaps close elections are more likely to stimulate high levels of online political engagement than those in which the result seems to be a foregone

conclusion. Bimber and Copeland (2013) suggest that it is the nature of the political content that people encounter online which drives their behaviour, rather than the technology itself. However, there appears to be little extant research directly addressing the question of what kind of political content most drives people's behaviour. The research presented here addresses this gap in the particular context of tweets.

Researchers have consistently found a bivariate relationship between some measure of a party or candidate's Twitter presence and election results (e.g. McKelvey *et al.,* 2014, Connor *et al.,* 2010, Tumasjan *et al.,* 2010). In the UK, research from consultancy Tweetminister examining the 2010 General Election (Tweetminister, 2011) found a strong correlation between the number of votes cast for a party and the number of mentions that it had on Twitter. In 69% of seats where each main party had a candidate on Twitter the most mentioned candidate won. Individual seat predictions based on Twitter were correct 69% of the time. This rose to 87.5% when looking at regional party performance and 90.5% when looking at national share of vote predictions. This suggests that political discussion on Twitter does fairly closely mirror political discussion in real life, but cannot be taken as evidence that Twitter is influencing voting behaviour rather than simply reflecting it. However, Gayo-Avello (2011) argues that the biases inherent within Twitter data and the diversity of methods used by researchers in this field means that solid evidence of Twitter's effectiveness at predicting election results is thin.

Whilst Tweetminister (2010) and Tumasjan *et al.,* (2010) base their research on the total volume of tweets on a particular topic, regardless of who sent them, Vergeer *et al.,* (2010) take a different approach and focus just on politicians' tweets, examining whether there is a correlation between politicians' volume of tweeting and their final share of the vote. Their research finds a positive correlation between how often candidates tweeted and the number of votes they received. Moreover, candidates who increased their tweeting volume as election day approached got more votes, as did candidates who sent higher numbers of tweets directed to specific people, suggesting that perhaps interactivity on Twitter might pay dividends. Effing *et al.*

(2011) calculated a social media indicator (SMI) score for each politician, taking into account the volume of their activity across a range of social media (not only Twitter) and the extent to which they interacted with their followers. In just over half of cases there was a positive correlation between the politicians' SMI score and their eventual share of the vote, suggesting that if there is a relationship between tweet volume and share of vote then it is not a strong one. This is likely to be because of the number of other variables that influence the extent to which politicians tweet, such as the perceived closeness of the election, whether one is a challenger or an incumbent, the ethos of the party one represents and so on.

All this suggests that there is a link between Twitter use and political popularity but it does not explain in which direction the relationship works. Are candidates popular because they are on Twitter or are they on Twitter because they are popular? Jackson and Lilleker's (2011) research suggests that it is the latter, showing that the MPs with the most Twitter followers are those who are best known, thus providing further support for the normalisation hypothesis that Twitter gives an additional channel to those who are already influential. However, a look at MPs on Twitter presents a slightly more nuanced picture. Table 2 shows the ten most followed British MPs on Twitter[9] at the time of writing. Most of the names in this list are clearly big hitters offline as well, being the leaders of the main parties and prominent frontbench MPs. However, there are a few people on this list – Chukka Umunna and perhaps also Caroline Lucas – whose Twitter popularity is beyond their 'real world' political influence, suggesting that perhaps canny politicians can use Twitter as a way of generating influence beyond that which their position would suggest.

---

[9] Data extracted from http://www.mpsontwitter.co.uk/ on 11 May 2016

Table 2 - The ten most followed MPs in British politics

MP	Followers
David Cameron	1,443,971
Ed Miliband	519,690
Jeremy Corbyn	454,904
Nick Clegg	269,474
Tom Watson	215,612
Alex Salmond	204,364
Sadiq Khan	181,410
George Osborne	171,091
Caroline Lucas	170,923
Chukka Umunna	146,118

People may process Twitter communications in a different way than they process more traditional media communications, and this may influence their attitudes to politicians or stated willingness to act after engaging with them via Twitter. For example, Lee and Shin (2014) find that people exposed to a particular (male) politician's Twitter page felt a greater sense of having had a direct conversation with him than those who read a newspaper interview with the same politician. Additionally, people exposed to the Twitter feed reported more favourable impressions of the politician and a greater intention to vote for him but, crucially, only if they were already predisposed towards him. In contrast, those who were only exposed to the interview reported greater awareness of the politicians' position on policy issues and made much less use of source-centred message processing.

Thus it can be seen that political use of social media is growing and that there are a number of ways in which social media communication is changing aspects of the political process and the nature of the relationships between politicians and citizens.

The next section of this review delves into the literature on politics and Twitter specifically in more depth.

2.5.   Mapping extant political Twitter research

Research into political Twitter use generally falls into one of two broad camps (as shown in Figure 7) according to the sampling strategies used. One group of researchers select all tweets sent by particular people (for example candidates in an election, or MPs) irrespective of the topic (e.g. Hemphill *et al.,* 2013; Golbeck *et al.,* 2010; Jackson and Lilleker, 2011), and the other group select all tweets on a particular subject (generally by focusing tweets which include a particular hashtag), irrespective of who sent them (e.g. Burgess and Bruns, 2012; Small, 2012; Larsson and Hallvard, 2011).

Figure 7 - Categorisation of extant literature on political use of Twitter
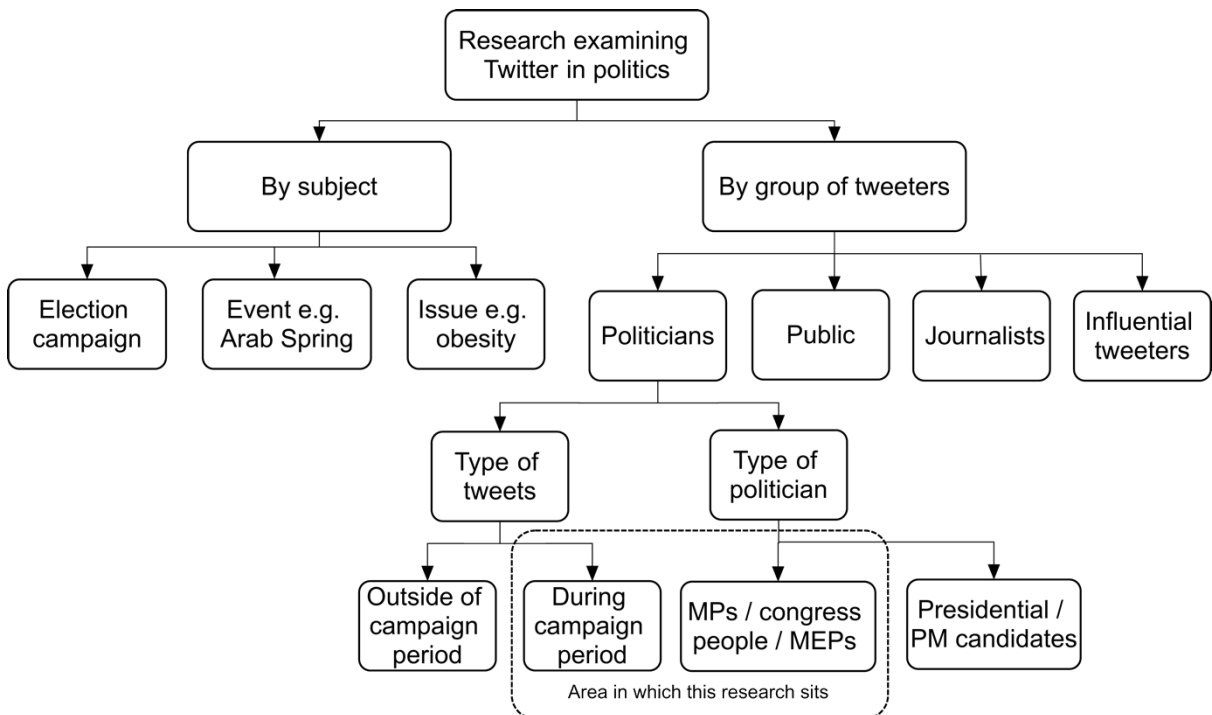
People-focused research generally focuses on one of four groups: politicians, voters, journalists or influential tweeters. Who counts as an influential tweeter is defined in

different ways by different researchers, and this group can contain a mix of all the other types of individuals.

Within the second category of subject-based research there are three broad groups:

- Conversations about election campaigns
- Conversations about political events e.g. the Arab spring or Tunisian revolution
- Conversations about particular political issues

The subset of extant research most relevant to this thesis is that which takes the tweets of politicians as its focus. Again, this research can be further subdivided, depending on whether the focus is on the type of tweets (campaigning or non-campaigning) or on the type of politician (presidential candidates and others campaigning at a national level or MPs and other local representatives).

The research presented here focuses on the tweets sent by politicians during a campaign period and so the bulk of the extant research reviewed here is from the same category. However, where papers from other categories help to illuminate some aspect of political activity on Twitter that is relevant to this research, they have been included in the review as well. A summary of key papers is given in appendix four.

All extant research reviewed focuses on how the politicians use Twitter – there is virtually no consideration of how the public responds to their tweets. The fact that Twitter use amongst politicians is almost ubiquitous suggests they must believe that tweeting has some value, and we can assume that their overall aims are broadly similar to their aims when engaging in other forms of campaigning. However, to date there has been little examination of how well politicians' tweets engage people. There is little point in tweeting if the tweets do not have any impact on those who see them, a gap that Bode and Dalrymple (2014) acknowledge, saying:

> "We cannot fully comprehend how [politicians] use Twitter if we have no information as to who is interacting with them in this medium, *as well as how and why they are doing so.*" (p6, italics added).

The research presented here identifies which tweets successfully engage people and determines the characteristics that those tweets share, thus shedding some light on the gap identified by Bode and Dalrymple (2014): the question of how and why people interact with politicians' tweets. To this end, a retweet is assumed to be indicative of some level of engagement with a tweet (as will be further discussed in the following chapter). This is consistent with the approach taken by much extant research in which retweeting is generally assumed to indicate some form of engagement (e.g. Lotan *et al.,* 2011; Ahn and Park, 2015).

2.6.   How politicians use social media

Interviews with politicians and their campaign strategists suggest that they see social media as having three main purposes: personal marketing through revealing something of their private selves; a way of mobilising followers to take some kind of action; and as an opportunity to enter into a dialogue with voters (Enli and Skogerbø, 2013; Ross and Burger, 2014) and there is no shortage of advice from campaign strategists recommending how politicians should best use Twitter (e.g. Agranoff and Tabin, 2011; Davies, 2014; Cook, 2013). However, there may be substantial differences between how politicians *say* they use social media and how they *actually* use it and so it is important to consider the range of research which examines what politicians actually do on Twitter as well as what they say they do.

Of particular interest is the fact that politicians are advised that effective use of Twitter means securing a wider audience for their messages and capitalising on the 'multiplier effect' through the medium of retweets (Davies, 2014). However, to date, there appears to be little research examining the extent to which politicians' tweets are actually picked up and retweeted and strategists' documents are very light on advice regarding how best to achieve this. In order for politicians to really be able to benefit from Twitter's multiplier effect they need to understand which tweets get retweeted and why. The research presented in this thesis addresses this gap.

2.6.1. How many politicians tweet?

In the UK over 85% of MPs are on Twitter (at the time of writing) and the equivalent figure for members of the US House of Representatives is almost 100%. Analysis of Twitter activity patterns, for both politicians and non-politicians, tends to show a small number of highly active people and a long tail of people who are barely active at all. During the 2010 UK General Election 18% of the tweeting MPs were responsible for two thirds of all the tweets sent (Graham *et al.,* 2013). These findings are repeated across the literature and in different political contexts (e.g. Adi *et al.,* 2013; Grant *et al.,* 2010) so we would expect the range of Twitter activity within the group of MPs who tweeted during the 2015 General Election campaign to be very broad and also to be skewed with a few 'super-users' and a long tail of people who barely tweeted at all.

2.6.2. Effectiveness at attracting followers

If politicians wish to achieve anything by being on Twitter they need to build an audience of followers otherwise they will be tweeting into a void. However, this aspect of political behaviour on Twitter – who follows politicians, how effective politicians are at gaining followers and what those followers do in response to politicians' tweets – appears to be little considered in extant research. Whilst there is a growing volume of research looking at how politicians behave on Twitter (what they tweet about, how often they tweet and so on), there appears to be relatively little which looks at the extent to which members of the public engage with politicians' content.

In one of the few pieces of research considering how politicians and the public interact on Twitter, Nielsen and Vaccari (2013) find that although most people in the US are online, as are most politicians, there is relatively little connection between the two groups. A few popular politicians have high numbers of Twitter followers but the rest have very few – 95% of the candidates examined by Nielsen and Vaccari (2013) had fewer than 3,500 followers. Neilsen and Vaccari present this as a failure of the candidates to engage successfully with supporters via Twitter however, when compared to the mean number of followers across all Twitter users, we can see that having 3,500 followers would actually put someone in the top 1% of Twitter users

(Cha *et al.*, 2012) by follower numbers, suggesting that perhaps politicians do better than average when it comes to attracting followers.

As well as considering the size of politicians' follower networks, it is also useful to consider how these networks are structured. To this end, Vergeer *et al.* (2010) examine the extent to which there is overlap between individual politicians' Twitter networks and find that their networks are relatively homophilous and disconnected from one another. Very few people, it seems, follow more than one candidate, providing further evidence in support of the techno-pessimist view that people use Twitter as a way of engaging with those politicians with whom they agree. Of course, one is not only exposed to the tweets of the people one follows, but also to the tweets that they retweet, so an individual's exposure to a range of different perspectives may be wider than a simple analysis of the network of people they follow would suggest. Further study of retweets would develop understanding in this area.

In addition to looking at follower numbers, it is also interesting to look at the ratio of followers to 'followees' as this statistic reveals something about not only the size of the politician's network but also how interactive that politician is (e.g. Jackson and Lilleker, 2011; Vergeer *et al.,* 2011). Jackson and Lilleker (2011) find that MPs' follower / following ratios ranged from one (following the same number as is followed by) to 745 (745 followers for every one person following). They suggest that a ratio of 10 or under indicates that an MP is in principle willing to listen to the views of others whereas 50 or over suggests that they view Twitter as a one-way megaphone. The more people who follow a candidate, the less likely the candidate is to follow back (Vergeer *et al.,* 2010). This is in line with wider research showing that mass media tweeters – those with 100,000 or more followers – rarely reciprocate follows (Cha *et al.,* 2012). Politicians now have many more followers than was the case when Jackson and Lilleker (2011) did their research, with many having hundreds of thousands or even millions of followers. If one has many followers it is not possible to follow them all without one's Twitter timeline becoming cluttered with the random musings of thousands of people that one does not personally know.

2.6.3. Tweets as broadcast messages or personal interaction?

In view of the debate between techno-optimists and techno-pessimists, it is useful to consider the extent to which politicians' online behaviour mirrors their offline behaviour (the techno-pessimist perspective) or whether indeed social media facilitates a new kind of communication which differs substantively from offline communication (the techno-optimist view). Here again the extent to which politicians use Twitter in 'broadcast mode' is relevant as this suggests that politicians are essentially mirroring how offline communications such as leaflets, letters and party political broadcasts work, whereas greater use of the more interactive features of Twitter suggests moving beyond traditional broadcast communication and engaging in some level of two-way communication.

There is no agreement between researchers regarding what constitutes interactive behaviour for politicians on Twitter. Several studies have examined how politicians use Twitter in terms of levels of interactivity as measured through the use of Twitter conventions such as hashtags, including links, @replies and @mentions, and retweets (e.g. Parmelee and Bichard, 2012; Graham *et al.,* 2013; Hemphill *et al.,* 2013; Adams and McCorkindale, 2013). Others have engaged in a more detailed content analysis of tweets and tried to draw conclusions about levels of interactivity based on that (e.g. Sæbø, 2011; Golbeck *et al.,* 2010). Whatever the method used, the consensus is that politicians primarily use Twitter in broadcast mode rather than to interact with people.

Whilst Twitter offers the potential for direct communication with voters, very few politicians actually use it in that way. For example, Golbeck *et al.* (2010) examine over 6,000 tweets from members of the US Congress and find that they primarily use the medium as a way of transmitting news about themselves and to report on their daily activities. Similar patterns were found in a comparative sample of tweets from UK MPs (Golbeck *et al.,* 2010). Hemphill *et al.* (2013) also examine how Congresspeople used Twitter and they too find little evidence of any personal interaction or of Twitter being used to increase the transparency of the political process. Additionally, even when there is interactivity between politicians and their followers, it is hard to

establish whether those followers are actually constituents of the politician in question.

Research examining political conversations on Twitter has found that less than 10% of tweets using particular political hashtags are conversational (Small, 2012). It may be that few people take advantage of Twitter's interactive possibilities and perhaps the bulk of Twitter communication, irrespective of who it is from, is a one-way transmission of information. Indeed, boyd *et al.* (2010) find that only 3% of their random sample of tweets were retweets, and Suh *et al.* (2010) come up with a similar number – 2.19%. However, much of this research was conducted when Twitter use was in its infancy. By 2014 just under 30% of all tweets were retweets (Liu *et al.,* 2014) showing that Twitter behaviour has indeed changed over time and it may be that now politicians are more adept at making use of Twitter's interactive possibilities.

Most extant research on politicians' Twitter use has an American context. However Jackson and Lilleker (2011) analyse the Twitter use of UK MPs during June 2009 and find similar patterns, showing that UK MPs use Twitter for two main purposes: to publicise their achievements and indulge in impression management, and to present themselves as good constituency MPs thus maximising their chances of benefitting from the personal vote. UK MPs tend not to use Twitter for partisan promotion of their parties or for attacks on opponents, suggesting that they view it more as a medium for personal rather than party promotion and that perhaps there is a difference between UK and US politicians in this regard. Indeed, Lawless (2012) finds that, in contrast, American members of Congress very rarely included any personal information in their tweets – only 5% of the tweets she examines contained any mention of personal thoughts, feelings or experiences.

The nature of the political system in which one operates might make a difference as to how one uses Twitter. There may be less need for American presidential candidates in a two horse race to present themselves as real people with hinterland than there is for a candidate in a marginal UK constituency who is trying to maximise their personal vote. That said, both Lawless (2012) and Jackson and Lilleker (2011) conducted their

research outside election periods and it may be that the character of tweets sent by politicians would be different during an election campaign. Perhaps the MPs studied by Jackson and Lilleker (2011) would have made more use of Twitter for partisan promotion and attacks on their opponents during election periods.

With this in mind it is useful to consider Graham *et al.*'s (2013) research examining the content of UK politicians' tweets sent during the 2010 General Election campaign. They find that 80% of all tweets were about the campaign and party affairs with virtually no policy discussion. Overall 68% of tweets demonstrated some form of broadcast behaviour. The authors suggest that this may be because of the particular nature of tweets sent during an election campaign and that a sample of tweets collected outside of campaign time may be more interactive. The authors also suggest that broadcast tweets are not necessarily a bad thing – they may serve a useful purpose in disseminating information to the public and giving politicians some measure of control over what they say about themselves.

Another non-American study is that conducted by Momoc (2012), examining the use of Twitter by candidates in the 2009 Romanian presidential election. Seven of the twelve candidates were active on Twitter in the month before the election, using Twitter to promote their offline activities such as TV appearances or news articles, and to mobilise voters. None included any information about their personal lives, nor did any of them debate issues or policies, or enter into any form of dialogue with voters. In this regard they had more in common with the US presidential candidates than with the UK MPs, providing further support for the idea that Twitter use may vary depending on one's position within the political system in which one operates, and the nature of that system. For example, the American system is candidate-centred – individual politicians are the focus of the campaign, rather than the party. By comparison, in countries using proportional voting systems campaigns tend to be more party-focused and the role of individual candidates is downplayed. These differences may then influence the ways in which candidates in different political systems make use of social networking (Enli and Skogerbø, 2013) – however there do not appear to be any extant comparative studies examining this question.

Another way to approach the question of the extent to which politicians use Twitter to engage with citizens is to examine broader political conversations which are taking place on Twitter and consider whether politicians contribute to them. Research in Canada (Small, 2012) examines the use of one particular hashtag (#cdnpoli) which flags tweets about Canadian politics. A content analysis of all tweets using this hashtag shows that that less than 2% of them are from politicians with the remaining 98% being from the general public and journalists. Small (2012) suggests that this is because the #cdnpoli hashtag focuses on discussing topics and disseminating useful information rather than on status updates, and therefore is not so relevant to politicians who are much more likely to use Twitter as a way of broadcasting what they are doing rather than as a way of engaging in political debate. It is also important to bear in mind that politicians make up a very small percentage of Twitter users and Small's research does not suggest what percentage of the #cdnpoli tweets *should* be from politicians if they were tweeting at the same rate as other Twitter users. It may be that 2% of the total actually represents a reasonably high level of participation from politicians.

There are of course examples of individual politicians who make effective use of Twitter as a form of two-way communication. For example, Ottawa mayoral candidate Clive Doucet used the #AskClive hashtag to engage with voters during his campaign, generating much more exposure for himself than his limited presence in offline channels would suggest and leading to the eventual winner of the race copying this approach to interaction once in office (Raynauld and Greenberg, 2014). Enli and Skogerbø's (2013) research examining the Twitter behaviour of candidates in the 2011 Norwegian local elections finds that only 38% of tweets were in broadcast mode with all the rest including some form of interaction. Whilst the evidence points to the majority of politicians making relatively limited use of Twitter or sticking to broadcast mode, there are exceptions to this.

An additional challenge when trying to draw conclusions about how interactive politicians are on Twitter is that different researchers have defined interactivity in different ways. Some count numbers of retweets, @mentions or @replies as evidence

of interactivity (e.g. Parmelee and Bichard, 2012; Small, 2012; Vergeer *et al.*, 2010) whilst others present subtler content analyses examining whether tweets are asking respondents to take action or asking for feedback as evidence of interactivity (e.g. Golbeck *et al.*, 2010; Hemphill *et al.*, 2013; Graham *et al.*, 2013). Using retweets, @mentions and @replies as measures of interactivity has the advantage that these measures can be compared across different pieces of research, as definitions of retweets, @mentions or @replies do not differ significantly between contexts[10]. That said, when researchers use some form of content analysis to determine interactivity this adds a level of richness that simple counts of retweets and so on cannot provide, however to date there is no agreed coding scheme for determining tweet interactivity. Each set of researchers has developed their own coding scheme, making it hard to compare like with like or to draw general conclusions across contexts.

Another limitation of research examining politicians' use of conventions such as retweeting and @mentions is that while it gives an insight into the ways in which MPs tweet, it does not consider how citizens respond to these tweets, for example in terms of how likely particular tweets are to be retweeted, to generate an @reply or to influence the recipients' attitudes or behaviour in any other way. An analysis of the content of MPs' tweets may shed some light on the ways in which MPs *intended* their tweets to be received but tells us nothing about whether their tweets were *actually* received in that way.

It is also important to bear in mind that this is a field in which things move very quickly indeed. In the UK the two main works looking at MPs' use of Twitter – Graham *et al.* (2013) and Jackson and Lilleker (2011) – are based on tweets from 2010 and 2009 respectively. The 2010 General Election was the first British election in which social media played any role and took place only a few years after Twitter's launch so it may be that all Twitter users, not just politicians, were making relatively unsophisticated

---

[10] Note, there are some differences of approach regarding measuring retweets. This research, in common with the majority of other researchers, counts only those retweets made using Twitter's automated retweet facility. More detail of the other approaches, along with a justification for the approach taken here, can be found in appendix three.

use of the medium then whereas if the same research were to be conducted now it might find that politicians are using Twitter in different ways as their understanding of its potential for citizen engagement develops over time and they become more sophisticated Twitter users.

Research examining the behaviour of Australian politicians on Twitter found that those politicians who were most retweeted were those who used Twitter for two-way conversational communication rather than for one-way broadcast communication (Grant, Moon and Busby Grant, 2010), suggesting that perhaps the nature of politicians' relationships with the people that follow them is as important as the nature of their tweets in determining which tweets get retweeted.

2.6.4. Use of retweets and @mentions

Previous researchers have examined the extent to which politicians make use of retweets and @mentions in their tweets but there is a great diversity amongst the findings, making it hard to draw any general conclusions about patterns of behaviour. Golbeck *et al.* (2010) examine tweets sent by members of the US Congress and find that only 5 out of their sample of 4,626 were retweets. In contrast, Parmelee and Bichard (2012), looking at the tweets of 26 candidates in US congressional races in 2010, find that 12.4% were retweets and 25% contained @mentions. Graham *et al.* (2013), looking at candidates' tweets in the 2010 UK General Election find that 32% of their sample of politicians' tweets contained @mentions and 18% were retweets. Enli and Skogerbø (2013) examine the Twitter behaviour of a sample of Norwegian politicians during the 2011 local elections and find that substantially more of them included an @mention than was the case with either the British or American politicians examined by Graham *et al.* (2013) or Parmelee and Bichard (2012) (44% compared to 32% for the British politicians and 25% for the Americans). 18% of the Norwegian politicians' tweets were retweets.

Any number of factors might account for these differences. Perhaps Norwegian politics is substantially different in nature to British or American politics, perhaps local

elections are different from general elections, or perhaps Twitter use amongst all politicians has changed between 2010 and 2011. Given the diversity of different methodological approaches, sampling techniques, timescales and countries examined, it is virtually impossible to draw any general conclusions about the extent to which politicians make use of retweets and @mentions in their tweeting.

2.6.5. Valence of political tweets

The effect of the valence of traditional political advertising messages is much debated in the literature and so the valence of political tweets seems worthy of study. Although researchers use many different definitions of negative political advertising (Lau *et al.*, 1999) there is general agreement that "a candidate using a positive advertisement tries to communicate 'what's good about me' [whereas] a candidate using a negative advertisement communicates 'what's bad about my opponent'" (Homer and Batra, 1994, p164). Most agree that voters do not like negative political advertising although, as Lau, Sigelman and Rovner (2007) point out, there is little evidence that they much like positive political advertising either.

Research examining the effect of negative political advertising tends to fall into one of two categories. There are those who view it as a bad for democracy or ineffective. The main proponents of this view are Ansolabehere *et al.* (1994). Their demobilisation theory argues that negative advertising makes voters feel disengaged and disenchanted with the political process, hence reducing political participation and damaging democracy. However, many researchers take the opposite view, arguing that negative political ads actually increase voters' knowledge of candidates, encourage participation and are good for democracy (Finkel and Geer, 1998; Wattenberg and Brians, 1999; Ridout *et al.,* 2004). There is no consensus on this issue. Lau, Sigelman and Rovner's 2007 meta-analysis of research into the effects of negative political advertising shows a complete lack of consistency in research findings.

Measuring the extent to which the valence of a political tweet influences its chances of being retweeted would be one way of shining further light onto the question of whether negative political advertising reduces political participation or not. However,

to date there does not appear to be any significant body of extant research examining the valence of political tweets in general terms, let alone considering the role of valence in stimulating retweets. In one of the few studies of politicians' tweets that does consider valence, Momoc (2012) examines the valence of tweets sent by candidates in the 2009 Romanian presidential elections. He predicts that candidates would be more likely to conduct negative campaigns only due to the lack of regulation of online political communication, however in fact only one candidate out of 12 conducted a predominantly negative Twitter campaign. However, this research is based on a small sample of 12 presidential candidates and so its generalisability is limited.

More broadly, there is some research which examines campaign websites to see whether negative campaigning strategies online match or differ from those offline (Druckman *et al.*, 2010; Klotz, 1998). These papers show that candidates are much less likely to go negative on their websites than they are on TV, suggesting that perhaps the internet encourages or facilitates a different kind of political communication. However, things change fast in this field. Klotz's research is almost 20 years old and Druckman et al's research is based on a content analysis of sites between 2002 and 2006. A further weakness of Momoc (2012), Druckman *et al.* (2010) and Klotz (1998) is that none of them move beyond describing what politicians did to tell us anything about how effective these positive or negative communication strategies were in terms of stimulating voter response. As the internet develops, so the way it is used for political communication changes (discussed in more depth in appendix one), and this is particularly the case with social media so further research which considers how social media is used for negative campaigning is needed to address this gap.

There are good reasons why politicians might want to go negative in their campaign communications. Whilst negative political marketing is unpopular with voters (Merritt, 1984) the bulk of evidence suggests that it is effective nonetheless and candidates tend to use negative advertising because they believe that it works (Yoon, Pinkleton and Ko, 2005). In particular, negative adverts seem to have a higher effect on supporters of the candidate who publishes the advert than on supporters of the

advert's target (Faber, Tims and Schmitt, 1993). This would provide a reason for politicians to campaign negatively on Twitter since people who follow them on Twitter are likely to be their existing supporters. There is also some evidence that voters who are more engaged are more heavily influenced by negative political advertising (Faber, Tims and Schmitt, 1993) and this would also provide incentive to go negative on Twitter as much evidence suggests that people tend to follow those that they agree with, as already discussed. However, the jury is still out on this question as there is a competing body of literature showing that negative political advertising works better on low involvement voters (Dermody and Scullion, 2000). The research presented in this thesis may shed some light onto this question by showing whether negative political tweets are more or less likely to be retweeted. Whilst retweeting is clearly not the same as voting, it is still a measure of endorsement, so if negative tweets were more likely to be retweeted that would still tell us something about how voters might respond to negative campaigning tactics. One might argue that a retweet does not always signify agreement with a message – someone might retweet something in order to comment critically on it. That's certainly true in the case of modified retweets when users quote part of a tweet and add their own commentary. However, this research does not consider modified retweets, only 'straight' retweets where someone uses Twitter's retweet functionality without adding any additional commentary of their own. In these cases, it is much safer to assume that a retweet indicates agreement.

Research examining whether negative political advertising influences voters' attitudes or behaviour tends to share a significant limitation, namely reliance on self-reports of behaviour. We know what voters *tell us* about how negative advertising influenced their attitudes or behaviour rather than how it *actually* influenced their attitudes or behaviour. One benefit of social media research such as that being presented here is that it tells us something about people's real behaviour. We can measure the extent to which the valence of a politician's tweet influences the chances of it being retweeted.

2.6.6. Influence of candidates' positions

There is evidence suggesting that candidates' positions in electoral races may influence the extent to which they use Twitter. For example research in America and Scandinavia (Larsson and Kalsnes, 2014) shows that challenger candidates tend to be earlier adopters of Twitter than are incumbents and that they make greater use of it. In party list systems such as Norway's, a candidate's position on their party's list as well as the party's chances of gaining a mandate both influence the extent to which, and the way in which, the candidate will make use of social media (Enli and Skogerbø, 2013). Graham *et al.* (2013), in their examination of the 2010 General Election, also find that challengers are more likely to tweet than incumbents.

Candidates from minor parties with no chance of winning a seat (or indeed incumbents in safe seats with no chance of losing their seat) may have significantly less incentive to tweet than do candidates in marginal constituencies for whom every vote counts. Jackson and Lilleker (2011) suggest that MPs from major parties are more likely to tweet than those from minor parties, but that marginality of a seat has less impact on tweeting behaviour. However, their research only examines the tweets of 51 MPs and was not conducted during an election campaign period, at which point marginality may become more of an influence.

A politician's electoral position may not only influence whether they tweet in the first place but also the kind of tweets that they send. Research examining tweets sent by American politicians during the 2010 US elections (Parmelee and Bichard, 2012) finds that challengers are more likely to retweet or to use @replies or @mentions than are incumbents. In comparison, Graham *et al*.'s (2013) analysis of the Twitter usage of candidates in the 2010 UK General Election finds that the pattern of usage between incumbents and challengers is reversed, with Labour (the incumbents) and Liberal Democrat candidates used Twitter much more interactively. This suggests that it may not be merely the fact of being a challenger or an incumbent that influences Twitter use but that other factors like the closeness of the election or the nature of the parties may play a role. Extant research suggests that there is some kind of relationship between party and retweeting / @mentioning behaviour, but the exact nature of this

relationship is unclear. The research presented here contributes to this discussion by examining patterns of tweeting behaviour of UK MPs according to party, marginality of seat and election outcome.

2.6.7. Networking with peers

There is some evidence suggesting that politicians are under social pressure to connect with each other on Twitter and that some may use the way in which they interact with other politicians as a way of boosting their own profile (Yoon and Park, 2014). Yoon and Park's analysis of the Twitter behaviour of South Korean politicians shows a high degree of network density in the follower / following networks meaning that following between politicians tends to be reciprocated. They also suggest that politicians mention other prominent colleagues in their tweets as a way of increasing their own profile on Twitter as well as providing support to the mentioned politician. They find a clear relationship between the number of followers and tweets that a politician has and the number of mentions they receive from other members of their party – more followers and tweets is associated with more mentions.

Research examining the Twitter behaviour of German politicians also finds that they are highly likely to have reciprocated following relationships with other members of the same party and to mention each other positively in tweets on a regular basis (Plotkowiak and Stanoevska-Slabeva, 2013). There is also evidence suggesting that politicians from particular parties retweet each other as part of concerted efforts to generate more presence on Twitter for their party (Grant, Moon and Busby Grant, 2010), so a politician's retweet of colleagues' tweets cannot be assumed to mean that the retweeter has particularly engaged with the content of the original tweet. Of course this could also be the case with non-politicians – the more visible you are on Twitter, the more people are likely to interact with you.

2.6.8. What kinds of politicians' tweets do people find engaging?

There has been some research examining what kinds of politicians' tweets people find most engaging or persuasive. For example, Parmalee and Bichard (2012) suggest that tweets from politicians with whom people agree (homophilous sources) tend to be

more influential (in terms of stimulating people to take any form of political action) than those from politicians with whom they do not agree (heterophilous sources). One major limitation of this research, however, is that it is based on interviewing people about their behaviour on Twitter retrospectively, asking them which tweets they remember engaging with and why. Asking people to remember what they have previously done is not as effective as measuring what people actually do, and there is evidence that people cannot always accurately recall political actions (Himmelweit, Biberian and Stockdale, 1978).

That said, Parmelee and Bichard's (2012) research provides evidence that people who follow politicians on Twitter *believe* that some aspects of their political behaviour have been influenced by what politicians have tweeted. They find that 61% of people in their sample claimed to have taken an action such as signing a petition or donating to a campaign as a result of a tweet from a politician; 89% said that they had been motivated to look for more information about an issue as a result of a politician's tweet; 61% claimed to have retweeted a politician's tweet; and 35% said that they had replied to a politician's tweet. This suggests that Twitter can be a powerful tool in a politician's communications arsenal in terms of its ability to influence people to take action of some kind. The research presented here contributes to this discussion by considering which tweets are most effective when it comes to stimulating one clearly measureable form of action – the retweet.

2.7.   Chapter conclusion

This chapter has examined extant research considering politicians' use of Twitter. As can be seen, research on political Twitter use to date has tended to focus either on how citizens talk about politics on Twitter or on how politicians use Twitter but there appears to be very little examining the ways in which the two interact on Twitter. In particular little is known about how politicians can most effectively stimulate engagement with voters on Twitter through the medium of the retweet, a gap which this thesis addresses. The review of political tweeting literature in this chapter has suggested some factors that may influence retweeting, namely the size of the MP's

follower and following networks; the ratio of followers to followees; their political position in terms of party, marginality of seat, incumbency and the nature of the campaign they are in; the valence of their tweets and the content. The next chapter moves on from a discussion focused on politicians' use of social media and Twitter to consider what is currently known about why people retweet things more generally, within the context of literature on the ways in which ideas spread, with the aim of further identifying factors that might influence whether politicians' tweets get retweeted.

Chapter 3 Why do people retweet things?

3.1.   Chapter introduction

The aim of this research is to predict which of the MPs' tweets will get retweeted and to identify the factors that most influence the chances of a particular tweet getting retweeted. In the most general terms, any research aiming to predict retweets is, at its core, about how information is spread. Twitter is a social network built around the sharing and passing on of information, and retweets are the core mechanism through which this is done. Figure 8 shows the key areas of extant literature that are relevant to a theoretical understanding of this question.

Figure 8 - How extant literature informs this research
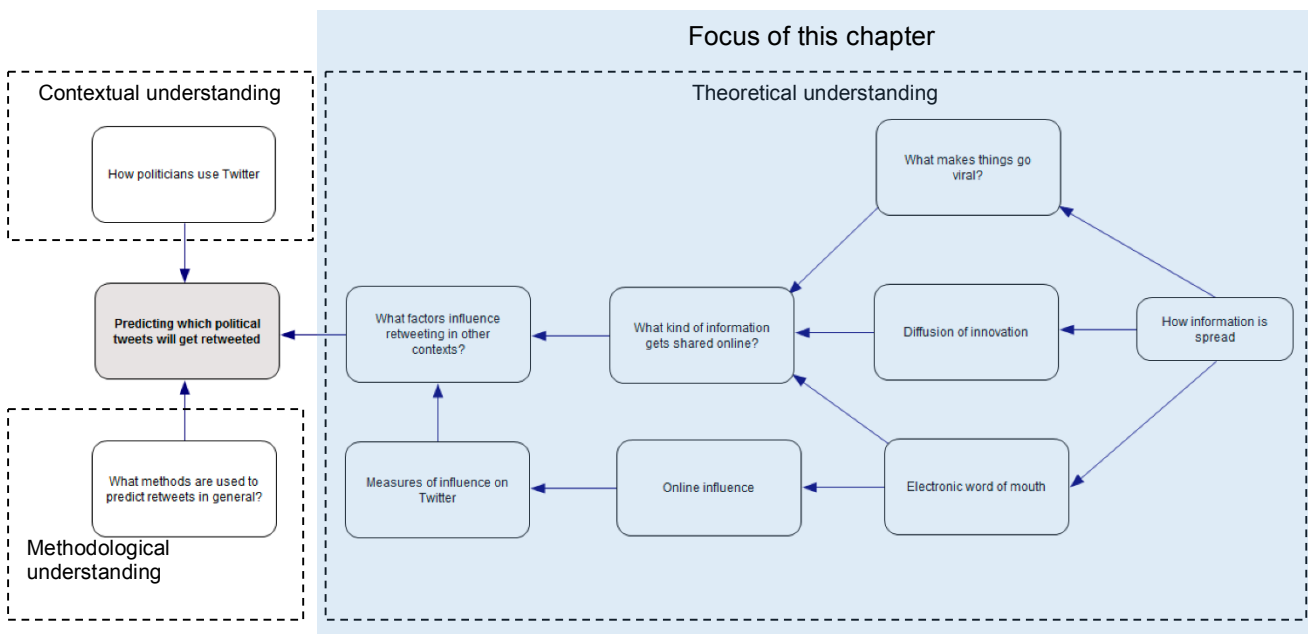
Literature on how information is spread divides into three subgroups that are particularly relevant to this research. Each of the three will be briefly outlined below before being explored in more detail in the chapter that follows. They are as follows:

- **Diffusion of innovation**

  Diffusion of innovation theory focuses on how new products gain acceptance in the market, however, ideas based on it provide a starting point for understanding how people influence one another's thinking more generally and also for understanding why some messages go viral on the internet whilst others do not. This research examines which political messages get picked up and are passed on and why, which is ultimately about how politicians can use social media to build their influence, so an understanding of diffusion of innovation is relevant.

- **Online virality**

  There is a growing body of literature examining how and why particular pieces of content 'go viral' online. There is no commonly accepted definition of what counts as 'going viral' but generally one knows it when one sees it. This thesis is not concerned directly with virality – a politician's tweet that gets picked up and retweeted even a hundred times cannot really be said to have gone viral – however theories of virality shed light on why people share things online and what kinds of materials are most likely to get shared, which helps inform an understanding of what might motivate people to retweet something.

- **Electronic word-of-mouth communication (eWOM)**

  Theories of word-of-mouth (WOM) communication initially considered how and why people share their experiences of products and services with each other through face-to-face communication with people that they know. With the advent of internet marketing and, more recently, social media marketing a branch of word-of-mouth theory has developed which specifically considers electronic word-of-mouth (eWOM) and the ways in which people influence

each other's purchasing decisions through online communication via media such as review sites and social networks. This is relevant because retweeting, replying to or mentioning a politician's tweet is a form of political eWOM.

Information from these three areas feeds into two strands of thought particularly relevant to this thesis – what kind of information generally gets shared online (where the focus is on the characteristics of the information), and who has influence online so is more likely to get their information shared (with the focus on the characteristics of the individual providing the information to be shared). Relevant here too is a small but growing body of literature specifically looking at the factors that influence Twitter retweets. Having outlined the areas of theory most relevant here, each will now be discussed in more detail with the aim for identifying further factors that might influence the retweeting of MPs' tweets.

3.2. Electronic word-of-mouth and online virality

The concepts of eWOM and online virality both relate to aspects of how and why people share information online. However, although the two concepts are closely related, not all eWOM messages go viral and not all viral communications are examples of eWOM. In this section first online virality and then eWOM will be defined and their relevance to this thesis explained, before a further discussion of how they differ and of where retweeting fits between the two.

3.2.1. Online virality

Definitions of virality are thin on the ground, however Goel *et al.* (2015) suggest that "When a piece of online media content…is said to have 'gone viral' it is generally understood not only to have become rapidly popular but also to have attained its popularity through some process of person-to-person contagion, analogous to the spread of a biological virus" (p1), however, despite a growing body of literature examining different aspects of why online content goes viral, from emails (Phelps *et al.*, 2004) to news articles (Berger and Milkman, 2010) to videos (Eckler and Bolls,

2011), there seems to be little or no discussion of how many times something needs to have been shared, or how quickly it needs to be spread, in order for us to be able to say that it has gone viral.

In Twitter one can think of the retweet as a measure of virality (Hansen *et al.*, 2011; Petrovic *et al.*, 2011; Kupavskii *et al.*, 2012; Goel *et al.*, 2015) but clearly not all tweets that get retweeted have gone viral. Most people would agree that a single retweet would not constitute virality, whereas a hundred thousand retweets probably would, however there does not appear to be any real discussion in the literature, let alone a consensus, of where the line should be set. Another related limitation of much of the extant literature on virality is that assumes that something going viral is always to the benefit of the content's originator. This is because most of this literature comes from a marketing perspective where the aim is to try and understand how and why things to go viral in order for brands to be able to capitalise upon this. However, in many contexts when things go viral it is not to the benefit of the originator. This is particularly the case on Twitter when virality often comes about because the tweet's author has spectacularly misjudged the tone or content of their tweet. Jon Ronson's book, *So You've Been Publically Shamed* (Ronson, 2015), presents many cautionary tales of people whose lives have been ruined by a careless tweet going viral. Certainly, this is something that politicians need to be particularly careful of as ill-judged tweets can be career-limiting and, in most cases, when a politicians' tweet goes viral it is not a positive thing.

An additional consideration when examining the extant literature is that researchers have looked at a multiplicity of different media when considering virality, from emails (e.g. Phelps *et al.,* 2004) to video ads (e.g. Eckler and Bolls, 2011) to news articles (e.g. Berger and Milkman, 2010) and indeed to Twitter (e.g. Goel *et al.,* 2015; Hansen *et al.,* 2011). This makes generalisability difficult as the factors that influence virality of email, videos or news articles may differ from those that influence the virality of tweets. For example, it is probably reasonable to assume that if a video ad goes viral that means people are sharing it because they like it. Likewise, if a news article is shared many times then the act of sharing is generally intended as a form of

endorsement of the content. However, when tweets go viral the sharers often intend to mock the original tweet. Politicians then need to understand what factors will influence people to share their tweets in a positive way, whilst avoiding kinds of mistakes that would lead to negative virality. Passing online content on to others with an implied endorsement of the content is a form of word-of-mouth communication and so literature on electronic word-of-mouth is relevant here to help develop a better understanding of why people might share online content.

3.2.2. Electronic word-of-mouth

Marketers have long been aware of the powerful effect that word-of-mouth can have on people's purchasing decisions (Lazarsfeld and Katz, 1955) and the development of the internet has moved word-of-mouth beyond face-to-face communication with someone that you know to a much broader range of communication options via electronic word-of-mouth (eWOM). eWOM can be defined as "any positive or negative statement made by potential, actual, or former customers about a product or company, which is made available to a multitude of people and institutions via the internet" (Hennig-Thurau *et al*., 2004 p39). This is a broad definition which could include customers writing about products on their own blogs, or making statements about products on review sites such as TripAdvisor or Amazon, or talking about products on social media. Tweets about products or companies are a form of eWOM and customers often take to Twitter in order to alert companies to dissatisfaction with service, as per the example in Figure 9.

Figure 9 - Twitter being used as a form of eWOM

Although Hennig-Thurau *et al.*'s (2004) definition of eWOM focuses on commercial marketing it could be extended to include political marketing. If we treat political communication as a form of marketing, as first suggested by Kotler *et al.* (1969), then Twitter interactions with politicians and the things that people say about politics on Twitter can be seen as a form of eWOM too (Jansen and Zhang, 2009). Opinion leaders play a role in influencing how people vote, perhaps more so than the mass media (Lazarsfeld, Berelson and Gaudet, 1968) however, the potential for influence is now much wider than it was in the 1950s, largely because of the power of the internet and social media. There is clear evidence that people's buying decisions are influenced by electronic word-of-mouth (Chevalier and Mayzlin, 2006; Cui *et al.,* 2012; Ye *et al.,* 2011) so perhaps their voting decisions could be as well. Hence an understanding of eWOM in commercial marketing may shed some light on how it could be relevant to politics.

3.2.2.1.   How is eWOM different from traditional WOM?

eWOM is fundamentally different from traditional offline WOM (King, Racherla and Bush, 2014). Offline WOM relies on two people communicating synchronously with one and other, generally face to face (as per Lazarsfeld and Katz's (1955) assumption that opinion leaders must be people whom one has met in person). Therefore the potential reach of a WOM message is limited by the number of people that the message's transmitter can meet and talk to, and that those recipients can then meet and talk to. Thus WOM moves slowly and people are exposed to a limited range of viewpoints. In contrast, eWOM is asynchronous – the message can be passed on without the transmitter and recipient needing to be present in the same place at the same time. This gives eWOM messages a much greater reach than traditional WOM and allows messages to spread beyond the network of people whom an individual knows personally. Weak ties can be highly influential in the world of online eWOM, in a way that would not be possible with traditional WOM. As Jeff Bezos of Amazon writes "If you make customers unhappy in the physical world, they might tell six people, their closest friends. If you make customers unhappy on the internet, they can each tell 6,000 people" (quoted in Agranoff and Tabin, 2011 p26).

eWOM can also take place over a large number of different platforms, from social media and blogs to review sites and discussion forums. This means that not only can messages travel faster and more widely but that, as King *et al.* (2014) argue, it can be hard for brands to track what people are saying about them across all platforms. However, the fact that eWOM offers the possibility of tracking what people say about your brand at all also distinguishes it from traditional WOM. Pre-internet, brands had no way of knowing what people were saying about them to each other. Now any number of social media listening tools[11] exist offering brands (and indeed politicians) the potential to track, measure and monitor all mentions of them on social media and the wider internet.

eWOM is generally a written communication, although the rising popularity of video bloggers, live video streaming services such as Periscope, and unboxing[12] sites suggests that this is changing. Indeed, the DisneyCollector unboxing YouTube channel has more than 2 million subscribers in the United States, making it one of the most popular YouTube channels in the world (Prince, 2014) and a hugely important source of eWOM for Disney. Generating eWOM, whether written or filmed, requires greater effort on the part of consumers than traditional WOM, which could just be a passing comment in a conversation. That said, eWOM's largely written nature means that it has greater longevity than traditional WOM, continuing to influence consumers over time and reaching a much wider range of people.

The fact that eWOM encourages connections between weak ties has a drawback as it allows for the possibility for anonymity and deception, both on the part of companies and consumers. Consumers may give great weight to reviews on Amazon (Cui, Lui and Guo, 2012) when making buying decisions but companies can 'seed' Amazon with fake reviews and other positive comments, a practice known as 'astroturfing' (Bienkov, 2012). Customers can air vexatious or unreasonable complaints via public social media forums, a criticism that many hotel owners and other hospitality businesses have of

---

[11] See, for example www.brandwatch.com, www.hootsuite.com, www.socialmention.com
[12] Unboxing is the act of filming oneself receiving a new product and unpacking it.

the TripAdvisor review site. In the political arena there have been well-publicised cases of both government and party officials attempting to influence online discussions on contentious topics or editing their own Wikipedia entries to present themselves in a more favourable light. For example, in 2007 the Conservative Chairman, Grant Shapps, posed as a Liberal Democrat online in an attempt to discredit his rivals, and in 2012 computers from his constituency office were used to remove mention of this from his Wikipedia entry (Ramesh, 2012), a practice known as 'sock puppeting'. However, whilst astroturfing and sock puppeting do take place, consumers do not necessarily swallow such content whole without any evaluation. They do consider the source of eWOM when deciding how much weight to give it in their decision-making, meaning that, for example, product reviews on small blog sites are less influential than independent review sites (Parmelee and Bichard, 2012).

Whilst companies are increasingly recognising that eWOM cannot be ignored, politicians and political parties still make relatively unsophisticated attempts to engage with the eWOM conversation, if they attempt it at all. For example, in 2014 Ukip's attempts to soften its image via the #whyImVotingUkip campaign on Twitter backfired when the hashtag was hijacked with hundreds of spoof responses (see Figure 10) and David Cameron's attempts to engage via social media were routinely spoofed in similar ways. In 2014 he tweeted a photograph of himself on the phone to Barack Obama and was subjected to a barrage of mockery from celebrities and members of the public as shown in Figure 11.

Figure 10 - Hijacking of the #WhyImVotingUkip Twitter campaign

Elizabeth Day ✔
@elizabday

Follow

#WhyImVotingUkip Because I'm fed-up of foreigners coming over & denying hard-working British people the right to be Nigel Farage's wife.

10:15 AM - 21 May 2014

272 RETWEETS 139 FAVORITES

Vikki Stone
@vikkistone

Follow

#WhyImVotingUkip Because the weather's really starting to pick up, and I don't want it ruined by gays.

10:04 AM - 21 May 2014

2,917 RETWEETS 1,777 FAVORITES

Nathaniel John
@Nathaniel_John

Follow

#WhyImVotingUkip because an Oriental gentleman just pushed into the front of the bus queue and got on first

8:37 AM - 21 May 2014

38 RETWEETS 22 FAVORITES

Figure 11 - The David Cameron Twitter meme

3.2.2.2.   Twitter as a form of eWOM

Twitter is often referred to as a form of 'micro blogging' and thus some researchers underpin their work with literature on blogging in general (e.g. Larsson and Hallvard, 2011). However, it will be argued here that Twitter should actually be treated as a form of eWOM as it has more in common with this mode of communication than it does with traditional long form blogging. King *et al.* (2014) identify the core characteristics of eWOM and, whilst their focus is on product review websites and they do not mention Twitter, it is clear that Twitter matches most of the characteristics of eWOM that they identify, as follows:

- **Asynchronous** – eWOM communication is asynchronous meaning that 'conversations' do not generally take place in real time. Twitter is also largely an asynchronous medium. Although the possibility of real time conversations does exist, most Twitter threads develop asynchronously.
- **Dispersed over many platforms** – compared to traditional WOM, eWOM takes place over many different platforms. Whilst King *et al.* (2014) focus on online review sites there is no reason to exclude Twitter from this list of platforms.

- **Persistent and observable** – eWOM has longevity. A review can remain on a website for many years and nothing is ever truly deleted from the internet. Twitter may be less persistent than other eWOM sources, such as review sites, because the volume of tweets is such that Twitter users cannot generally engage with every tweet that flows through their timeline. It is also possible for Twitter users to delete tweets that they have previously sent. However, tweets which quote the deleted tweet will not be deleted nor will retweets with an added comment, and it is always possible to take a screenshot of a tweet and save it that way. It is also possible to search the Twitter archive for tweets referencing particular topics and indeed there are websites devoted to keeping track of the tweets that politicians delete (e.g. www.politwoops.co.uk), so nothing is ever truly deleted. Thus tweets do have persistence and long-term observability.

- **Enables anonymity and deception** – when King *et al.* (2014) talk about anonymity and deception they are referring to the possibility that companies can exploit the eWOM by posting fake product reviews and 'seeding' discussion forums with positive comments about their products. Along similar lines, company employees can set up fake Twitter accounts in order to tweet positively about brands and products.

- **Offers firms the possibility of community engagement** – as already discussed above, companies are beginning to understand Twitter's potential as a way of communicating with customers and of encouraging the development of stronger relationships between customers and brands. For example, the supermarket chain Lidl used customers' tweets in its #LidlSurprises advertising campaign (Lewis, 2014).

Therefore, tweets can be viewed as a form of eWOM. However, not all brand-related tweets sent by consumers are eWOM. Research in consumer marketing shows that the majority of tweets (over 80%) in which people mention brands are seeking information or asking questions, indicating that much brand-related Twitter use is for the sharing and requesting of general information rather than for spreading eWOM

(Jansen and Zhang, 2009). That said, we cannot assume that the same is true when voters tweet about politicians. Intuitively it seems more likely that if someone mentions a politician in a tweet it will generally be with a view to expressing an opinion, positive or negative, about that politician. Politicians are less likely to be mentioned in the context of information-seeking tweets of the same kind that form the bulk of brands' Twitter mentions (requests for information about opening hours, stock availability and so on).

3.2.3. What does sharing content mean?

The factors that influence how persuasive someone finds a product review may not be the same as those that influence how likely they are to share it. Sharing content is not the same as endorsing it and does not necessarily imply agreement or mean that the sharer found the original content persuasive. For example, someone might share a product review on Amazon because they find it amusing rather than because they endorse its content. There are a number of well-publicised cases of this kind of content going viral, such as the Amazon reviews for Bic black ballpoint pens, a selection of which can be seen in Figure 12.

Figure 12 - Spoof reviews on Amazon[13]
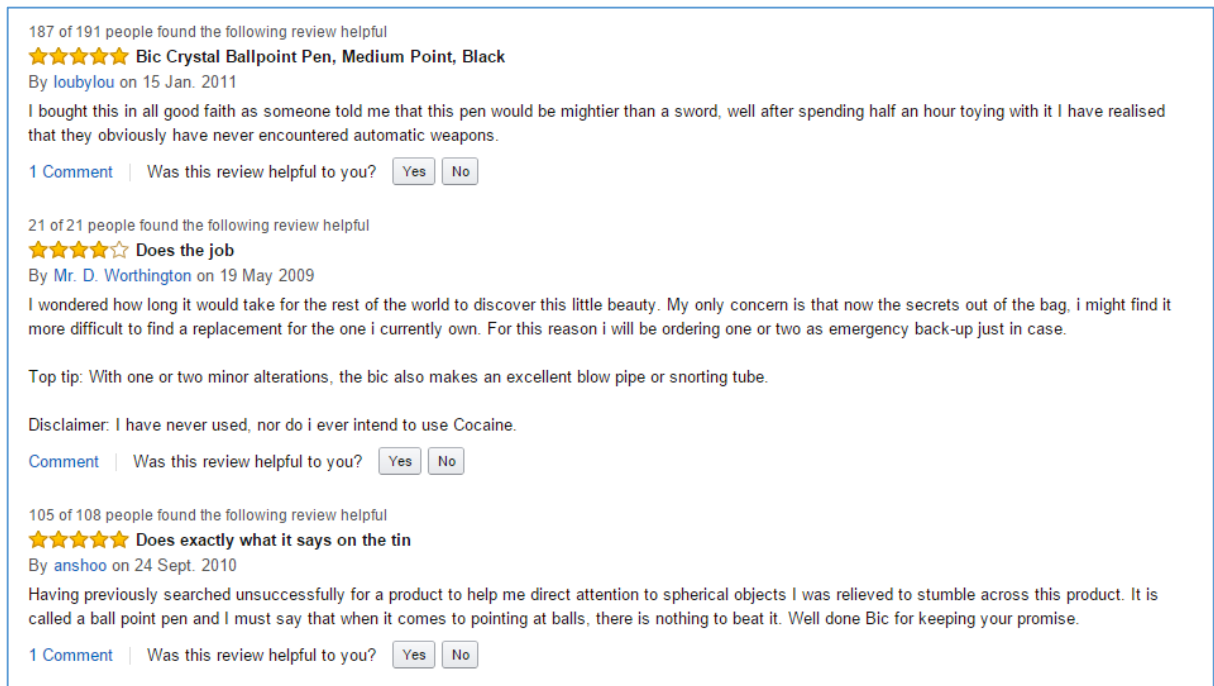
These reviews are clearly not intended as serious eWOM and so the number of people engaging with the content (by indicating that they have found it helpful) is not a useful measure of eWOM power. That said, a straight retweet without comment is likely to indicate endorsement of whatever is being retweeted. If the sender's intention was to mock then they would most probably use a modified retweet so they could add some commentary of their own to the tweet. On Twitter it appears that retweets and @mentions have different meanings, with retweets signalling endorsement of the original content and being about dissemination whilst @mentions generally include more nuanced discussion and hence perform the function of commentary (Bruns and Burgess, 2012). Conover *et al.* (2010) examine 250,000 tweets containing at least one politically relevant hashtag and find that the retweet and @mention networks have different structures and mean different things. The retweet network is almost entirely comprised of users retweeting things that they agree with and so forms a relatively

---

[13] http://www.amazon.co.uk/Crystal-Ballpoint-Medium-Point-Black/dp/B000JTOYLS

homogenous network of people who largely agreed with each other, providing further support for the argument that a retweet can generally be taken as an endorsement of the original content. In contrast the @mention network is much more politically heterogeneous, containing more debate between users at different ends of the political spectrum. @mention tweets form a bridge between the two sides of the political debate made by politically motivated users annotating tweets with commentary of their own, so it seems that @mention tweets indicate engagement with the original content but cannot be taken as a sign of agreement with it.

This provides further evidence that the debate between techno-optimists and techno-pessimists about whether social media enables people to purely engage with content they agree with is more nuanced than much research presents it. Whilst people may not necessarily retweet things they do not agree with, that does not mean that they are not exposed to the material in the first place. However, Conover *et al.* (2010) argue that this does not solve the problem of political polarisation – whilst people may mention others from across the divide they are very unlikely to share information from the 'other side' with their own networks.

3.2.3.1.   Retweets as eWOM

The retweet is the feature of Twitter that makes it particularly powerful as a form of eWOM. As Suh *et al.* (2010) put it, "Retweeting is the key mechanism for information diffusion on Twitter" (p178). When someone retweets they pass on the original message to their own followers, who in turn may pick it up and tweet it further. This amplifies the original message and can extend its reach many times, not necessarily to the benefit of the original sender. A tweet going viral can have extremely serious consequences for the individual who originally sent it, as in the case of MP Emily Thornberry and the other politicians already discussed, and joke or fake tweets can also have unforeseen repercussions once the tweet is picked up by the wider world. In 2013, hackers took over the Associated Press Twitter account and sent tweets falsely claiming that there had been an explosion at the White House and the President was hurt. In the time before the account was taken offline the original message was

retweeted more than 3,000 times and the US stock market dropped by 143 points before recovering (Kelly, 2014).

Given the potential power and reach of the retweet, it is useful to understand more about what factors influence people's decisions regarding whether to retweet a particular tweet or not. People are motivated to retweet for a wide variety of reasons: to spread information to a new audience, to entertain or inform, to start a conversation by commenting on someone's tweet, to ensure that one is visible as a listener, to agree with someone publicly, to validate the thoughts of the original tweet's author, as an act of friendship or homage to the original author, to give recognition to less popular tweeters, to gain followers or in the expectation of a return favour, and to save tweets for future personal access (boyd *et al.,* 2010). Although boyd *et al*.'s (2010) research is based on a non-random sample of Twitter users' self-reported motivations, it broadly supports the idea that retweets can generally be taken to indicate some form of agreement and engagement with the original content. No one suggested that they retweet as a way of bringing attention to content with which they do not agree.

A retweet differs from other forms of engagement with social media content insofar as it is a much more public form of endorsement. When I retweet something, that retweet appears in the timelines of all of my Twitter followers and I am then publically associated with the content of the tweet that I have retweeted. Essentially I am saying to my followers that on some level I endorse this content and consider it worthy of their consideration. This is not true with a Facebook or Instagram 'like' or indeed a Twitter 'like'. The only person who is guaranteed to see a 'like' is the person who shared the original piece of content. I may occasionally see content that other people 'like' in my Facebook feed but I certainly don't see everything that all my Facebook friends like because Facebook is selective about what it shows me. This is not the case with Twitter. As things currently stand my Twitter feed contains all the tweets that the people I follow have sent, although perhaps not in the order in which they were originally sent (Pierce, 2016) and of course with no guarantee that I will actually see them all.

The particular status accorded to retweets is shown by the fact that politicians regularly get into trouble not just for the content of their own tweets but for the content of the tweets that they retweet. For example, during the 2015 General Election campaign, Rosemary Healy, a Labour Party councillor, was reprimanded by Ed Miliband for retweeting a doctored version of a Conservative campaign poster (Figure 13) that drew parallels between the Conservatives' economic policy and that of the Nazis.

Figure 13 - Rosemary Healy's misjudged retweet

More recently, Andrea Leadsom, the Conservative leadership contender and prominent campaigner to leave the EU ran into trouble after she retweeted a message about the country being 'overrun by foreigners' (Figure 14).

Figure 14 - Andrea Leadsom's problematic retweet

These retweets are problematic precisely because the act of retweeting implies agreement with the content and turns it into a form of eWOM. Agreement with the content is one reason why someone might retweet a tweet but people generally don't retweet everything that they agree with so literature on eWOM alone cannot explain which tweets get retweeted and why. However, there is some literature looking specifically at what factors influence retweeting rather than any other forms of eWOM, and this will be discussed next.

3.3.   Factors influencing whether tweets get retweeted

There is a small but growing body of research that examines retweeting behaviour. Mostly this comes from the fields of computer science or statistics research where the researcher's aim is to build the most predictive model possible rather than necessarily to consider the factors that determine the 'predictiveness' of said model, however this research does still shed some light on what factors tend to be influential when it comes to determining whether a tweet gets retweeted or not and so is relevant for consideration here.

Unsurprisingly, the factors that influence whether a tweet gets retweeted are broadly similar to the factors that influence whether a piece of eWOM gets passed on. In both cases the decision to pass on the information or not is based on a combination of

three main factors: the characteristics of the information itself (in this case, the tweet), the characteristics of the sender of the information and the characteristics of the receiver of the information.

3.3.1. How message content influences online sharing

The chance of a piece of content being passed on depends in large part on the nature of the content itself. Broadly, there are two aspects of content that influence its chances of being passed on: the valence of the information in terms of whether it is positive or negative, and the nature of the information contained within the content, for example the topic of the content or its perceived usefulness. There is some extant research examining how the valence and content of eWOM messages influence their impact, however there is little consensus on the question of whether positive or negative messages are more likely to go viral (King, Racherla and Bush, 2014).

Research examining retweets of news information suggests that negative information enhances virality when it comes to news information whilst positive information makes non-news tweets more likely to be retweeted (Hansen *et al.*, 2011). This might suggest that politicians would do better to tweet negative news-based messages should they wish to be retweeted. However Berger and Milkman (2012) disagree, contending that in fact positive online news content is more likely to go viral than negative content. They also add an additional factor – emotional arousal – into the mix, arguing that high arousal emotional content, whether positive or negative, is more likely to go viral than low arousal emotional content. These findings suggest that politicians wanting to be widely retweeted would do best either tweeting positive content or ensuring that any negative tweets they send are designed to arouse high emotions. Stieglitz and Dang-Xuan (2013a) support the idea that tweets arousing strong emotion get more retweets than those with less emotional content, and, their research is of particular relevance here because it focuses on political tweets, albeit those discussing particular elections rather than those sent by politicians. However they state that they "found almost no support for the notion of negativity bias regarding retweet quantity and retweet speed (i.e., people do not tend to pass along

negative content more and at a faster pace than positive content)" (Stieglitz and Dang-Xuan, 2013a p241)

Whilst the evidence is mixed regarding whether people are more or less likely to pass on negative information, research suggests that they may still give more weight to it in their decision-making. For example, research examining the effect of product reviews on buyer behaviour indicates that negative information in product reviews prompts stronger responses in consumers in terms of changed beliefs about product performance and affect towards the product (Mizerski, 1982), suggesting that perhaps negative information carries more weight when it comes to influencing people's behaviour. Indeed, the percentage of negative reviews that a product has on Amazon correlates more closely with product sales than does the percentage of positive reviews (Cui, Lui and Guo, 2012). Similarly, negative tweets about movies appear to have more effect on the number of people who subsequently go and see a movie than do positive tweets (Hennig-Thurau, Wiertz and Feldhaus, 2014).

It also appears that whether a person is likely to pass information online not only depends on valence but also in part on the strength of the sharer's ties with their network. People in a less dense social network are more likely to share positive information whereas those in a tightly knit network are equally likely to share positive or negative information. In short, people are more likely to share positive information with acquaintances but will share both positive and negative with closer friends (Sohn, 2009). Twitter networks tend to be less dense, particularly those that cluster around politicians, and so this would suggest that people may be more likely to share positive eWOM information on Twitter than they would be to share negative.

Thus it seems reasonable to assume that tweet valence might play a role in determining how likely a tweet is to get retweeted but what that role might be is little considered in extant literature. Kupavskii *et al.* (2012) do include some consideration of valence in their retweet prediction model, however they extrapolate valence automatically based on the appearance of positive or negative terms and smileys in each tweet and so will miss many of the subtleties of valence (e.g. sarcasm) that a

human coder would pick up. Lemahieu *et al*. (2015) examine which tweets were most likely to be retweeted from 1,000 most recent tweets from 500 top Twitter users in eight different categories. They find that tweet sentiment is the least predictive of the factors that they consider, however their sentiment score is a measure of the overall strength of sentiment in a tweet – its emotionality - rather than whether the particular sentiment expressed is positive or negative. Additionally the sentiment of the tweets is determined by computer analysis rather than human coding, and so is likely to be much less accurate (Conway, 2006).

Dang-Xuan *et al*. (2013) analyze the valence of the tweets sent by a selection of tweeters in the run up to the Berlin state parliament election of 2011 and find that two content-related factors are significant in determining whether tweets get retweeted or not – the level of emotionality in the tweet and its content. The higher the level of emotionality, more likely it is to be retweeted. Emotionality in this context is about the strength of the emotion displayed in the tweet rather than about the valence specifically – strongly expressed positive or negative tweets are more likely to get retweeted than weakly expressed positive or negative tweets (in line with Berger and Milkman's (2012) contention that highly emotionally arousing messages are more likely to get passed on, irrespective of valence). Messages which contain some kind of appraisal of politicians or political parties are also more likely to be retweeted (Dang-Xuan *et al.*, 2013) showing that the topic of the tweet can also play a role in determining retweets.

When it comes to content, most extant retweet research focuses the structural elements of the tweet's content – for example the use of hashtags, inclusion of URLs and so on. Evidence on the role that hashtags and URLs play in determining retweets is mixed. For example Suh *et al*. (2010) and boyd *et al.*, (2010) both find that including hashtags or URLs in tweets significantly boosts their chances of getting retweeted. Dang-Xuan *et al.* (2013) agree that tweets containing hashtags are more likely to get retweeted, but find that including a URL in the tweet has no effect on retweet chances. However Liu *et al.*, (2012) find that including a URL in a tweet actually has a negative effect on the chances of it being retweeted although including multimedia

information such as pictures or videos has a positive impact. However, their research is based on examining which tweets get retweeted in cases of public emergency so they speculate that passing on tweets with URLs in them may be less relevant in this specific context than in many others. That said, Malhotra, Malhotra and See (2012) find that neither URLs nor hashtags increased retweets. However, their paper is practitioner-focused with very little detail regarding what methodology they used to come to this conclusion. The jury is still out on the role of hashtags and links, so it seems worth considering further the role they might play in determining the retweet levels of particular kinds of tweets (politicians' tweets, in this case).

Moving beyond an automated consideration of content elements such as hashtags and URLs towards a detailed consideration of the topic and purpose of the tweet, it seems reasonable to assume that one factor that might be relevant in determining whether a tweet is retweeted is what the tweet is actually about. Whilst there is a considerable body of literature examining the content and purpose of politicians' tweets (e.g. Tumasjan *et al.,* 2011; Golbeck *et al*., 2010; Hemphill *et al.,* 2013), there is very little research that moves beyond simply categorizing politicians' tweets and identifying topics that they typically tweet about to considering how effective these different types of tweets are at stimulating retweets. There is some existing research that considers content and sentiment when predicting retweets (of tweets in general, rather than specifically of politicians' tweets) but this is largely based automated calculation of content and sentiment variables using word counting or lexical similarity rather than on human consideration of the actual topic of the tweet or the intent of its author (e.g. Uysal and Croft, 2011; Kupavskii *et al*., 2012). The research presented here addresses this gap by presenting a detailed consideration of the role that both content and sentiment play in influencing retweets, using both machine coding and manual coding.

3.3.2. How the source of the content influences sharing

It has long been known that people give more weight to information that comes from sources they view as expert and credible (Sternthal, Dholakia and Leavitt, 1978). This also applies online. Who the original tweet is from plays a role in determining how

much weight someone gives to the content of a tweet and how likely they are to retweet it. For example, the number of followers a celebrity has influences consumers' perceptions of the credibility of that celebrity as a source of eWOM, with celebrities with low follower numbers having much less influence on consumers' product purchase intentions than those with more followers (Jin and Phua, 2014). However, the relationship is not a simple case of more followers equals more weight given to the content of tweets. Consumers are more likely to retweet negative eWOM tweets, particularly when they come from a celebrity with fewer followers. Jin and Phua (2014) speculate that this may be because negative information is seen as more reliable because it seems less self-serving, and that people may be more inclined to retweet such information when it comes from celebrities with a small number of followers because they reason that there is less likelihood of people in their network having already seen the information.

In Twitter terms there is a growing body of literature suggesting that the characteristics of the sender of the original tweet play an important role in determining whether or not a particular tweet gets retweeted. Previous research shows that source credibility – "the extent to which an information source is perceived to be believable, competent and trustworthy by information recipients" (Petty and Cacioppo (1986) cited in Liu, Liu and Li, 2012) – is used as a heuristic by members of online communities to assess the value of information posted. On Twitter, one can make a judgment of a sender's credibility based on the source's identity (if known), the number of followers that they have (on the basis that if someone has a lot of followers that means that many other people considered them to be worth following), whether their Twitter account is verified (Twitter gives verified status to the accounts of well-known people, verifying their identity), the number of tweets the person has previously sent (someone who is active on Twitter is likely to be viewed as more credible within the context of the Twitter ecosystem than someone who very rarely tweets), and the number of previous retweets someone has achieved (on the basis that retweets are a measure of how useful other people found that person's tweets to be).

Extant research in this area consistently finds that the number of followers that the message sender has is highly predictive of how likely their tweets are to be retweeted (e.g. Suh *et al.*, 2010; Zhang, Xu and Yang, 2012). However, Westerman, Spence and Van Der Heide (2012) find a curvilinear relationship between number of followers and source credibility, suggesting that having too many followers can be as damaging to perceptions of credibility as having too few. However their research does not examine how perceptions of source credibility then influenced retweets.

Another measure of source credibility is whether one's account is verified or not. Twitter applies verified status to the accounts of well-known people once it is confident that they are who they say they are. Thus account verification is a reflection of status in the offline world and seems likely to play a role in determining source credibility[14]. Indeed Petrovic *et al.* (2011) find that 91% of verified users' tweets are retweeted compared to just 6% of the tweets of non-verified users. However, this could also be because verified users will tend to be well known and hence have more followers, and it could be the follower numbers that are driving the retweets rather than the verification itself.

Extant research predicting retweets tends to take a random sample of tweets from across Twitter and examine how likely each is to get retweeted, rather than to consider a particular corpus of tweets. But it seems likely that the factors which influence source credibility might differ according to different types of senders – the factors that make a brand credible on Twitter may differ from those that make a politician credible. In a political context the relationship between Twitter credibility cues and the actual perception of a particular politician's credibility is further complicated by the fact that most people know who the politicians they are following are and have additional information about them beyond simply how many followers they have and so on. There are many other factors not directly related to their Twitter

---

[14] Note, in July 2016 Twitter announced that it will offer verified account status to anyone who wishes to apply for it and who is happy to tweet under their real name, so Twitter verification will no longer be limited just to public figures and will no longer be an indication of offline influence. However, at the time at which this research was conducted only celebrities and other well-known people had verified accounts and so account verification online could be taken as an indication of offline status.

profile that might influence a politician's source credibility, for example gender, age, how long they have been in parliament, party affiliation and the marginality of their seat. Indeed, there is evidence that these things influence how politicians use Twitter (e.g. Larsson and Kalsnes, 2014, Jackson and Lilleker, 2011) so it seems worth investigating whether they also influence how followers respond to politicians' tweets.

Thus it can be seen that the research presented in this thesis goes significantly beyond extant research on the role played by source effects in determining retweets by bringing in consideration of additional source characteristics which a politician's Twitter followers may be aware of but which go beyond simply information gleaned from their Twitter profile.

3.3.3. How the recipient of the content influences online sharing

Most extant retweet prediction research does not consider the role played by the characteristics of the tweet's recipient in determining how likely a tweet is to get retweeted, largely because information about the tweet's recipients is not easily accessible. Thus research considering why people retweet particular tweets and not others tends not to have any predictive element but instead to be conducted using methods such as survey research or interviews.

Individuals considering passing on eWOM information, for example by retweeting a tweet, may have a number of motivations to do so. Hennig-Thurau *et al.* (2004) identify several reasons why consumers might be motivated to participate in eWOM, which they define specifically as giving product feedback on review sites. These reasons range from relatively altruistic motives such as concern for other consumers and a desire to help the company through to more selfish motives such as self-enhancement, the reward associated with feeling part of a community, and a desire to seek redress when dissatisfied with a product or service. The same seems likely to be true of retweeting. Someone may be motivated to retweet something purely because they want to pass a useful piece of information on to their network, or because they

want to be seen as helpful by people within that network, or because they want to 'punish' an organisation or individual by passing on a negative message about them.

Retweeting has become a complex act, with many layers of hidden meaning. Retweeting no longer operates purely as an easy way to indicate endorsement of content. As this quote from satirist Henry Alford indicates, the decision of what to retweet and when can be a very complex one with many factors to consider:

> "Every time someone retweets one of my jokes, it sets off a spate of fretting about reciprocity… If the person is a total stranger whose feed I do not follow, then I will look at this feed and consider climbing aboard. I'll look at the ratio of how many tweets to how many followers that person has: if it exceeds 10 to 1, then I may suddenly feel shy. Because this person is unknown to me, I will feel no compunction to retweet a post of hers, though I may be tempted to 'favourite' one… [If I am followed by someone I know] suddenly the pressure mounts. I'll proceed to follow her, of course, if I don't already. Then I'll start feeling very guilty if I don't retweet one of her posts." (Silverman, 2015)

The literature suggests that whether or not someone retweets a tweet is likely to be a function of two factors: the extent to which they find the tweet personally engaging, and the extent to which they perceive a utility associated with the act of retweeting. Uses and gratifications theory suggests that people engage actively with different media, using those that most help them meet their needs. Uses and gratifications theory is commonly used as the basis for social media research (e.g. Chen, 2011; Dunne *et al.,* 2010; Hennig-Thurau *et al.,* 2004; Park, 2013; Parmelee and Bichard, 2012) and applies well to Twitter where users choose whom to follow and, by doing so, curate a personal news feed of content that interests them. They can then decide the extent to which they engage with the material they see based on how useful they find it and how well it meets their needs. Twitter users are also thinking about what message retweeting a particular tweet will send to their own followers – will the content be of interest or use to them (boyd *et al.,* 2010)?

The uses and gratifications approach has also been used in research examining broader word-of-mouth communication, suggesting that people pass information on for a variety of reasons including not only utility to the intended recipient but also utility to themselves (Hennig-Thurau *et al.*, 2004). Someone considering a retweet may therefore be thinking not only about how useful its content will be to other people but also about the extent to which that retweet will make them appear useful, helpful, insightful or knowledgeable to the rest of their network.

Whether people find particular tweets to be engaging and believe that retweeting will be useful to them in some way seems likely to be a function of three factors, which map closely onto the three factors that influence eWOM and viral sharing as already discussed. Firstly, the personal characteristics of the tweet recipient themselves such as the extent to which they are or perceive themselves to be an opinion leader in their network (Ma, Lee and Goh, 2013) and how closely they feel tied to the other members of their network (Sohn, 2009). Secondly, the nature of the tweet including, as already discussed, such considerations as whether it is positive or negative in valence (Hansen *et al.*, 2011; King, Racherla and Bush, 2014) and how emotionally arousing the content is (Berger and Milkman, 2012; Dang-Xuan *et al.*, 2013). Thirdly, the potential retweeter's perception of the tweet's original sender including the extent to which the sender is viewed as an opinion leader and the number of followers that they have (Jin and Phua, 2014).

3.4.   Conceptual model of factors influencing retweeting

As explained above, the literature suggests that once a tweet is sent three factors influence its chances of being retweeted: the characteristics of the tweet itself, the characteristics of the sender of the tweet and the characteristics of the recipient of the tweet. The discussion of existing literature on retweeting, eWOM, online virality and political Twitter use suggests a number of factors that might influence whether politicians' tweets get retweeted, each of which fits into one of these three categories. These are shown in the conceptual model in Figure 15.

Figure 15 - Conceptual model of factors influencing retweeting of politicians' tweets
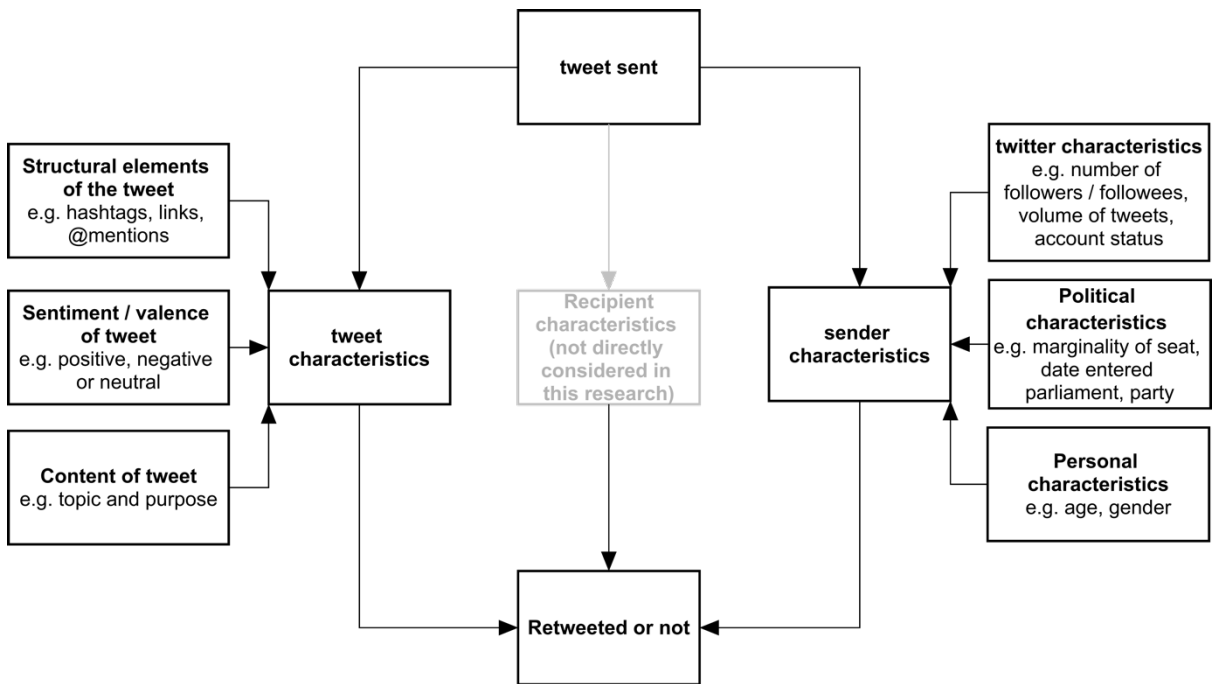
This research directly considers only the first two factors – tweet characteristics and sender characteristics. This approach is in common with extant retweet prediction research, which focuses almost exclusively on the structural elements of the tweet and the characteristics of the author (e.g. Petrovic *et al.*, 2011; Suh *et al.*, 2010), as this information can be easily extracted from the Twitter API whereas meaningful information about the tweet's recipients is not so easily available. The research presented here builds a number of predictive models to determine the chances of tweets getting retweeted, and a predictive model can only be built with variables that are known to the modeller. In this case, variables relating to the characteristics of the tweets and their senders are easily accessible, whilst variables relating to the recipients of the tweets are not. Indeed, the prime focus is on the content of the tweet, which is justified because content is the aspect of the tweet that the author can most directly control and adapt when trying to boost the tweet's popularity (Lemahieu *et al.*, 2015). One of the aims of the research is to generate practical recommendations which politicians can use to improve their effectiveness on Twitter, thus the research is restricted to considering only those factors over which the sender

of the tweet has some measure of influence or control. Politicians have no way of knowing who is going to follow them on Twitter and cannot control who their followers are. Equally, once a tweet is sent politicians have no control over who notices it in their timeline

3.5.   Gaps in the literature

Current literature makes it hard to draw any conclusions about the results or otherwise of social media use in political communication. This is largely because researchers have tended to examine different aspects of social media using different contexts and different methodologies, and so there is little consistency either in approach or in results generated.

At its heart, this thesis is about the nature of the relationships between politicians and citizens on Twitter and, specifically, about how citizens respond to politicians' tweets. Whilst there is a growing body of research examining how politicians tweet, there is a gap in the literature when it comes to understanding how citizens respond to politicians' tweets (Golbeck *et al.,* 2010). Extant research on politicians' Twitter use tends to focus on identifying and describing the different ways in which politicians' use Twitter in terms of the content of their tweets, but does not tell us anything about which of these tweets people respond to and why. The research presented here addresses this gap by moving beyond simply describing the different ways in which politicians use Twitter to identify how successfully those different kinds of tweets stimulate engagement. Just delineating the ways in which politicians use Twitter is interesting but only tells us half the story. Tweets on their own are meaningless if they do not stimulate some kind of response in the people who see them. An understanding of which tweets are most likely to lead to retweets moves research examining politicians on Twitter to the next level.

A limitation of extant research into predicting retweets is that it tends to overlook the role played by content and sentiment, except to the extent that these can be determined by machine analysis, and does not tell us anything about the chances of

specific types of tweets (such as those sent by politicians) getting retweeted. Machine sentiment analysis has substantial weaknesses when it comes to interpreting Twitter data, particularly as political tweets require an understanding of humour and sarcasm before their sentiment can be accurately categorised (Gayo-Avello, 2012). Thus, this research also moves retweet prediction research on by including a consideration of content and valence that is not simply based on machine coding and word counting but on a human coder determining the nature of each tweet.

There is a limited amount of literature looking at the context of UK politicians on Twitter and even less that specifically focuses on their Twitter behaviour during campaign times. Indeed there seems to be only one other paper that specifically addresses this context (Graham *et al.*, 2013). This paper identifies the different ways in which politicians use Twitter themselves but shares the limitation of other similar papers already discussed, insofar as it does not include any consideration of how voters responded to those tweets. There is a gap in the literature for specific consideration of not just the nature of MPs' campaign tweets but the effectiveness of those tweets. That is the gap which the research presented in this thesis addresses.

Research in this field dates quickly because of the fast-paced nature of change on the internet (Karpf, 2012). Extant research examining Twitter in election campaigns is largely based around the UK General Election of 2010 (Graham *et al.*, 2013) or the US Presidential Election of 2012 (Adams and McCorkindale, 2013). There is also some research looking at the Romanian Presidential Elections of 2009 (Momoc, 2010, 2012) and 2014 (Kereszturi, 2014), the European elections of 2009 (Vergeer, Hermans and Sams, 2010), the German Bundestag elections of 2009 (Plotkowiak and Stanoevska-Slabeva, 2013), the Spanish General Election of 2011 (Zamora Medina and Zurutuza Munoz, 2014) and the Norwegian General Election of 2013 (Larsson and Kalsnes, 2014). The UK General Election of 2015 provides an opportunity to bring this research up to date by considering a recent election that took place once Twitter was very firmly established as an important communication tool used by almost all politicians.

In summary, the research presented in this thesis directly addresses the clear gap between understanding how politicians tweet and understanding which tweets get retweeted. It identifies not only the different types of tweets that MPs send but also moves beyond that to determine the extent to which different types of tweets engage people.

3.6.   Chapter conclusion

As things currently stand we know a good amount about how politicians behave on Twitter. We also know a reasonable amount about how citizens use Twitter politically. However, we know very little indeed about how the two groups interact. This chapter builds on the research on political Twitter use presented in chapter two to show how literature on diffusion of innovation and electronic word-of-mouth has informed the current state of knowledge regarding which tweets get retweeted and why. Together these two bodies of literature form the basis of the conceptual model presented in Figure 15. Two of the three factors – the content of the tweet and the characteristics of its sender – will be examined further in the research presented here in a bid to address some of the gaps in this literature as it currently stands. Chapter four presents a brief discussion of the methods that have been used in extant research aimed at predicting retweets, as context for the discussion of the methods used in this research presented in chapter five.

Chapter 4   Methods used to predict retweets in extant research

4.1.   Chapter introduction

There is a small but growing body of literature in which researchers aim to build
models or use statistical techniques in order to predict either whether or not tweets
get retweeted or how many times they get retweeted. However, there does not
appear to be any research that uses the CHAID method of prediction presented in this
thesis. This chapter presents a brief discussion of this literature (Figure 16), largely
with a view to building an understanding of the methods that have previously been
used in order to support the claim that the method used in this research does not
appear to have been used before. There follows a discussion of the benefits of CHAID
analysis as a tool to predict retweets, as compared to other methods commonly used.

Figure 16 - How extant literature informs this research
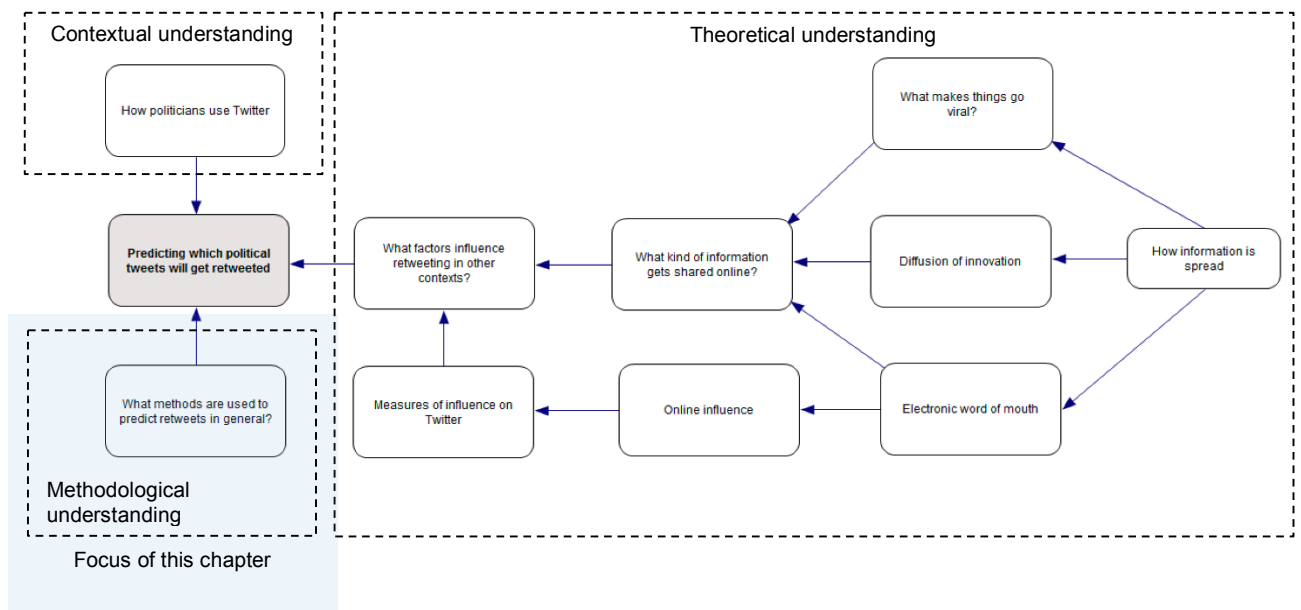
4.2.   Existing retweet prediction literature

Existing research on retweeting can broadly be divided into two categories – that
which considers what motivates an individual to engage in retweeting and the reasons
that people have for retweeting others' tweets (which can be thought of as audience

effects) (e.g. boyd, Golder and Lotan, 2010; Kim, Sung and Kang, 2014; Rudat, Buder and Hesse, 2014; Yang *et al.*, 2010), and that which considers how the characteristics of the tweet itself (message effects) and / or of the tweet's author (source effects) influence the chances of a tweet being retweeted (e.g. Bakshy *et al.*, 2011; Kupavskii *et al.*, 2012; Lemahieu *et al.*, 2015; Petrovic *et al.*, 2011; Suh *et al.*, 2010). The research presented here falls into the second category. Its focus is on better understanding how the message effects and source effects of politicians' tweets influence retweeting.

Research considering why people retweet has already been discussed in the preceding chapter so will not be discussed again here. Rather this chapter examines the small but growing body of research aimed at predicting retweets, with a particular focus on the prediction methods used, so as to explain how this literature influenced the methodological choices made for the research presented here. The bulk of retweet prediction research at present comes from the fields of computer science, statistics or mathematics rather than marketing or social science. The research presented here aims to go some way towards bridging the gap between these two worlds.

The emphasis of extant retweet prediction research is on the method and the process of prediction rather than on what can be learnt in practical terms from the results of the prediction. This research tends to divide into two groups that differ according to the dependent variable that they aim to predict: either whether or not the tweet has been retweeted (a categorical variable) or the number of retweets (a continuous variable), although there are a small number of papers that aim to do both (e.g. Hong *et al.*, 2011).

A significant limitation of extant retweeting research is that it tends to make very limited use of variables relating to the content of the tweets. Most researchers in this field share a common emphasis on using variables that can be simply extracted from Twitter as their predictors. This divides into two types: structural variables relating to the characteristics of the tweet itself (number of @mentions, use of hashtags, inclusion of URLs and so on), and descriptive variables relating to the author of the

tweet (size of Twitter network, number of tweets sent, whether their account is verified and so on). These variables are appealing because they can easily be extracted from Twitter and are readily available for analysis, however, as the literature on eWOM and online virality makes clear, other less easily measured factors relating to the content of the tweet are likely to have a significant influence on whether the tweet is retweeted and so analysis that purely focuses on variables readily available from Twitter is overlooking a significant part of the picture.

That said, there is some research which does consider content-related variables or which tries to move beyond the easily available Twitter data to gain a deeper understanding of the role played by the tweet's content or author characteristics. For example, Hong *et al*. (2011), Petrovic *et al.* (2011), Uysal and Croft (2011) and Kupavskii *et al.* (2012) (amongst others) all include some kind of measure of the content of the tweet in their modelling, but in every case this has been machine-calculated. For example, Hong *et al*. (2011) consider the topic distributions for each tweet simply by calculating how many times particular words appear in the tweets. Petrovic *et al.* (2011) also use variables relating to the appearance of particular words in tweets as well as considering the novelty of the tweet, calculated by comparing occurrences of words in the tweet with the same words in adjacent tweets in a user's timeline. Uysal and Croft (2011) calculate each tweet's novelty using a similar method, and its unexpectedness based on its similarity to other tweets by the same author, as well as including the appearance of content-related elements such as exclamation marks, first person pronouns, quotation marks and emoticons. Kupavskii *et al.* (2012) take this approach further by also including consideration of the tweet's valence, arousal and dominance, each computed automatically. Some other researchers also include consideration of the sentiment of the tweet (e.g. Lemahieu *et al.*, 2015; Liu, Liu and Li, 2012; Mahmud, Chen and Nichols, 2014) but in each case this is based on a machine-calculated sentiment score. The limitations of machine-based content and sentiment analysis will be discussed in more depth in chapter five. Suffice to say here that machine-coded sentiment scores often differ significantly from the scores generated by human coders (Canhoto and Padmanabhan, 2015), and machine-based

content analysis cannot accurately detect sarcasm or irony, nor accurately determine the intent of the tweet's author (Krippendorff, 2013). There is one paper in which human coders are used to determine the interestingness of a sample of tweets (Bakshy *et al.*, 2011) however this research only focuses on tweets that include a URL and the interestingness scores were generated by following the URL and assessing how interesting the content it linked to was, rather than assessing the interestingness of the tweet itself. As not all tweets contain links, this method does not give a true picture of the retweeting chances of all tweets, only those that have links.

4.3. Methods used to predict retweets

Because the focus of most extant retweet predicting literature is to build the most accurate model, there is a great diversity of methods used as each group of researchers tests a different method in order to demonstrate its effectiveness as a retweet predictor. However, these researchers all share a common aim which is to build a classification model – one which predicts a variable using one or more predictors. In this case the variable to be predicted is either whether or not a tweet gets retweeted or how many times a tweet gets retweeted.

There are three different approaches one can take to building classification models – rule induction models, traditional statistical models and machine learning models such as neural networks (IBM, 2014). Rule induction models use a set of observations to build formal rules that explain a particular phenomenon under study. The output of such models is a clear set of rules that the analyst can then apply to predict the outcomes of future observations. Decision tree methods are an example of the rule induction approach in action as they generate a set of rules which can then be used to classify new cases. Traditional statistical methods (such as logistic or linear regression) differ from rule induction models and machine learning models in that they tend to make more assumptions about the nature of the data (for example regarding distribution), and that their output takes the form of equations. Finally, machine learning models work by learning patterns in the data to predict outcomes, simply

producing a score for the likelihood of a particular outcome, with no rules or equation to explain how that score was reached.

Table 3 shows the methods used in extant literature. Papers that use multiple methods are shown in more than one table row. As Table 3 shows, statistical models and machine learning models are the most commonly used tools, with logistic regression being used more than any other method, followed by support vector machines and then a selection of different machine learning methods. There are a handful of papers that use decision tree approaches but none uses the CHAID (chi-square automatic interaction detector) method presented in this thesis.

Table 3 - Summary of retweet prediction methods

Type of classification model	Specific method	Literature
Statistical models	Logistic regression	Yang *et al.*, 2010 Liu *et al.*, 2012 Mahmud *et al.*, 2014 Lemahieu *et al.*, 2015 Yuan Huang and Zhang, 2015
	Principal components analysis	Suh *et al.*, 2010 Morchid *et al.*, 2014
	Autoregressive moving average model	Luo *et al.*, 2012
	OLS regression	Malhotra *et al.*, 2012
Machine learning models	Support Vector Machines	Yang *et al.*, 2010 Zhang *et al.*, 2012 Liu *et al.*, 2012 Gransee *et al.*, n.d. Morchid *et al.*, 2014
	Machine learning techniques (unclear which)	Webberley *et al.*, 2016
	Naïve Bayes Classifier	Morchid *et al.*, 2014
	Binary and multiclass classifiers	Hong *et al.*, 2011
	Passive aggressive machine learning algorithm	Petrovic *et al.*, 2011
	Neural networks	Morchid *et al.*, 2014
Rule induction models	J48 decision tree classifier	Uysal and Croft, 2011
	Gradient boosted decision tree model	Kupavskii *et al.*, 2012
	Regression tree modelling	Bakshy *et al.*, 2011

4.4.   Predictive modelling of retweets using CHAID analysis

CHAID is an example of rule induction modelling approach, in that it proposes a set of rules that can be used to describe groups within the data as they relate to the target variable. CHAID works by building a decision tree using relevant statistical tests – either chi-square or *F* tests depending on the nature of the target used (chi-square test in the case of categorical targets and *F* tests in the case of continuous targets). CHAID has an advantage over other rule induction models in that it can be used to predict any kind of target, either categorical or continuous, whereas most rule induction algorithms will look at one or other of these but not both. Although the research presented in this thesis focuses on predicting whether or not a tweet is retweeted, it is useful to know that the same method could in principle be used to predict how many times a tweet is retweeted, should that be desired.

4.4.1.   Why CHAID rather than statistical modelling or neural networks?

One could also use neural networks or traditional statistical models to predict whether or not a particular tweet would be retweeted. For example, logistic regression is commonly used to predict categorical targets (retweeted or not), and much of the existing retweet prediction research uses this method (see Table 3). However, this approach has a significant disadvantage. Logistic regression produces an equation which can be used to predict outcomes but cannot easily be translated into a simple set of business rules to be used to inform operational practice. This is also a disadvantage of using neural networks, which use a network of hidden connections to draw their conclusions and do not tell the analyst what they are doing (Struhl, 2015), making the extrapolation of meaningful operational business rules from the model virtually impossible. This is of no concern to computer science and statistics researchers whose goal is simply to build the most predictive model possible, but it is of relevance to marketing and business researchers who want to be able to understand the factors that influence retweets in order to be able to make operational decisions.

An additional benefit of using decision trees rather than statistical modelling or neural networks is that decision trees automatically ignore variables that do not add significantly to the predictive power of the model, meaning only important attributes are included in the final model. This contributes to making decision tree models simpler to interpret and gives a clear picture of which variables are important and which are not. It also means that if there are two variables that include very similar information, the CHAID model will only include one of them, making the model more efficient and meaning that the analyst does not need to worry about potential overlap between variables, something that is likely to be common in exploratory modelling such as this where there are lots of variables all measuring similar things being thrown into the models.

The choice of modelling approach should be driven by both business needs and by the nature of the data. In this case the business need is to provide politicians with workable rules that they can use to maximise the chances of being successful at stimulating engagement on Twitter. The building of a predictive model is an activity designed to enable this goal, rather than the ultimate goal itself, as it is unlikely that a politician will spend time feeding draft tweets into a predictive model in order to predict their chances of being retweeted. Therefore, a rule induction approach is suitable as this provides clear guidelines for effective tweets which can be easily operationalised. This is not the case with either the formula that a logistic regression produces nor with the somewhat opaque output of a neural network.

Another advantage of CHAID is that decision trees are less affected by extreme values than are neural networks or logistic regression (IBM, 2012). This is particularly relevant when using Twitter data as many aspects of Twitter use – number of followers, number of tweets sent, number of retweets per tweet – tend to have highly skewed long tail distributions with many extreme values. This is the case with the dataset analysed here, as a few politicians are extremely active on Twitter and thus have many more followers than average, or send very many more tweets. Similarly, a small number of tweets attract hundreds of retweets but the vast majority attract very few.

Two other substantial advantages of CHAID are that its output is highly visual and hence easy to understand, and that it works on non-parametric data (unlike alternative approaches such as multiple regression). This is important because the dataset analysed here includes a substantial number of non-parametric fields (all those that contain ordinal or nominal data) and many fields which do not follow a normal distribution.

4.4.2. How CHAID works

CHAID is a decision tree approach to predictive analytics which uses significance testing to identify the factors which are relevant in predicting a particular outcome (in this case a tweet being retweeted or not) and the order of their significance. Working with a categorical target (the variable to be predicted), a CHAID model classifies cases into one of two groups – in this case tweets that are retweeted and those that are not retweeted. Variables are flagged as either the target (the outcome to be predicted) or predictors (the inputs that will be used to build the predictive model). Cases are split into two groups – training data and testing data – before the modelling begins. The initial model is trained on data where the outcome is known and then tested by being applied to a clean sample of data that it has not seen before.

The CHAID algorithm begins by evaluating each predictor in turn and selecting whichever one best splits the training data into two groups based on the values of the predictor (retweeted yes or no). Its aim is to identify the two purest subgroups possible in relation to the outcome of interest. Thus one group will have as many retweeted tweets are possible and the other will have as many non-retweeted tweets. The algorithm then looks at each of the two groups and splits them according to the same principles. This continues until no further subgroups can be usefully generated (or until the maximum number of splits that the analyst specified has been reached).

Unlike other decision tree algorithms, CHAID can perform multiple splits in a single step so that the resulting tree is not binary (with only two subsamples coming from each node). Multiple subsamples can be present at each split making the resulting tree wider rather than taller. There are other decision tree algorithms available but

this ability to split variables into as many groups as are shown to be significant gives CHAID a big advantage over them and gives it its particular analytical strength (Struhl, 2015).

There are of course some limitations to the CHAID method. Recursive partitioning algorithms such as CHAID can sometimes be unstable, meaning that different results can be obtained even when running the algorithm on the same dataset. This is because if the algorithm gets to the point where two or more variables are pretty much equal in terms of their predictive value, then it will choose between them randomly and this can then significantly change the way in which the dataset splits from that point forward (Fast *et al.*, 2012). This is not a problem if the main goal is to build an optimally predictive model, but it means that the model is only likely to have high accuracy on the particular dataset for which it was developed and is likely to be much less accurate if used on other datasets (Fast *et al.*, 2012). This is not a particular problem for this research as the aim is to demonstrate the effectiveness of this general approach to retweet prediction rather than to build a model that would hold true across multiple datasets.

Building a predictive model using CHAID is an art rather than a science – there is no such thing as the optimum decision tree – different analysts will come up with different trees. This is because small changes such as removing one record from a dataset can have a big impact on how the model looks (Struhl, 2015). The analyst controls how many variables to include and what rules to set for the model (for example by specifying a maximum number of levels to the tree) and this kind of analytical modelling is an iterative process. The analyst continues until she feels that she has the best model amongst the alternatives available (Fast *et al.*, 2012). 'Best' in this instance means "[creating] a model that has the most useful information and that has very good predictive power." (Struhl, 2015 p191). Generally the aim is to keep the model efficient and stop growing it at the point at which additional changes only yield very small increases in explained variance (when predicting continuous variables) or correct classification (when predicting categorical variables).

An important part of this kind of modelling is dividing the data into two sets – a training dataset and a testing dataset. Cases can be randomly allocated to one of the two groups by the software being used (SPSS Modeler in this case). The dataset is partitioned so that a proportion of the cases are used to train the model with the remaining cases used to test it. Doing this means that the reliability of the model can be determined by seeing how it performs on testing data which it has not seen before and which was not used in the development of the initial model.

4.4.3.   CHAID and text analytics

This research uses text mining[15] and predictive modelling. The purpose of text mining "is to process unstructured (textual) information, extract meaningful numeric indices from the text and thus make the information contained in the text accessible to various data mining (statistical and machine learning) algorithms." (Dell, 2013). In summary then, text analytics is about turning text (in this case tweets) into numbers and enabling those numbers to be analysed using traditional data mining techniques.

Decision trees in one form or another have been used for text analytics since 1994 (Apté, Damerau and Weiss, 1994). Using this approach, the accuracy of predictive models can be improved through combining the structured numeric information (in this case tweet metadata) with unstructured text (the content of the tweets) which has been transformed into numerical variables. The predictive model uses the variables generated from the text analytics process as just another set of input variables that can be used to improve the accuracy of the prediction. The use of this technique takes the research presented here beyond extant retweet prediction research which makes very limited use of the text of the tweets when predicting whether or not they will get retweeted.

4.5.   Chapter conclusion

Whilst there is a growing body of research addressing the problem of predicting retweets, to date there is limited focus on determining which factors are most

---

[15] The terms 'text mining' and 'text analytics' will be used interchangeably.

influential with a view to then enabling people to make more effective use of Twitter. This research extends existing retweet prediction research in two ways. One is by demonstrating the use of a new method – CHAID analysis – which, as far as can be determined, has not previously been used this way. The other is by including consideration of the content and valence of the tweets in the predictive modelling process, using variables based on human coding rather than machine coding. These variables capture a richness of understanding of the purpose and nuance of the tweets that machine coding cannot achieve and so add another dimension to understand of the factors that drive retweeting. In the following chapter the specific choices of method will be explained in more detail.

Chapter 5 Methodology

5.1. Chapter introduction

This chapter discusses the methodological approach taken by this research and places the specific methods used within the context of extant research in this field. This research is, at its heart, a predictive analytics project. Specifically, it uses two forms of text mining – content analysis and sentiment analysis – in conjunction with CHAID to build predictive models that estimate the likelihood of tweets being retweeted, based on two types of predictors: the characteristics of the tweet and of the tweet's sender, as per the conceptual model shown again in Figure 17. The aim is to identify factors that increase the chances of a tweet being retweeted in order to develop a set of practical guidelines for politicians wishing to improve their effectiveness on Twitter.

Figure 17 - Conceptual model underpinning this research
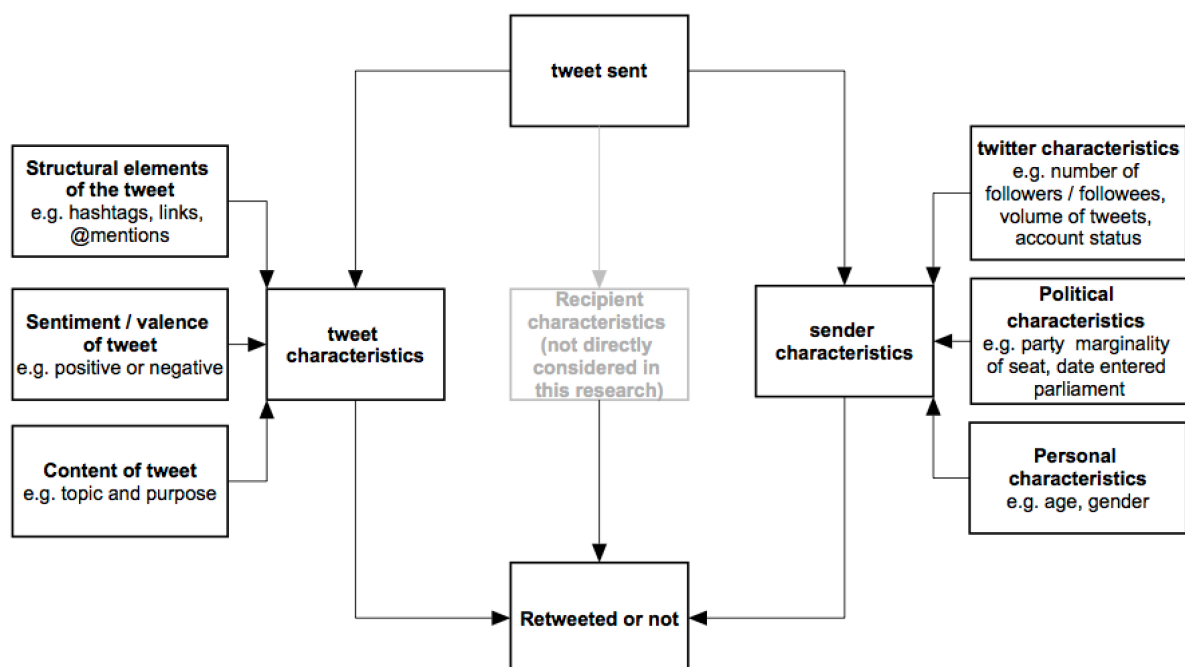
This chapter presents a discussion of the ontological and epistemological assumptions underpinning the research, followed by a detailed overview of the methods chosen and the reasons for their selection. Finally, it concludes by outlining the stages of data collection and analysis involved required in order to examine the relationships suggested by the conceptual model and answer the research questions posed in chapter one.

5.2.   Research objectives

As outlined in chapter one, the overarching research question to be addressed in this thesis is: **what factors influence whether politicians' tweets are retweeted?** This question breaks down into a series of sub-objectives, shown in Table 4 along with a summary of the methods used to address them.

Table 4 - Summary of research objectives and methods used to address them

Objective	Method
1. To identify the factors that extant literature suggests will determine whether politicians' tweets are retweeted.	Review of existing literature on retweeting in order to determine what is already known about factors that influence retweeting.
2. To identify other factors which might also play a role in influencing whether politicians' tweets are retweeted.	Review of literature on social media and politics, online virality and electronic word of mouth to identify additional factors that may be relevant in determining political retweets but are not considered in extant retweeting research.
3. To propose a typology of the tweets sent by UK politicians during the 2015 General Election campaign and identify which most effectively generate retweets.	A review of existing tweet typologies formed the basis of manual content and sentiment analysis of a sample of tweets, with additional codes being added during the process as required. From this, a typology of political campaign tweets was developed.
4. To test the extent to which the factors identified in objectives 1 and 2 do indeed determine the chances of a tweet getting retweeted by building predictive models using CHAID.	A dataset of relevant tweets was obtained and additional variables appended to it, based on findings from the literature review. Manual content and sentiment analysis were performed on a sample of tweets and relevant content-related and valence variables added to the data. This enriched dataset was then subjected to CHAID predictive analysis.
5. To demonstrate a new method of predicting retweets – CHAID analysis – which could be of use to social media researchers and marketers in other fields.	A literature review of existing retweet research was performed to demonstrate that CHAID analysis has not been used in this way before. Using SPSS Modeler, a number of CHAID models were built which demonstrate how the technique can effectively be used to predict retweets.
6. To provide practical advice for politicians regarding how to best to use Twitter to engage with citizens as part of a campaign communication strategy.	Recommendations written up based on findings of research outlined.

5.3.   Ontology and epistemology

This research is based on the critical realist ontological perspective. This section briefly reviews other research perspectives commonly used in marketing research before outlining the ontological and epistemological assumptions underpinning critical

realism and explaining why the critical realist approach is particularly well-suited to social media research of this nature.

The dominant research paradigm in marketing is positivism (Hunt, 1990) and almost all extant research analysing aspects of Twitter behaviour takes this approach. Although the ontological and epistemological assumptions underlying such research are almost never explicitly identified, the preponderance of research based around statistical measurement shows that positivism dominates. However, whilst a positivist approach can tell us a lot about *what* is happening it is limited when it comes to explaining *why* things are as they are. Positivism also assumes an externally knowable objective reality and ignores the complexity of relationships and social phenomena.

Social media, by its very nature, is a social phenomenon involving highly complex networks of relationships, so whilst a positivist approach is useful to a point, it is limited in the extent to which it can explain observed phenomena in social networks. Positivist research is generally concerned with applying statistical methods to samples of data with a view to making generalisations that hold true for the wider population. However such generalisations can be problematic if the phenomena under study are context-bound and historically or culturally specific (Buch-Hansen, 2013). This is certainly the case with any analysis of Twitter because Twitter operates in 'internet time' (Karpf, 2012) and hence is in a constant state of flux – what is true today on Twitter will not necessarily be true tomorrow – and we cannot make generalisations about other kinds of social networks or other Twitter contexts based on the findings from research examining Twitter and politics.

The limitations of positivism as a way of understanding marketing phenomena could be taken to mean that a constructivist approach would be more useful for conducting research into marketing. However, Sobh and Perry (2006) argue that this approach too is of limited use to marketers. Marketing is about managing transactions taking place in an external marketplace. Whilst this is partially influenced by individuals' constructed meanings of things such as brands, it is also heavily influenced by external structures and mechanisms such as the economy which are clearly not constructs in

the minds of the individual participants in a transaction but have some kind of objective reality beyond them. Likewise, a social network such as Twitter exists as an objective reality beyond the minds of its members, and there are many externally measurable mechanisms and structures associated with it (retweets, @mentions, number of followers and so on) which social media marketers treat as important measures of the success of their campaigns.

Thus it seems that neither of the two dominant paradigms – positivism or constructivism – are particularly suited marketing research (Reige, 1998), particularly not research examining social media. Positivism ignores the complexity of the relationships within social networks whereas constructivism ignores the influence of clearly measurable external phenomena which affect the ways in which they operate. In marketing both these things matter and so a 'middle way' is needed. This middle way is critical realism (Ackroyd and Fleetwood, 2000).

Critical realism (Bhaskar, 2008) combines a realist ontology with a subjectivist epistemology. Observable social phenomena exist separately from the individuals within them. Whilst critical realists believe that there are objects and structures which exist in the world independently of our knowledge of them, they accept that it is not possible to know this reality objectively. The best that can be hoped for is that each iteration of research will shed some more light on a phenomenon. Pragmatically, the best theory is the one which does the best job of describing how things seem to be (Collier, 1994).

This feels like an accurate description of the social media and political worlds as they really are but not as they are represented in much positivist research. As Hirschl *et al.* (2009) point out, models of political behaviour tend to treat society as being the sum of the individuals within it, each person in society acting with individual agency. However, society is ontologically distinct from the individuals within it and individuals are socialised, both materially and mentally, to behave in particular ways. Thus, social phenomena have an existence beyond the individuals within them. This perspective seems particularly relevant to the context of social media networks which, in terms of

power and meaning, are clearly comprised of much more than just the sum of the individuals within them. Individuals interact through social networks in ways that can produce unobservable social structures about which they themselves may not be aware (Buch-Hansen, 2013).

Critical realists view individuals' actions as being both constrained and facilitated by pre-existing social structures, and that these structures in turn are developed or transformed through social interaction (Bhaskar, 2008). There are certainly parallels between this view and the way that a social network like Twitter operates. One could argue that pre-existing social structures (such as notions of power and celebrity in the 'real world') heavily influence behaviour on Twitter. People's behaviour on Twitter is also governed by the social structures of the 'Twitterverse' which has its own unwritten rules and etiquette but these rules and structures are constantly being changed by the collective behaviour of people within the network. A good example of this is the development of the hashtag convention, something which developed spontaneously from within the network rather than being introduced by Twitter as a feature.

Critical realists agree with postmodernists that the world is socially constructed; however, they refute the idea that these social constructs are all that exists. For critical realists, unlike postmodernists, social phenomena exist independently of the actors within them. The task of social scientists is to try and get beyond the surface appearance of experiences and perceptions in order to uncover the so-called generative mechanisms – the underlying causal principles behind a phenomenon (Ackroyd and Fleetwood, 2000). Whilst positivists might focus on trying to predict future events, critical realists would generally argue that the open character of social systems means that exact predictions of the future are not possible, so researchers would do better to focus on explaining past and current phenomenon (Buch-Hansen, 2013). That is the case with this research which is primarily focused on explaining the factors which influenced engagement with politicians' past tweets. Generalisability beyond this is limited and the research is not aimed at building a predictive model of future retweets but rather aims to build understanding of what has happened in the

past with a view to developing general guidelines as to the kind of thing which might work in the future.

Zinkhan and Hirschheim (1992) argue that a critical realist approach is well suited to the study of marketing phenomena, many of which act as enabling or inhibiting agents rather than as primary causes. Adopting a critical realist approach to marketing research moves one beyond theories which simply predict behaviour towards an understanding of the underlying generative mechanisms which explain it. Rather than just counting retweets and building a statistical model which identifies those tweets that have the greatest chance of being retweeted, this research uses a more detailed content analysis and sentiment analysis to try and uncover the generative mechanisms which explain *why* particular tweets are more likely to be retweeted, something which has been largely overlooked in Twitter research to date.

5.4.   Research design

Critical realists typically use a retroductive research design, in which the research process is modified and developed after each stage in the light of new information. The retroductive research process starts with the construction of a conceptual model (Figure 17) which could explain an observed regularity. The researcher then tries to establish the existence of the structures and mechanisms underlying the model, generally using a variety of methods. Critical realists are not so much concerned with uncovering a single objective truth as with uncovering as many different aspects of their phenomenon of interest to build understanding. Hence, mixed methods are commonly used as such an approach enables triangulation of findings, and different methods are suited to uncovering different aspects of a phenomenon. Here, the research mixes descriptive statistical analysis, machine-based content and sentiment analysis, and manual content and sentiment analysis. This retroductive approach works well in text analytics projects like this one because in text analytics the analyst generally does not start with an a priori hypothesis or even clear ideas regarding which aspects of the text will be most helpful when it comes to addressing the research question (IBM, 2013). The aim is to uncover some patterns or regularities in

the data and from that move towards a theory of retweeting, rather than to start with theory of retweeting to then be tested.

Quantitative methods were used throughout this research project, in line with the dominant approach taken in the field. There has been some qualitative research in this area, generally taking the form of interviews or free text questionnaires asking people to talk about how engaged they feel they have been by Twitter content (e.g. boyd *et al.,* 2010; Parmelee and Bichard, 2012). However, the vast majority of research in this field is quantitative, and the qualitative approaches taken have significant limitations. Whilst there is value in understanding people's perceptions of how they were or were not influenced by Twitter content, in order to understand better what types of content really influence them to act one needs to examine actual behaviour rather than stated behaviour, in this case retweeting.

All the 154,565 tweets sent by sitting MPs during the 2015 UK General Election campaign (the period between the dissolution of Parliament on 20 March through to polling day on 7 May) were collected using Brandwatch. 366 MPs sent at least one tweet during this period. Replies to other people or retweets of other people's content were removed, leaving 42,444 original tweets as the basis of the analysis. All the analysis was performed using SPSS Statistics, SPSS Modeler and SPSS Text Analytics (the choice of these tools is explained further in section 5.7.1.). The stages of the data analysis were as follows:

1. **Descriptive statistical analysis** of the 42,444 original tweets to identify patterns in the MPs' tweeting behaviour, enable comparison of the behaviour of this group of politicians with others reported in extant literature, and to develop an understanding of the factors that might most influence retweeting.

2. **A smaller sample of tweets was created**, matching the 6,510 tweets that were not retweeted with a random sample of the same number of tweets that were retweeted, thus creating a sample of 13,020 tweets in which exactly half were retweeted and the other half were not.

3. **A series of predictive CHAID models were built** and run on the sample of 13,020 tweets. In every case, the sample was split into two groups – training data and testing data. The models were built using the training data and then tested on the testing data. Different types of variables were used in each model in order to better understand how each influenced retweeting. These models were as follows:

   a) Model based on variables related to the structural elements of the tweets

   b) Model based on variables related to the authors of the tweets

4. **Machine-based content and sentiment analysis** were performed on the sample of 13,020 tweets and new content and sentiment-related variables were used as the basis of a further set of CHAID models.

5. **Manual content and sentiment analysis** were performed on a smaller sample of 1,212 tweets (evenly split between those that were retweeted and those that were not). The resulting variables were used as the basis of a further set of CHAID models.

6. **A 'master model' was built** bringing all the predictive variables together.

A visual representation of the research design is shown in Figure 18, along with an indication of which section of the thesis covers each stage of the process.

Figure 18 - Research design
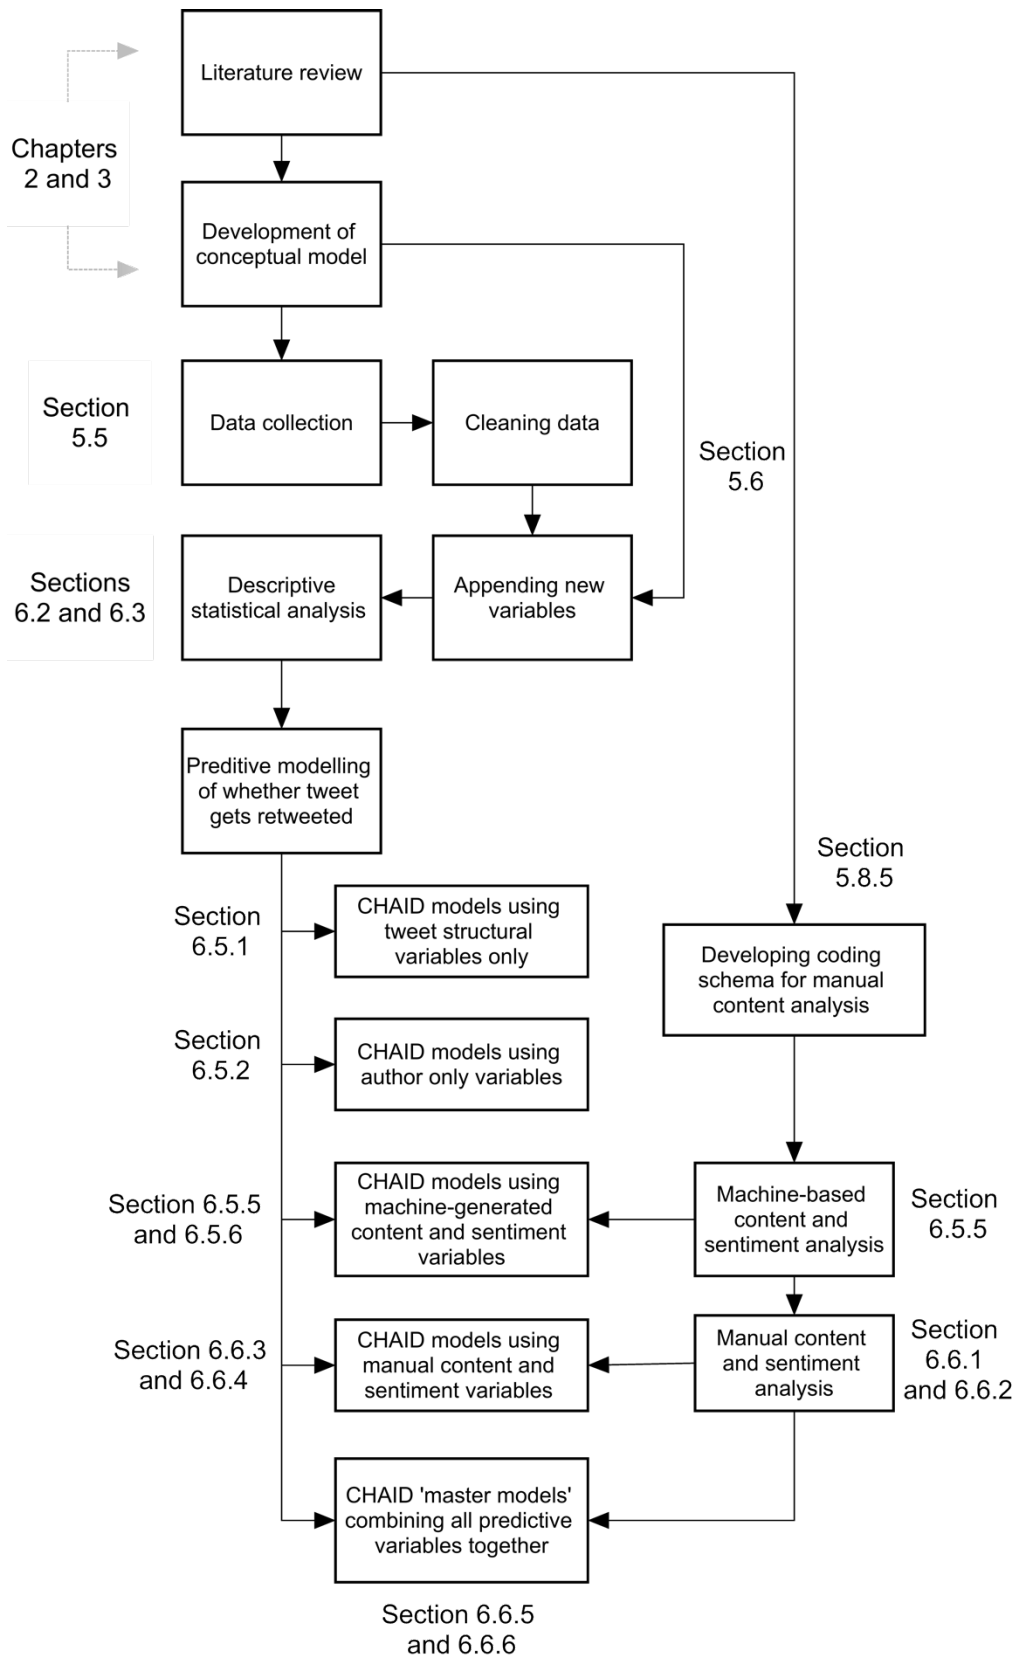
5.5. Sampling at each stage of the research

Phase one of the research – the descriptive analysis – was based on the full dataset of 154,565 tweets in order to build as full an understanding as possible of the tweeting behaviour of MPs. Once the analysis moved to the predictive modelling phases, all the MPs' retweets and replies were removed from the dataset leaving 42,444 original tweets. Of these, 35,934 had been retweeted at least once (85%), skewing the data in favour of retweets. This meant that a predictive model would have been right almost 85% of the time simply by predicting that every tweet would get retweeted. To correct for this, the 6,510 tweets that had not been retweeted were matched by a random sample of an additional 6,510 tweets that had been retweeted, giving a new dataset of 13,020 tweets in which each tweet had exactly a 50% chance of having been retweeted. These 13,020 tweets were used as the basis for the bulk of the predictive modelling. The only exception to this is the final stage of the research – the manual content and sentiment analysis. It was not possible for the researcher to manually code 13,020 tweets in the time available so, in common with much similar analysis in the field (e.g. Hemphill *et al.,* 2013; Dang-Xuan *et al.,* 2013; McKelvey *et al.,* 2014), a randomly selected sample of 1,212 tweets was coded, also evenly weighted between retweeted and not retweeted. The final CHAID models which included the manual content and sentiment variables were built using just this sample of tweets. A summary of which tweets were used at which stages of the analysis is shown in Figure 19.

Figure 19 - Summary of samples of data used at each stage of the analysis
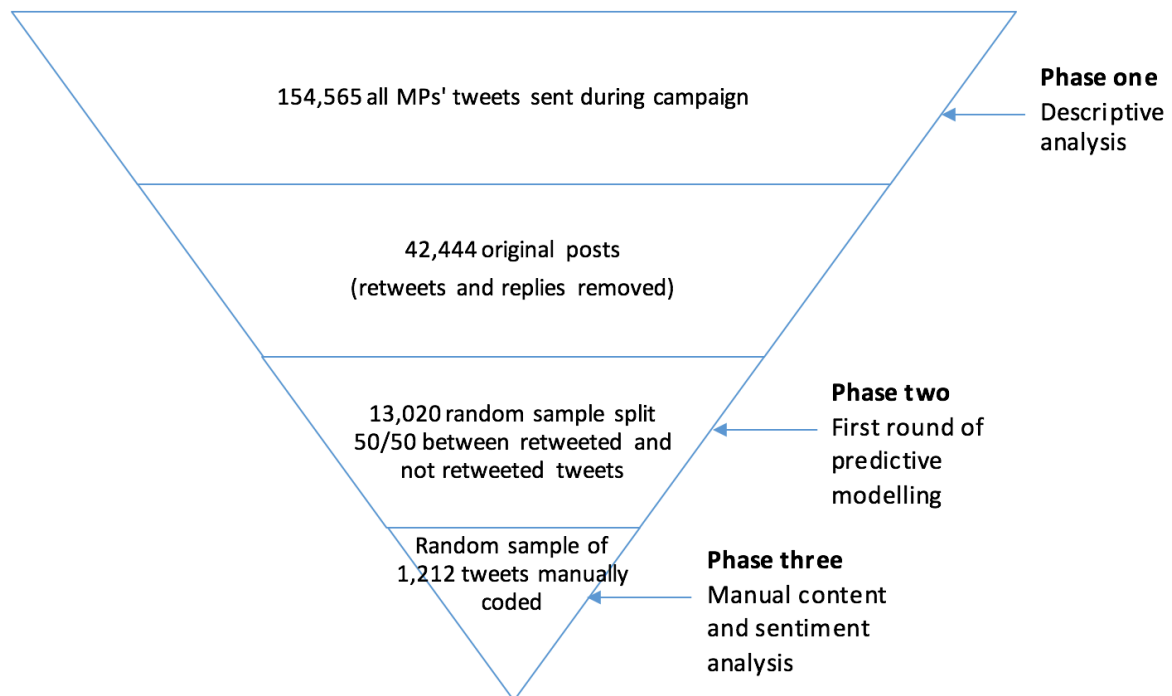
5.6.   Cleaning and preparing the data for analysis

The first stage of the research was a lengthy period of data cleaning and formatting in order to prepare the dataset for analysis. During this phase the accuracy of the Twitter-generated variables was checked (for example to ensure that the genders Twitter had assigned to authors were correct) and missing information was completed manually, for example by adding a gender designation for MPs to whom Twitter had not indicated one, including confirming the gender of MPs with ambiguous names (e.g. Alex, Chris) via their websites.

The variables included in the analysis were sourced in one of the following ways.

- **Extracted from Twitter** – the Twitter API provides a host of information both about each tweet and about its sender, for example the number of retweets generated, number of followers the sender has, whether the tweet includes an @mention or a link and so on. It also includes some variables that Twitter has calculated based on the characteristics of the tweet and its sender such as

each tweet's reach, calculated by adding the sender's number of followers to the number of followers that each person who retweeted the tweet has, giving an estimate of the total number of people who could have seen the tweet.

- **Manually researched and added** – the variables downloaded from Twitter did not include any political information about the authors of the tweets. However, the characteristics of the author of a tweet are likely to influence whether that tweet is retweeted and in the case of politicians, extant research suggests that information about their party affiliation, political position and so on could be relevant. Thus, additional information about each MP was researched and manually added to the dataset, for example their party affiliation, the year they entered parliament, their age, and the marginality of their seat going into the 2015 campaign.

- **Calculated from variables provided by Twitter** – the core variables extracted from Twitter were used to calculate a number of new variables that extant research suggested might be relevant. For example, each MP's ratio of followers to 'followees' was calculated, along with other measures of their Twitter behaviour such as the total number of tweets they sent during the campaign, the percentage of their tweets that were retweets, the mean number of retweets per tweet and so on. Variables were added to record whether tweets contained hashtags and, if so, how many, based on a count of the number of times the # character appeared in each tweet. Variables were also created to record whether tweets contained any of the particularly popular election-related hashtags (e.g. #GE2015, #leadersdebate).

- **Appended by Brandwatch** – Brandwatch adds value for its clients by appending additional variables of its own to the data extracted from Twitter. For example, it generates a sentiment score for each tweet (positive, negative or neutral)[16] as well as appending influence and outreach scores for each author using Kred.

---

[16] As Brandwatch is a commercial organisation, the exact method it uses to assess sentiment is proprietary information and hence not revealed.

- **Sentiment and content analysis** – both machine and manual sentiment and content analysis were performed (to be discussed further later in this chapter), during which process the sentiment of each tweet was assessed (positive, negative or neutral) and each tweet was coded based on its content. The development of this coding scheme will also be discussed later in this chapter. As a result of this process 19 new variables were added to the dataset.

No information about the receivers of the tweets is included in this dataset as this is not readily available either from Twitter or from Brandwatch. Additionally, the focus of this research is on helping politicians to understand how they can manipulate the things they can control (the content of their tweets and their own behaviour on Twitter) and so receiver information is not relevant, thus all the variables included in this analysis relate either to the tweets themselves or to the tweets' authors.

Broadly speaking, these variables divide into three types:

- Those that relate to the author of the tweet
- Those that relate to the content of the tweet and its valence (what the tweet is about, whether it is positive or negative)
- Those that relate to the structural features of the tweet (whether it includes hashtags, links, mentions of other people or other structural elements).

A full list of all variables included in the analysis along with details of what each one means and how it was sourced is provided in appendix two.

5.7. Data analysis

The cleaned dataset was imported into SPSS Modeler and SPSS Text Analytics for analysis. This section briefly introduces these tools and explains why they were chosen.

5.7.1. Why SPSS Modeler and Text Analytics were used

This project uses SPSS Modeler and SPSS Text Analytics. To some extent, this choice is pragmatic as I already have some knowledge of these tools and access to a network of

expert users who could provide further advice. However SPSS is also, by some margin, the most commonly used advanced analytics tool in academic research – between 1995 and 2013 it was used in more than twice as many academic articles as its nearest competitor SAS (Muenchen, 2014) and other researchers examining political use of Twitter have also used it (e.g. Larsson and Hallvard, 2011; Parmelee and Bichard, 2012; Ma *et al.,* 2013; Hermida *et al.,* 2012 amongst others).

SPSS offers substantial advantages for academic researchers, particularly that it does not require the researcher to learn a programming language in order to analyse her data (as would be the case with the open source software R) and there is no limit on the size of dataset that can be analysed. As commercial software tools, SPSS Modeler and Text Analytics offer more reliability and support than do open source alternatives such as R. Freeware alternatives tend to offer only basic functionality or are focused on specific application niches that make them inappropriate for this research (Fast *et al.*, 2012).

SPSS's Text Analytics module is designed, amongst other things, for analysis of social media data. It uses natural language processing, can handle extremely large datasets and enables integration between structured and unstructured data (such as that represented by tweets). It goes beyond merely counting words and instead uses an understanding of sentence structure, context and meaning in order to group concepts intelligently (to a point) as well as identifying mentions of entities such as people, places and organisations.

Another substantial benefit is the integration between SPSS Text Analytics and SPSS Modeler. The researcher can build a core predictive model in Modeler using all the data available before the content analysis takes place. She can then code her text data in Text Analytics (or allow Text Analytics to do this automatically) and use the results of that analysis to create variables which can then be seamlessly fed back into the initial model. This can be run both with and without the text data to see how the text variables affect the model's predictive power.

5.8.   Content analysis

This research goes beyond extant research in the field by including variables derived from content analysis, both machine and manual, in the predictive models generated. "Content analysis is a research technique for making replicable and valid inferences from texts (or other meaningful matter) to the contexts of their use." (Krippendorff, 2013 p24). In this case, the texts are tweets and the objective of the research is to be able to make replicable and valid inferences about the conditions which maximise the chances of a politician's tweets being retweeted.

Content analysis involves "the systematic assignment of communication content to categories according to rules, and the analysis of relationships involving those categories using statistical methods" (Riffe, Lacy and Fico, 2014, p3). Although other content analysis definitions have been proposed, there is a general agreement that content analysis requires the objective and systematic application of explicitly laid-out rules and procedures, central to Riffe *et al.*'s definition. The rules and procedures which were developed for this piece of research will be discussed in more detail later in this chapter.

5.8.1. Content analysis and Twitter

Content analysis is commonly used to analyse Twitter data but its full potential as far as predicting retweets is concerned has not yet been realised. Around two thirds of academic papers focusing on Twitter use some form of content analysis as their main method and around 16% use sentiment analysis (Zimmer and Proferes, 2014). It is common for Twitter researchers to combine methods in a single piece of research, as is the case in this thesis which combines content and sentiment analysis with predictive modelling.

Content analysis offers four key benefits to researchers - it is unobtrusive, it enables the analysis of unstructured material, it is sensitive to context and it can cope with substantial volumes of data (Krippendorff, 2013). These four benefits particularly apply to Twitter research. Tweets are unstructured, require an understanding of context and offer researchers a huge volume of information. Analysis of Twitter

communications also enables researchers to build a greater understanding of how people really behave rather than, as is the case in much social science research, of how people say that they behave, and to do so in an unobtrusive way. People do not always tell the truth to interviewers or researchers when asked about their political behaviour. For example, in Northern Ireland polls regularly under-report the level of support for Sinn Fein because people are reluctant to reveal their true behaviour to researchers (Coakley, 2008). However, the chances of people's behaviour on Twitter being influenced by the fact that academic researchers may access their tweets are minimal (Bruns and Burgess, 2012).

Of course, this does not mean that there is no bias in Twitter data. Social networks such as Twitter do not provide a 'window' through which one can view people's true selves. People actively manage how they present themselves on social media and their social media personas are likely to include some elements of the real them alongside some fictitious elements (Manovich, 2012). However, what we can say is that it is highly unlikely that people are consciously changing their behaviour on Twitter or actively managing how they present themselves specifically in response to the actions of academic researchers.

An additional benefit of using Twitter for content analysis research is the immediacy of the data collection. Tweets meeting pre-determined criteria can be collected virtually in real time and the analysis can begin straight away. Relatively little data preparation is required compared to that which would be needed for analysis of other kinds of content. The tweets are collected complete and already in digital form. They do not require transcription or digitisation, simplifying the process considerably (as well as reducing the cost) compared to more traditional content analysis of, for example, newspaper articles or party election broadcasts.

Content analysis research should have one of three aims: to describe the nature of a communication, to make inferences about its antecedents or to make inferences about its effects (Holsti, 1969). In emerging fields it is common for researchers to focus primarily on the first of these, concentrating their efforts on describing the

characteristics of a particular kind of communication (McMillan, 2000). Twitter research is very much an emerging field and so the bulk of extant research indeed concentrates on describing the nature of different aspects of Twitter communications such as, for example, the tweets sent by politicians. The research described in this thesis takes describing the characteristics of tweets as its starting point but then moves beyond that to consider both the antecedents of a particular kind of communication (the retweet) and to make inferences about the effects of particular kinds of communications (what kinds of tweets are most likely to be retweeted?).

5.8.2. Computerised content analysis

Traditionally content analysis was done manually, with human coders working through the data and assigning cases to categories. Whilst computer-based content analysis has been around since the 1960s (Conway, 2006), the development of the internet and availability of more computing power as well as more sophisticated computer programmes over the last ten years or so have resulted in powerful computer-based content analysis tools being much more widely accessible to academic researchers.

Technological developments have affected the practice of content analysis in two main ways. Firstly, huge volumes of data are now available to researchers for the purposes of content analysis. In particular, the internet and social media have led to a huge proliferation of text-based data, most of which is relatively easily accessible to content analysts. At the same time, computing power has grown, enabling the automated content analysis of massive volumes of data to be performed extremely quickly. Analysis which would once have required access to a mainframe or super-computer can now be performed using sophisticated content analysis software on a user's desktop (Lewis, Zamith and Hermida, 2013).

Computers enable content analysis to be performed on whole populations of data rather than on samples, and on volumes of data that would be impossible for human coders to work through in any sensible timescale. For example, Mckelvey *et al.* (2010) based their examination of the 2010 US presidential election on analysis of

547,231,508 tweets, whilst Cha *et al.* (2010) based their analysis of Twitter following, retweeting and mentioning behaviour on a dataset of 1.7 billion tweets. Datasets of this size are not uncommon. Such massive volumes of data can only be analysed by computer and the combination of readily available data and computers with sufficient analytical power means that more and more content analysis research, particularly in the field of social media, is based on computational approaches.

When manual content analysts are faced with too much data they either have to hire more research assistants or apply sampling techniques which add complexity and provide another opportunity for error to creep into the research (Lewis, Zamith and Hermida, 2013). In contrast, computer-based content analysis enables the researcher to analyse entire populations of data without any sampling being necessary, as is the case in this research project. An additional benefit of computer-based content analysis is that it offers perfect reliability and reproducibility. Whereas human coders will vary the way in which they apply coding rules, a computer will apply such rules in a completely precise and unvarying way. The same data will be coded in the same way irrespective of when the computer codes it and the coding rules, once determined, will always be applied accurately and without variation (Leetaru, 2012), removing a substantial weakness of much human-based content analysis.

However, computer-based content analysis does have significant weaknesses itself. In particular, computers are not yet good at understanding the context of communications and are generally unable to pick up on sarcasm, jokes or hidden meanings (Canhoto and Padmanabhan, 2015). Tweets often require a huge amount of contextual information and understanding of subtext in order to make sense. The tweet which precipitated Emily Thornberry's resignation is an excellent example of this (see Figure 20). I recently found myself struggling to explain to a group of digital marketing MSc students, none of whom were from the UK, why this tweet caused such uproar. To comprehend this requires an understanding of the British class system, of the context of the Rochester by-election, of the relationship between MPs and voters, of the concept of the 'white van man' and the assumptions that go with that, of the significance of displaying the George Cross outside one's house as well as

a host of other factors. Without the relevant context, none of the students could understand what the problem was. A tweet like this can be tricky for humans to interpret. A computer would have no chance.

Figure 20 - Tweet which led to Emily Thornberry's resignation

For this reason, comparisons of the results of human content analysis with computerised content analysis have found significant differences between the two. For example, Conway (2006) found that computers and humans came up with very different results when tested on the same sample of data (news articles relating to the 2002 Texas gubernatorial primary). In particular, Conway found that whilst human coders could use complete assertions as their unit of analysis, the computer coders could not identify assertions and so were restricted to using individual words as their unit of analysis, leading to a substantially higher counting of issues and attributes than amongst the human coders. This is a particular weakness when the content analysis is

based on lengthy texts such as newspaper articles, when there could be ambiguity regarding what counts as the unit of analysis. This is less of a problem for Twitter research because tweets are short and make a convenient and unambiguous unit of analysis that both computers and human coders can easily identify.

Another area where computers have a weakness is with regard to identifying emotion and valence. Conway (2006) also found that the computer was much poorer at accurately identifying the affect and valence of mentions of candidate attributes and issues. The human coders were able to produce much more nuanced results than the computer. However, this result is unsurprising – as already discussed, a clear weakness of the computer-based approach is computers' lack of sophistication when it comes to picking up on subtext and nuance. The type of content analysis presented in Conway's (2006) research – searching the data for examples of campaign issues and candidate attributes and trying to determine how those attributes and issues are seen – does not play to the strengths of computer analysis so a comparison between computer coding results and human coding results is not a fair fight. In reality, it is more likely that researchers conducting this kind of research would use a combination of computer and human coding to help answer different parts of the research questions.

There are two elements in Twitter data – structural elements such as hashtags, retweets and mentions, and context-related elements such as the choice of words in the tweet itself and surrounding contextual information required in order to make full sense of the tweet. Whilst computer-based analysis can be extremely effective at analysing the structural elements of tweets and also at "[mapping] the overall data landscape" (Manovich, 2012 p469), human coding is generally still required in order to make sense of the context-related elements and so a hybrid approach offers the best of both worlds (Lewis, Zamith and Hermida, 2013). In particular, computer analysis is good at determining how the different structural elements of a tweet relate to each other. For example, are tweets including a picture more or less likely to be retweeted? Does including a hashtag influence one's chances of a retweet? If so, does the number of hashtags in a tweet make a difference?

However, there are also aspects of Twitter data which do not lend themselves to computer-based analysis. A particular challenge is that people tweet for an almost infinite variety of different reasons. This variety of intention is not usually a problem when analysing other types of content such as, for example, newspaper articles or advertising copy as these texts tend to be created for a common purpose and with a common audience in mind (Lewis, Zamith and Hermida, 2013). This is not the case with tweets and so human coding is the only way to begin to understand a tweet's context sufficiently to be able to assign meaning to a particular hashtag or other pieces of information contained within that tweet. It is for these reasons that a triangulation-based approach, with a combination of both computer and human content analysis, is used in this research. This blended approach has been used before in research focused on categorising different kinds of political tweets (e.g. Dang-Xuan *et al.,* 2013) but does not appear to have been used in research focused on predicting retweets where content analysis, if it is used at all, is generally machine-based. Thus, the research presented in this thesis extends the use of manual content analysis into a field in which it does not appear to have previously been used.

5.8.3. Sentiment analysis

Sentiment analysis is closely related to content analysis and focuses on "the extraction of positive or negative opinions from (unstructured) text" (Thelwall *et al.,* 2010 p2545). Machine-based sentiment analysis generally uses natural language processing to categorise content into one of three groups – positive, negative or neutral sentiment. In principle, any text can be used for sentiment analysis and researchers are increasingly turning to social media as an easily accessible pool of analysable data. Twitter sentiment analysis, in particular, is a fast-evolving field as shown by the growing number of commercial providers offering some form of sentiment analysis to clients (e.g. Brandwatch, Sysomos and SproutSocial, to name but a few). However, the techniques and algorithms that they use are based on proprietary software and are commercially sensitive so generally not revealed, meaning academic researchers wishing to be open about how their results were achieved must develop their own approaches.

Tweets present particular challenges for sentiment analysis (Canhoto and Padmanabhan, 2015). As more sophisticated techniques have developed over time, so the unit of analysis which can be subjected to sentiment analysis has grown more precise, from document-level classification to sentence-level to, more recently, phrase-level (Agarwal *et al.*, 2011). However, the short nature of tweets still presents a challenge for even the most sophisticated sentiment analysis algorithms. Tweets tend to contain abbreviations, acronyms, emoticons and spelling mistakes (often deliberate rather than unintentional, as many sentiment analysis models assume (Thelwall *et al.*, 2010)), and are often written in a very informal style, making the sentiment much harder for a machine to interpret (Achananuparp *et al.*, 2012). These problems can be addressed to some extent, for example by developing a dictionary of emoticons and acronyms along with their associated sentiment, such as that compiled by Agarwal *et al*. (2011). However, keeping such a dictionary up to date is a challenge as Twitter is constantly evolving and new acronyms and internet terminology are emerging all the time.

In the case of politicians' tweets, the use of informal language, emoticons, acronyms and emotional intensifiers such as exclamation marks and all capitals may be less of a problem than in a sample of tweets from the general population. Tweets are part of the way that politicians present themselves to the public and form part of their campaign communication so are likely to be more formal and 'traditional' in writing style than the tweets of another group might be. However, politicians' tweets do present another challenge to accurate sentiment analysis. Computer-based sentiment analysis generally assigns valence to content based on the presence of words that commonly signify negative or positive emotion, words such as 'yes' and 'no'. In the case of political campaigning, however, the words 'yes' and 'no' also have very particular campaign-specific meanings. For example, in the case of the 2014 Scottish Independence Referendum campaign, a tweet from someone who supports the no campaign, using the word 'no', cannot necessarily be assumed to be a negative tweet.

Extant research shows that sentiment analysis approaches which are usually applied to non-Twitter data generally need to be adapted in order to make sense of tweets.

For example, a common approach is for sentiment analysis researchers to manually code the valence of a sample of data for the purposes of training a computer-based model, which can then be deployed to analyse additional cases which have not been pre-coded. Kouloumpis *et al.*, (2011) take this approach in their sentiment analysis of tweets. They find that traditional part-of-speech analysis is of limited use when measuring Twitter sentiment. However, they develop fairly accurate models by using Twitter-specific features such as hashtags and emoticons and emotional intensifiers to collect training data. They build a corpus of emoticons and hashtags that generally signify positive or negative emotions, collect data featuring these signifiers, train their model using this data and then use it to identify the sentiment of other tweets which did not include the hashtags or emoticons. This is the approach most commonly taken in existing retweet prediction research. Those papers in the field that use sentiment analysis all use automated methods based on the calculation of valence according to the appearance of certain words, emoticons and signifiers of intensity that are assumed to be an accurate predictor of valence (e.g. Uysal and Croft, 2011; Kupavskii *et al*., 2012; Lemahieu *et al*., 2015) Go *et al.* (2009) also use a large sample of tweets containing emoticons as training data, working on the assumption that the presence of a positive emoticon in a tweet means that the tweet is positive whereas a negative emoticon means the tweet is negative. They then test this approach by applying it to a sample of tweets which may or may not include emoticons. A limitation of this approach is that it does not include a neutral option which the authors acknowledge is important in the real world. An additional limitation is, as already discussed, such signifiers of valence may be absent from the more formal tweets that we might expect MPs to send. Additionally, one cannot assume that a positive emoticon means that a tweet is positive. For example, most automated content analysis models would consider the 'smiley face' emoticon ☺ to be an indication that a tweet is positive, however that is not always the case. Take the example tweet from Eric Joyce in Figure 21. The intention behind this tweet is to criticise the Guardian – this is not a positive tweet although a computer sentiment analysis tool would most likely score it as such

due to the presence of the smiley face. Knowledge of context is necessary in order to be able to accurately score tweets for sentiment.

Figure 21 - Ambiguous use of emoticons
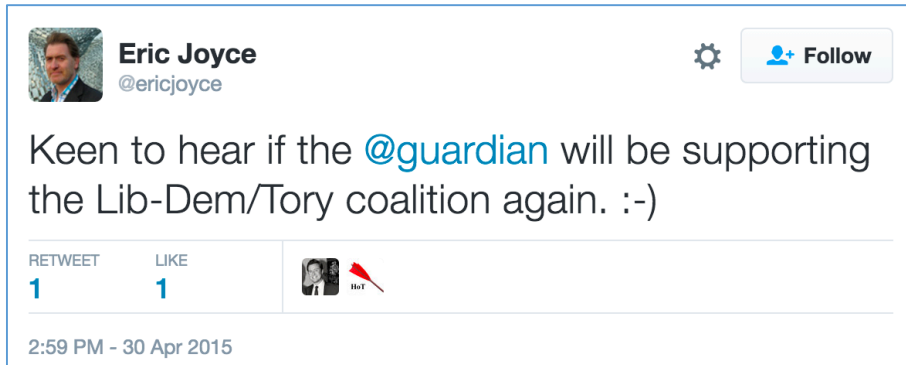
5.8.4. Content analysis units

The first step of a content analysis project is to determine the recording units – the basic unit of text which will be analysed. In this case, each individual tweet is a recording unit. This constitutes whole text coding. Coding a complete text as a single unit can make coding reliability hard to achieve (Weber, 1990) however in this case it is appropriate because tweets are extremely short, being limited to 140 characters and generally cannot be meaningfully broken down into smaller units.

5.8.5. Development of the coding scheme

A number of frameworks already exist for categorising types of political tweets (e.g. Tumasjan *et al.,* 2011; Golbeck *et al.,* 2010; Hemphill *et al.,* 2013) but there is no uniformity in the field and only very limited attempts to re-test existing frameworks in new contexts. Instead, researchers have tended to develop coding schema which are relevant to their particular context of interest. There is no pre-existing coding scheme that is directly suitable for application to the two political contexts being examined in this project. Existing coding schema tend to either relate to an American context (Golbeck, Grimes and Rogers, 2010; Hemphill, Otterbacher and Shapiro, 2013) or to have been developed for categorising non-campaigning tweets (Jackson and Lilleker, 2011). This is par for the course as far as text analytics projects go - "a common

problem [when extracting meaning from unstructured text] is that an a priori list of terms or themes is not available. Instead a heterogeneous corpus of text must be analysed to extract themes and meanings with respect to the dimensions of interest" (Fast *et al.*, 2012, p57).

Additionally, the way in which people use Twitter evolves as Twitter matures and so it is not unreasonable to assume that how MPs were tweeting during the 2010 General Election will differ in some way from how they are tweeting now. Thus it is important to develop a coding scheme that builds on what has gone before whilst taking into account the specific characteristics of current tweeting behaviour. Therefore, rather than starting completely from scratch, this research builds on coding schemes that already exist, adapting the elements that are applicable to this research context and modifying them where necessary.

It was clear when reading through the tweets that there were some which could not easily be fitted into the categories developed from existing literature and so new categories were developed as required. Prime amongst these were charity tweets – tweets in which MPs express support for a particular charity unrelated to their wider political campaign. No extant research includes such a category but it became clear when scanning the tweets that these kinds of tweets popped up fairly regularly and so deserved a category of their own.

Another new category developed during this research was 'achievements'. Most of the tweets in this category could also be referred to as humblebrags[17] - attempts by the MPs to draw attention to something good that they had done whilst not appearing to be too self-promoting such as the example in Figure 22.

---

[17] "An ostensibly modest or self-deprecating statement whose actual purpose is to draw attention to something of which one is proud" (OED)

Figure 22 - Example of a personal achievement tweet

Most pre-existing schema have some kind of category for tweets that inform people about the politician's ongoing activities (e.g. Sæbø, 2011) but the tweets in the 'achievement' category go beyond simply informing people of what the politician is doing towards taking credit for a specific achievement (more in line with Lawless's (2012) 'credit-taking' category). Also included in this category are tweets promoting the wider achievement of the MP's party rather than just the MP as an individual, such as the example in Figure 23.

Figure 23 - Example of party-based achievement tweet
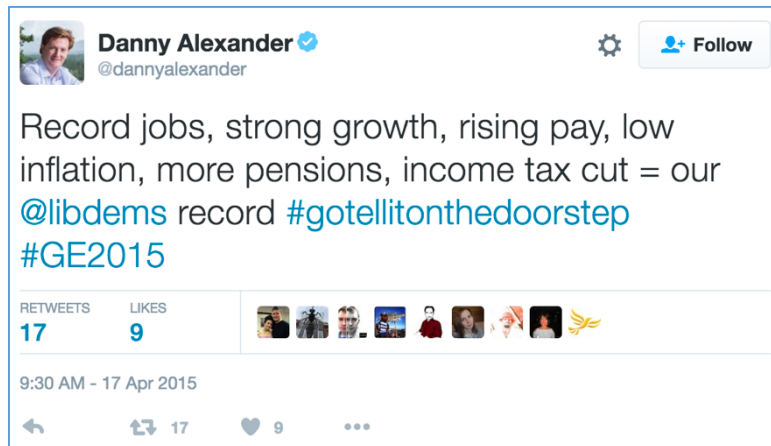
Existing schema make very little mention of attack tweets – those in which a politician explicitly critiques or attacks an opponent. This is a surprising omission given the volume of literature that exists considering various other aspects of negative communication in political marketing and campaigning (e.g. Fridkin and Kenney, 2011; King and McConnell, 2003; Chou *et al.,* 2011; Dermody and Scullion, 2000; Lau *et al.,* 1999; Wicks and Souley, 2003 amongst very many others). That said, Graham *et al.* (2012) include a category for tweets which critique or argue and Sæbø (2011) has a category for tweets which feature dialogue with other politicians – this category would include negative attacking tweets when sent as part of an ongoing policy discussion between politicians but not when sent as standalone tweets. In the research presented here, the category 'attack' includes all negative mentions of opposition parties in general or specific named politicians, whether sent as standalone tweets or as part of a wider conversation on Twitter.

Fear appeals are a particular form of negative or attacking tweet in which the sender aims to provoke fear in the message's recipient by warning of the dire consequences of not voting a particular way (Calantone and Warshaw, 1985). A fear appeal is "a persuasive communication attempting to arouse fear in order to promote precautionary motivation and self-protective action" (Ruiter *et al.,* 2001, p614). In order to change behaviour a fear appeal needs to contain two elements, firstly a

threat ("staying in the sun too long puts you at risk of skin cancer") and then a recommendation for action ("use sunscreen"). Researchers have tended to concentrate on the use of fear appeals in social marketing (e.g. Strong and Dubas, 1993; Dillard and Anderson, 2004) rather than in other contexts. The emphasis is generally on examining how fear appeals change behaviour and the role of the negative emotion of fear is to strengthen the ability of the message to change recipients' attitudes or behaviour (Antonetti, Baines and Walker, 2015).

The relevance of this to politics is clear as political advertising is designed to persuade people to change their voting behaviour (from not voting to voting, switching allegiance from one party to another, and even from voting to not voting) and messages designed to stimulate fear are common (Dermody and Scullion, 2000), being present in up to 34% of negative ads (Kaid and Johnston, 2001). In a political context fear appeals can be thought of as:

> "a particular type of emotional appeal that attempts to scare voters or raise their fears about specific issues or character traits. For this reason fear appeals are often found in negative ads and are frequently combined with other types of evidence. Almost one fifth of all presidential spots use some type of fear appeal." (Kaid and Johnston, 2001, p58)

Some of the most memorable political advertising campaigns of recent years make explicit use of them. The Conservative Party made extensive use of fear appeals during the 2010 election campaign, all designed to stimulate fear that a Labour government would damage the UK's chances of economic recovery (almost exactly the same central idea as the Labour Party used in their campaign, albeit the other way around). George Osborne expressed this idea during the campaign saying "The battle is going to be between hope and fear – hope that the change that the Conservatives can bring can get our economy off its back and the fear of a Labour Government that will throw all sorts of scare stories at us." (Porter, 2010). More recently, politicians on both sides of the Scottish Referendum campaign in 2014 and the EU Referendum campaign of 2016 were criticised for exploiting the fears of voters, with the EU remain campaign in

particular being characterised by its opponents as 'Project Fear' (Stewart and Asthana, 2016).

There is a small body of extant research examining the ways in which fear appeals are used within political campaigns (e.g. Dean, 2005; Brader, 2005; Calantone and Warshaw, 1985) however there does not appear to be any consideration to date of the use of fear appeals within political social media communications. Thus, it was decided to code fear appeal tweets into their own category. All fear appeal tweets are negative but not all negative tweets are fear appeals so it would be interesting to see if particular types of negativity are more or less successful at stimulating retweets.

Many of the attack tweets were sent in response to television appearances by opponents – Question Time, the Leaders' Debate and the Challengers' Debate in particular generated a lot of negative tweets. Indeed, it became clear when going through the tweets that a large number, both negative and positive, were sent in response to media appearances. Given the rise of 'second screening' in general (Curtis, 2014) it is not surprising that MPs tweet whilst watching television or listening to the radio. Thus, a category was included for tweets which include some mention of either television or radio.

In addition to negative attack tweets it was also necessary to have a new category for positive support tweets. These are tweets in which the MP praises a colleague, expresses support for something a colleague has done or makes some other positive statement about another politician as in Figure 24.

Figure 24 - Example support tweet mentioning party leader

A related category was developed for tweets in which the MP passes on supportive messages that they have received themselves as per the example in Figure 25.

Figure 25 - Example of tweet passing on support for self

As this research focuses on an election campaign, unsurprisingly the majority of the tweets were campaign-focused in some way. Several of the previous political tweet schema were developed outside of election times and so did not include many categories that were directly relevant here (e.g. Golbeck *et al.,* 2010; Hemphill *et al.,* 2013; Lawless, 2012). However, upon examination of the literature, it became clear that some codes could still be developed from these non-campaign-related pieces of research.

One of the first coding schemes developed (Golbeck *et al.,* 2010) does not distinguish between information-giving tweets and those which take a position on an issue. This was later critiqued by Hemphill *et al.* (2013) on the basis that too many tweets would then come into this category and that taking a position on an issue is different from merely providing information. This is the approach taken here. Informational tweets are those which simply state a piece of information in a dispassionate fashion with no additional commentary stating the MP's own views. Commonly, but not always, these tweets will provide a link to another website. Note, the links given in the tweets have not been followed up so it is possible, indeed likely, that neutral information-giving tweets can contain links to highly partisan position-taking sources. The focus of the coding is entirely on the tweet itself and not on any additional content to which it links.

In contrast a position-taking tweet is one in which the MP clearly states their own opinion with regard to the issue under discussion, or repeats the party's line on the issue. Lawless (2012) also uses two separate categories for tweets that contain neutral information and those that take up a position. Graham *et al.* (2013) make a further distinction between position-taking tweets that present a candidate's own view compared to those that present the party line. When coding these tweets however it was rarely possible to make that distinction as MPs tended to stick very closely to the party line and there were no examples of tweets that presented an MP's personal view as in any way different from the party line. This may be indicative of the increasingly professional and more tightly controlled use of social media by political parties during the 2015 election when compared to the 2010 election which was the focus of Graham *et al*.'s research (2013).

Most of the pre-existing Twitter coding schema include a category for personal tweets (e.g. Golbeck *et al.,* 2010; Jackson and Lilleker, 2011; Lawless, 2012; Graham *et al*., 2012; Sæbø, 2011). Some are quite specific in their definitions of what counts as a personal tweet or divide such tweets into multiple categories. For example, Jackson and Lilleker (2011) have categories for tweets giving details of the MP's personal life, those which identify their interests (sport, music etc.) and those which display a sense

of humour, however that level of granularity is not used here. The 'personal' category in this research is closest in intention to Sæbø's (2011) category of non-political content, which he defines as tweets that attempt to allow constituents to get to know the politicians. In his research that category is heavily dominated by discussion of sports and family life. In the sample of tweets coded here, these personal tweets include discussion of sport, music, non-political humour, activities unrelated to the campaign such as visiting a restaurant, and good wishes to other people such as 'happy birthday' or 'happy Easter' tweets.

One of the largest categories here is local tweets – those which make some mention of the MP's constituency or refer in some other way to local matters. The only pre-existing scheme which includes explicit consideration of localism in tweets is that developed by Jackson and Lilleker (2011). A substantial focus of their paper is on how MPs use Twitter as part of their impression management strategies and to present themselves as good constituency MPs and so they have several categories for different types of local tweets. That level of granularity is beyond what's required for this research and would have led to too many small categories with insufficient tweets in them for effective analysis (within the scope of the number of tweets that could realistically be hand-coded within the time available for this research) thus there is a single category for all tweets which include some mention of the MP's constituency or local area.

As this research focuses on the General Election it is unsurprising that there are various different kinds of campaign-related tweets in the data. These divide into three categories: tweets that relate to events that the MP is attending (hustings, campaign meetings, surgeries, business forums and so on); tweets that are about the campaign on the streets (leafleting, going door to door, putting up posters); and tweets that are about meeting people on the campaign trail (tweets that mention a person or group of people that the MP has met that day). Many of the pre-existing coding schema do not include any campaigning categories as they were developed outside of election times. One exception is Graham *et al*.'s (2012) schema which includes categories for updates from the campaign trail, campaign promotion and campaign action. These

three categories could not be accurately distinguished in the data analysed here, which fell more naturally into the three campaign-related categories already discussed.

Most pre-existing schema include a category related to requests for action. Golbeck *et al*. (2010) have one category into which all requests for action are allocated with the exception of fundraising requests which get their own category. This approach is further refined by Hemphill *et al.* (2013) who make a distinction between requests for action that require the recipient to do something meaningful such as vote compared to those that require less significant activity, such as signing a petition or reading something. Graham *et al.* (2013) have a similar category – mobilising and organising – into which they put tweets that make a request for direct action of some kind such as signing a petition or joining a campaign team. In the case of this research fundraising was not deemed to be a useful category as there were only a couple of examples of tweets soliciting funds. Instead, the requests for action were divided into two categories – calls to vote and calls for other actions. Tweets in the voting category were those explicitly directing people to vote. A simple #VoteConservative hashtag (or similar) would not be enough here – the primary focus of the tweet needs to be on exhorting people to vote. The tweets in this category tended to fall into two groups – general exhortations to vote ("the polls are open now – you've got till 10pm to show your support and vote") and requests for support directed to named individuals ("@VoterName, Just wanted to check you have voted today – would really appreciate your support!"). All other requests for action were coded into the broader 'calls to action' category. This included requests to sign petitions, tweets urging people to register to vote, invitations to attend events and to participate in other ways. The requests for funds were included within this category.

Both Hemphill *et al.* (2013) and Graham *et al.* (2013) include a category for thank you tweets and such a category is included here as well. A significant number of MPs use Twitter as a way of acknowledging the support of their campaign staff and volunteers in general or to say thank you to specific named individuals who have supported them in some way. General thanks to voters for their support are also included in this

category. A full list of the codes used in the manual content analysis is given in appendix five.

5.8.6. The coding process

The first stage was to read the tweets to get a feel for which codes from existing schema might apply. Coding then began, either applying the pre-determined codes or developing new ones as appropriate. A second round of coding was completed to refine the codes further, for example by blending categories together and removing codes that did not have sufficient cases. The researcher applied as many codes as were relevant to each tweet, in common with the approach taken by Golbeck *et al.* (2010); Jackson and Lilleker (2011); Lawless (2012); Gainous and Wagner (2014) and Hemphill *et al.* (2013) amongst others, as it was clear that there were cases where an individual tweet could serve more than one purpose. For example, a tweet such as "Happy St Georges Day to everyone celebrating in Southampton today!" would be coded as being both a personal tweet (as were all mentions of public holidays and other similar celebrations) as well as being a local tweet as it includes a mention of the MP's local area.

That said, the sentiment variables are considered to be mutually exclusive (Riffe, Lacy and Fico, 2014) meaning that the same piece of content cannot be considered to be both positive and negative. In this case, tweets are short and so it is assumed that individual tweets will generally be positive, negative or neutral but cannot be more than one of these at the same time. This approach has been taken by other researchers who have assumed that any emotion identified in a tweet will apply to the whole of the tweet (e.g. Pak and Paroubek, 2010). Clearly, a tweet can express both positive and negative sentiment but where that is the case the coder must make a decision regarding which predominates. That is the approach that has been taken here.

Content categories should be exhaustive (Riffe, Lacy and Fico, 2014), meaning every unit of content can be fitted into a relevant category. Here every tweet was able to be manually coded into at least one of the identified categories. Riffe, Lacy and Fico

(2014) guard against the use of an 'other' category but in this case it was unavoidable as there were some tweets which simply could not be fitted into any clearer category (e.g. Welsh MPs tweeting in Welsh, typos and mistakes, tweets that simply include a link with no clue what it is a link to and so on). However, the number of tweets coded as 'other' was kept to a minimum. After the first round of coding all the 'other' tweets were re-examined – in some cases they were added into pre-existing categories upon closer consideration, in others new categories were developed to take them.

5.9. Reliability

The reliability of a piece of research is about the extent to which the findings are repeatable and consistent. A measure is seen as reliable if it gives the same result time and again. In the case of content analysis, reliability is defined as agreement between coders regarding how to categorise items of content (Riffe, Lacy and Fico, 2014). How easy it is to develop reliable coding categories depends largely on the nature of the content being coded. For example, if one is counting the number of times that particular names are mentioned in a newspaper article (manifest content) then one can expect a very high level of intercoder agreement. However, if one's coding scheme requires the coders to make more complex judgements about the nature of a piece of content such as, for example, whether it is intended to be positive or negative, (latent content) then a high level of intercoder agreement can be much harder to achieve (Riffe, Lacy and Fico, 2014).

In content analysis projects there are three possible measures of reliability (Weber, 1990).

- **Stability** – the extent to which the results of the content classification do not vary over time. To determine stability, the same content can be coded more than once by the same coder. If this results in inconsistencies in the coding then that constitutes unreliability. Stability offers the weakest form of reliability.

- **Reproducibility** – also known as inter-coder reliability, reproducibility is the degree to which the same results are produced when material is independently coded by two or more coders. High reproducibility is a better indication of reliability than high stability. High stability indicates that one person's understanding of how to code the data has remained steady over time, whereas high reproducibility shows that two or more coders attribute shared, consistent meanings to the data.

- **Accuracy** – the extent to which the text classification is in line with a norm or standard. Although this offers the highest form of reliability, it is not commonly used by researchers as it requires a standard coding schema to have already been developed for a particular type of content.

In this case a sample of 15% of the coded tweets (196) were coded by a second independent coder. This coder was provided with a copy of the coding scheme and instructed to mark each tweet as either positive, negative or neutral and then to flag each content code as yes or no depending on whether it applied to the tweet in question or not. Generally, a minimum intercoder reliability of 80% is seen as the standard (Riffe, Lacy and Fico, 2014), although lower rates of agreement are sometimes seen as acceptable in exploratory research. In this case, the levels of agreement between the first coder and second coder are shown in Table 5.

Table 5 - Percentage of intercoder agreement by code

Code	Percentage agreed
Call to vote	99.5%
Media response	98%
Other	98%
Achievement	97.4%
Call to action	97.4%
Weather	97.4%
Information giving	96.9%
Attack	96.9%
Support for others	95.9%
Fear appeal	95.4%
Meeting people	95.4%
Event	95.4%
Thanking	94.4%
Personal	94.4%
Position-taking	93.4%
Support for self	93.9%
Local	91.4%
Campaign	90.8%
Sentiment	87.8%

As Table 5 shows, percentage agreement across all codes was substantially higher than the 80% minimum, showing that the coding scheme developed for this research is reliable.

5.10.   Validity

Research validity is about the extent to which the results presented accurately reflect reality (Easterby-Smith *et al.,* 2008). Generally in a content analysis project, if the content codes used are theoretically and conceptually sound and applied reliably, then the results of the study will be valid (Riffe, Lacy and Fico, 2014). Validity is relevant to content analysis in two ways – the extent to which the classification scheme used is valid, and the extent to which the final results are valid (Weber, 1990). If the classification scheme is valid this means that the variables or categories within it do represent the things they claim to represent. If the results of the entire research

project are valid, this means that the findings do not depend on the particular dataset, methods or measurements used in a specific study but can be generalised more widely. Krippendorff (2013 p329) suggests that:

> "A content analysis is valid if the inferences drawn from the available texts withstand the test of independently available evidence, or of being able to inform successful actions."

In this case, one of the purposes of the research is to be able to generate operational guidelines for practice, which politicians can use to improve their performance on Twitter and increase the number of retweets that they get, so according to Krippendorff's definition of validity, the emphasis is on being able to inform successful actions.

Researchers distinguish between internal and external validity. Internal validity is about the extent to which the effects observed in a study can be said to be due to a causal relationship between the independent variables and the dependent variable rather than to some other factor, whilst external validity is about the extent to which the results of the research can be generalised to the wider world.

Regarding internal validity, the aim the research is to be able to make causal inferences about the relationships between variables. In this case retweeting is something that takes place after a tweet has been written and sent, so the retweet can be seen as the dependent variable which is causally influenced by the characteristics of the tweet, its sender and the recipient. This relationship can only work in one direction – it makes no sense to say that the fact of a tweet being retweeted then causes it to have certain characteristics in terms of its content. In this regard then the internal validity of the research is high.

External validity can be enhanced by conducting research on a full data population rather than on a sample of the data, as is the case for much of this research. We can say with confidence that the results of the exploratory analysis can be generalised across MPs' tweets in the run up to the 2015 General Election because they are based

on an analysis of all such tweets. We cannot generalise beyond this and say that the results would apply in other sets of circumstances, such as tweets sent by people who are not MPs, or tweets sent by MPs outside of election times. The CHAID models and the manual content and sentiment analysis are conducted on samples of tweets but care has been taken to ensure that these samples are truly random and representative.

5.11.   Ethical considerations when using Twitter data

Researchers using Twitter data generally do not include any consideration of ethics in their published papers. Of 380 papers reporting Twitter research between 2006 and 2012 only 16 included any mention of ethical issues (Zimmer and Proferes, 2014). Of these, five took steps to ensure the anonymity of tweeters by changing their names whilst the remainder took the view that Twitter represents public data and hence does not present any particular ethical problems. None of the Twitter-based research papers reviewed here included any consideration of ethics and none took any steps to anonymise the politicians about whom they were writing. This is the approach taken in this research too – Twitter is public data and none of the politicians mentioned in this research could have any reasonable expectation of privacy when tweeting. This conclusion was reached after considerable investigation of the ethical issues associated with Twitter research, as will be outlined in this section.

Twitter is an inherently public-facing platform and its default mode is that accounts are public. Fewer than 10% of Twitter users have restricted access to their accounts (Liu, Kliman-Silver and Mislove, 2014) so one may assume that the other 90% are happy for their communications to be publically available. This is particularly the case when it comes to politicians who are explicitly using Twitter as part of their arsenal of campaigning tools. They can have no reasonable expectation that their tweets are private, particularly when one considers the numerous cases of political 'Twitter fails' which make the public nature of Twitter communication only too clear. Twitter itself makes data available to researchers and its terms and conditions include an

agreement to this and explicitly warn that users' public information will be shared with third parties such as universities for the purpose of research (Twitter, 2014).

That said, Twitter does enable the collection of a potentially highly personal data, connected to named individuals. Additionally, the possibility exists for users to retweet protected tweets sent by others (by copying and pasting them into the Twitter tweet composition box) in which case the expectations of privacy of the sender of the original tweet can be violated. It is also possible that an internet 'troll' could set up an account in the name of another person and send postings from that account (Swirsky, Hoop and Labott, 2014), something that public figures such as politicians might be particularly vulnerable to. In the case of very well-known public figures Twitter verifies that the owner of the account is who they say they are. In this case, 246 of the accounts under consideration (67%) are verified so we can be confident that they belong to the people whose names are on them. Additionally, the Twitter handles of all the politicians included in this research were confirmed by the researcher manually so the issue of trolling is not a significant risk.

In the US, federal rules require research to undergo ethical approval whenever human subjects are involved. Human subjects are defined as "living [individuals] about whom an investigator obtains data through interaction with the individual or identifiable private information" (Moreno *et al.,* 2013 p709). Twitter research is not considered to be research on human subjects because the information used is publically accessible on the internet to anyone, is not private and does not require any interaction with the author in order to access it. US regulations also exempt from ethical consideration research which involves the observation of public behaviour, which tweeting is. However, there is an exception to this exemption which is research that enables the identification of named individuals. In the general population of tweeters many are anonymous, using Twitter handles which give no clue as to their real names. However, in the case of politicians a named individual can be identified in every case. This restriction would mean that only general observations could be made about publically available tweets, rather than any analysis of the tweets of any named individual. However, the vast majority of Twitter research considering politicians' tweets

identifies the politicians concerned as named individuals, and in many cases the research would not make sense if this were not the case as knowing who the sender is sheds light on why the tweet may have been received in a particular way.

A parallel could be drawn between politicians' tweets and their other campaign communications such as speeches and flyers, all of which are in the public domain and have been used as the basis of academic research without any ethical concerns being raised. Social media is now another weapon in politicians' communications arsenals and as such is public-facing communication not put out with any meaningful expectation of privacy. Bruns and Burgess (2012) agree that identifying individual tweeters directly in research can be ethically problematic but suggest that an exception can be made for official Twitter accounts. The Twitter accounts of individual politicians in which they identify themselves by their role as an MP could be seen as analogous with official accounts representing organisations of other types, so consideration of individual politicians' behaviour on Twitter is less problematic than identifying other named individuals might be.

Research on Twitter tends to consider one of two main domains – either the sender of the tweet or the message sent (Williams, Terras and Warwick, 2013). In this case the prime unit of analysis is the tweet itself rather than the individual sending the tweet, although the personal characteristics of the individuals are relevant in general terms (gender, political position, party) and this may lead to individuals being personally identifiable (for example, there may only be one female MEP representing a particular party). However, the research does not involve detailed consideration of named individuals and so the ethical implications of this project are minimal.

The ESOMAR guidelines on social media research do not require any anonymisation of social media data where either explicit consent has been obtained or where the use of the data is in line with the terms of use of the network concerned, as is the case for this research (Esomar, 2011). However not all researchers agree that this is the end of the argument. For example, Oboler *et al.* (2012) suggest that the terms and conditions and privacy policies of social networks such as Twitter may themselves be unethical.

Gleibs (2014) also suggests that social network users may *feel* that they are in a private space, even though the network is public, because they use it primarily to communicate with friends and other close associates, and hence may have a strong expectation of privacy. Whilst this may be the case with private individuals, it is not the case for politicians who, for the reasons already discussed, can have no reasonable expectation of privacy when campaigning on Twitter. Should a politician wish to use Twitter in a personal, private capacity they can set up an account using an alias which does not identify them by name, or restrict access to their account. All of the accounts used as the basis for the research reported here are open to all to see and identify the politician concerned as a named individual either in the Twitter handle or in the Twitter biography. Hence, there are no substantive ethical issues associated with using the tweets generated by these accounts as the basis for research.

5.12.   Bias

Bias can creep into research in several ways. Table 6 identifies the most relevant sources of potential bias for this research and discusses the steps taken to limit their effect. A more detailed discussion of the research limitations is given in chapter eight.

Table 6 - Sources of bias and steps taken to minimise it

Source of bias	Steps taken to reduce
Sampling bias	Random samples generated using SPSS Modeler's random sampling function.
Measurement bias	MPs were unaware that their tweets would later be used for this research and so could not have altered their behaviour accordingly. Additionally, this research measures what people actually did (real retweets) rather than relying on self-reports of their behaviour.
Reporting bias	This is exploratory research which does not aim to prove any particular hypotheses regarding retweeting. The researcher has tried to present the research findings in as dispassionately as possible, neither ignoring negative results nor overstating positive ones.
Bias in selection of analytical techniques	Many different statistical and analytical techniques could have been used for this research and other researchers would no doubt have approached this task differently, including different variables in their models or building a different number of models.
Coding bias	A different researcher would almost certainly have come up with a different set of codes. Once the codes are determined the process of applying them to the tweets is also subject to researcher bias as different researchers would apply codes in different ways. However, two coders were used and the level of agreement between them was high, indicating that the coding scheme used was readily understandable by another person.

5.13. Chapter conclusion

The two main methods of this research – manual content analysis and sentiment analysis – do not appear to have been used by researchers aiming to predicting whether tweets will be retweeted, despite extant literature (as discussed in chapters two and three) strongly suggesting that the content of the tweet and its valence are likely to be important factors in determining retweets. Additionally, no extant research appears to have used CHAID modelling for retweet prediction. Thus this research makes two important methods-based contributions to research on retweeting. It integrates manual content and sentiment analysis into a series of predictive models using CHAID, with a view to demonstrating how this approach can effectively be used. The hope is that this approach can then be applied in other contexts. A secondary methods contribution is that this research compares the predictive effectiveness of manual content and sentiment variables and of machine calculated variables, feeding into a growing body of literature on the use of these techniques on Twitter data. The next chapter discusses how these methods were applied to the data collected and presents the findings of the research.

Chapter 6 Findings

6.1. Chapter introduction

The research presented in this thesis is designed to answer the question **what factors influence whether politicians' tweets are retweeted?** The objectives of the research, outlined in chapter one and presented again here for reference, are as follows:

1. To identify the factors that extant literature suggests are most likely to determine whether tweets are retweeted.

2. To identify other factors which might also play a role in influencing whether politicians' tweets are retweeted.

3. To propose a typology of the kinds of tweets sent by UK politicians during the 2015 General Election campaign and identify which work best in terms of stimulating retweets.

4. To test the extent to which the factors identified do indeed determine the chances of a tweet being retweeted by building predictive models using CHAID.

5. To demonstrate a new methodological approach to predicting retweets – CHAID analysis – which could be of use to social media researchers working in other fields or to commercial marketers.

6. To provide practical advice for political parties or individual politicians who wish to use Twitter as part of their campaign communication strategy, regarding how to best harness the power of Twitter to engage with citizens.

Objectives one and two were addressed in the literature discussions that took place in chapters two, three and four of this thesis. The current chapter presents a discussion of research findings and shows how objectives three, four and five have been addressed. Objective six is addressed as part of the recommendations made towards the end of the thesis in chapter eight. This chapter now begins with a discussion of the findings of the descriptive analysis phase of the analysis before moving on in the second half of the chapter to show how CHAID models were used to predict retweets.

6.2. Phase one – descriptive analysis

The aim of the descriptive analysis phase was to build an understanding of the underlying patterns in the MPs' tweeting behaviour in order to determine whether the tweeting behaviour of MPs in this dataset is in line with previous research on politicians' Twitter use, as discussed in chapters two and three. The descriptive analysis also sheds some light on the factors that might be in play when it comes to predicting retweets (research objective two), as well as helping to build a typology of the types of tweets that politicians send (research objective four).

The MPs sent 154,565 tweets during the campaign period of which 42,444 (27%) were original tweets, 32,597 (21%) were replies and 79,524 (51%) were retweets. The analysis presented in this thesis focuses only on the MPs' original tweets. Thus, replies and retweets were removed from the dataset leaving the 42,444 original tweets.

Of these 42,444 original tweets, 85% were retweeted at least once. Research by Stonetemple Consulting finds that only 36% of tweets across the entire Twitter corpus are retweeted (Enge, 2014), thus it would appear that MPs do considerably better than average when it comes to getting retweeted. The mean number of retweets that MPs achieve is 18.46 but a Kolmogorov-Smirnov test shows that distribution is highly skewed and non-normal (D(42,444)=0.449, p <.001, skewness 51.868, kurtosis 3724.323) so the median (four) gives a better sense of the true distribution[18]. Figure 26 shows the skewed distribution. In this stem and leaf plot each leaf represents 64 tweets. Of 42,444 tweets, 31,596 achieve 10 retweets or fewer (75%). The plot in Figure 26 treats anything over 27 retweets as an extreme case, and there are 4,843 of these (just over 11% of the total cases). So, a small number of MPs achieve high numbers of retweets whilst the majority get very few.

---

[18] The Kolmogorov-Smirnov test is used throughout this thesis to test for normality of distribution.

Figure 26 - Stem and leaf plot showing number of retweets

```
Frequency    Stem & Leaf

  6510.00     0 . 00000000000000000000000000000000000000000000000000000000000000000000000000000000000000000000000000000000
  5431.00     1 . 0000000000000000000000000000000000000000000000000000000000000000000000000000000000000
  4490.00     2 . 00000000000000000000000000000000000000000000000000000000000000000000000
  3563.00     3 . 0000000000000000000000000000000000000000000000000000000
  2711.00     4 . 000000000000000000000000000000000000000000
  2244.00     5 . 0000000000000000000000000000000000
  1862.00     6 . 0000000000000000000000000000
  1446.00     7 . 0000000000000000000000
  1300.00     8 . 00000000000000000000
  1132.00     9 . 000000000000000000
   907.00    10 . 00000000000000
   748.00    11 . 000000000000
   701.00    12 . 00000000000
   609.00    13 . 0000000000
   510.00    14 . 00000000
   439.00    15 . 0000000
   368.00    16 . 000000
   376.00    17 . 000000
   328.00    18 . 00000
   291.00    19 . 00000
   294.00    20 . 00000
   265.00    21 . 0000
   225.00    22 . 0000
   217.00    23 . 000
   209.00    24 . 000
   231.00    25 . 0000
   194.00    26 . 000
  4843.00 Extremes    (>=27.0)

Stem width:        1
Each leaf:        64 case(s)
```

The total number of retweets achieved for a single campaign tweet ranges from zero to 13,919[19], the most retweeted tweet being sent by Ed Miliband (Figure 27).

Figure 27 - Most retweeted tweet in General Election campaign

> **Ed Miliband** @Ed_Miliband ☑ +⚇ Follow
>
> David Cameron has decided not to attend tonight's debate. If you're applying for the job of PM, you should turn up to the job interview.

On the surface, who you are looks like the most important factor influencing large retweet volumes. There are 57 tweets in the dataset that achieved over 1,000

---

[19] This figure was correct at the time at which the data were collected. It is unlikely that it will have changed much since then but old tweets can of course still be retweeted so retweet numbers can change over time.

retweets, of which 86% were sent by either David Cameron or Ed Miliband. The remainder were sent by Nick Clegg (four), Diane Abbott (two) and Rachel Reeves and Tom Watson who manage one each. No MPs with low public profiles achieve significant volumes of retweets.

In order to enable comparison across MPs with varying tweet volumes, the number of retweets per campaign tweet was calculated for each MP. This ranges from zero through to 818 with a mean of 7.56 retweets per tweet. This distribution is also highly skewed due to a few MPs who achieve very high numbers of retweets per tweet (D(366)=0.438, $p < .001$, skewness 14.675, kurtosis 231.305) as shown by the fact that the median number of retweets per tweet is 1.86. Ed Miliband is the highest achieving MP, with 818 retweets per tweet, followed by David Cameron with 416 and Nick Clegg with 106. Everyone else has less than 50 retweets per tweet and six MPs did not generate any retweets.

This research focuses on how the characteristics of the sender and the characteristics of the tweet itself influence retweeting. In this descriptive phase of the analysis the sender characteristics and tweet characteristics will be considered in turn.

6.3.   Descriptive analysis of senders' characteristics

Sender characteristics can be further sub-divided into three groups – characteristics relating to the sender's Twitter presence, characteristics related to their political situation and their personal characteristics. Examples of each type of characteristic are given in the conceptual model (shown again in Figure 28) and each will now be discussed in more depth, along with a consideration of how they influence whether MPs' tweets are retweeted.

Figure 28 - Conceptual model highlighting this phase of the analysis
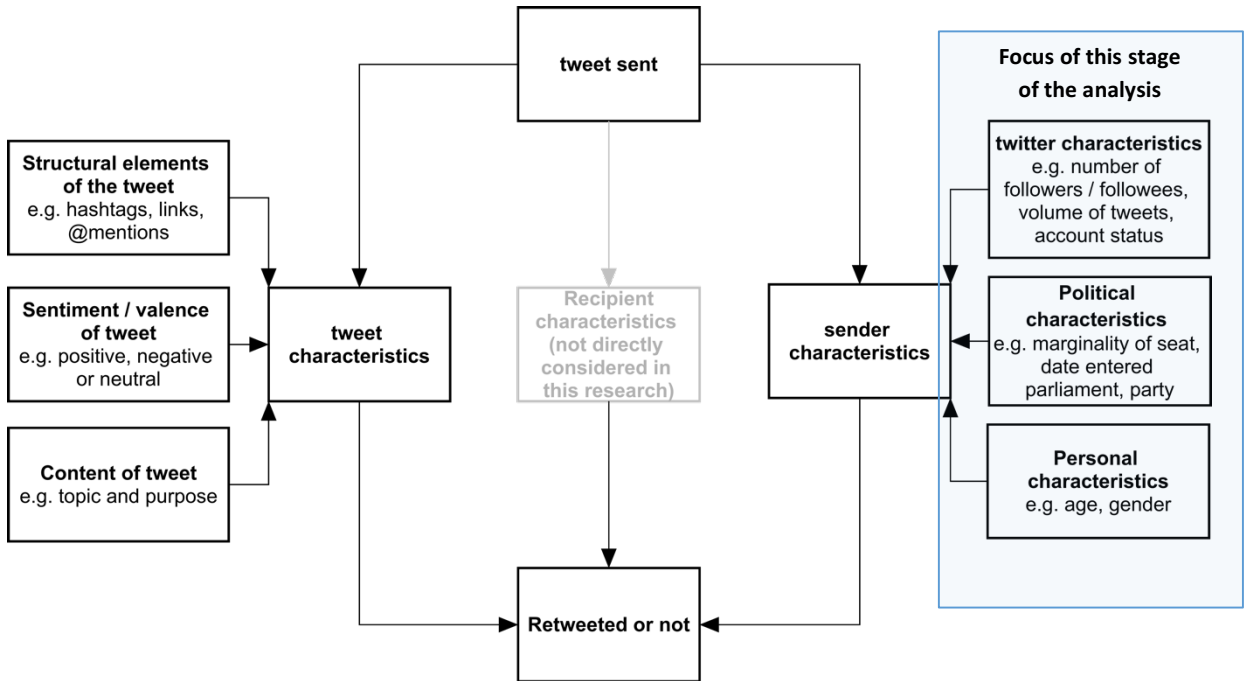
6.3.1. Senders' Twitter characteristics

This stage of the analysis considers variables that relate to each MP's Twitter profile or behaviour to understand more about the patterns of behaviour in this dataset in order to see how they compare to patterns identified in previous research. Some elements of the descriptive analysis can also be used to discover something about how some variables are related to retweeting. The key variables of interest here are the number of followers that the MPs have, the number of people they follow ('followees'), their ratio of followers to followees, the total number of tweets they have sent during their time on Twitter, the total number of tweets that they have sent during the campaign and whether their Twitter account is verified.

6.3.1.1. Number of tweets sent

The number of original tweets sent by individual politicians during the campaign ranges from 1 to 4,108 with a mean of 422.32 and a median of 245. The volume of tweets has increased very substantially since Graham *et al*'s (2013) examination of the 2010 General Election campaign during which the Labour Party averaged 62 tweets

per candidate (across all candidates – incumbents and challengers) compared to 44 per candidate for the Conservatives.

As can be seen in Figure 29, the pattern of MPs' campaign tweet volumes follows a long-tailed non-normal distribution (D(366)=0.223, p <0.001), as is common with many measures of Twitter activity, with the majority of MPs sending very few tweets whilst a small number are extremely active. The mean campaign tweets per MP is 422.32 however the data are not normally distributed and the mean is pulled upwards by a small number of outliers. The median is 245 and gives a truer picture of the distribution. The distribution has positive skewness of 3.334 indicating that it is skewed with a long tail to the right, and kurtosis of 15.262 indicating that the tail of the distribution is heavier than for a normal distribution. Both the skewness and kurtosis are accounted for by the small number of extremely active MPs at the top end of the scale, balanced against a much larger number of relatively inactive MPs.

Figure 29 - Distribution of campaign tweeting activity
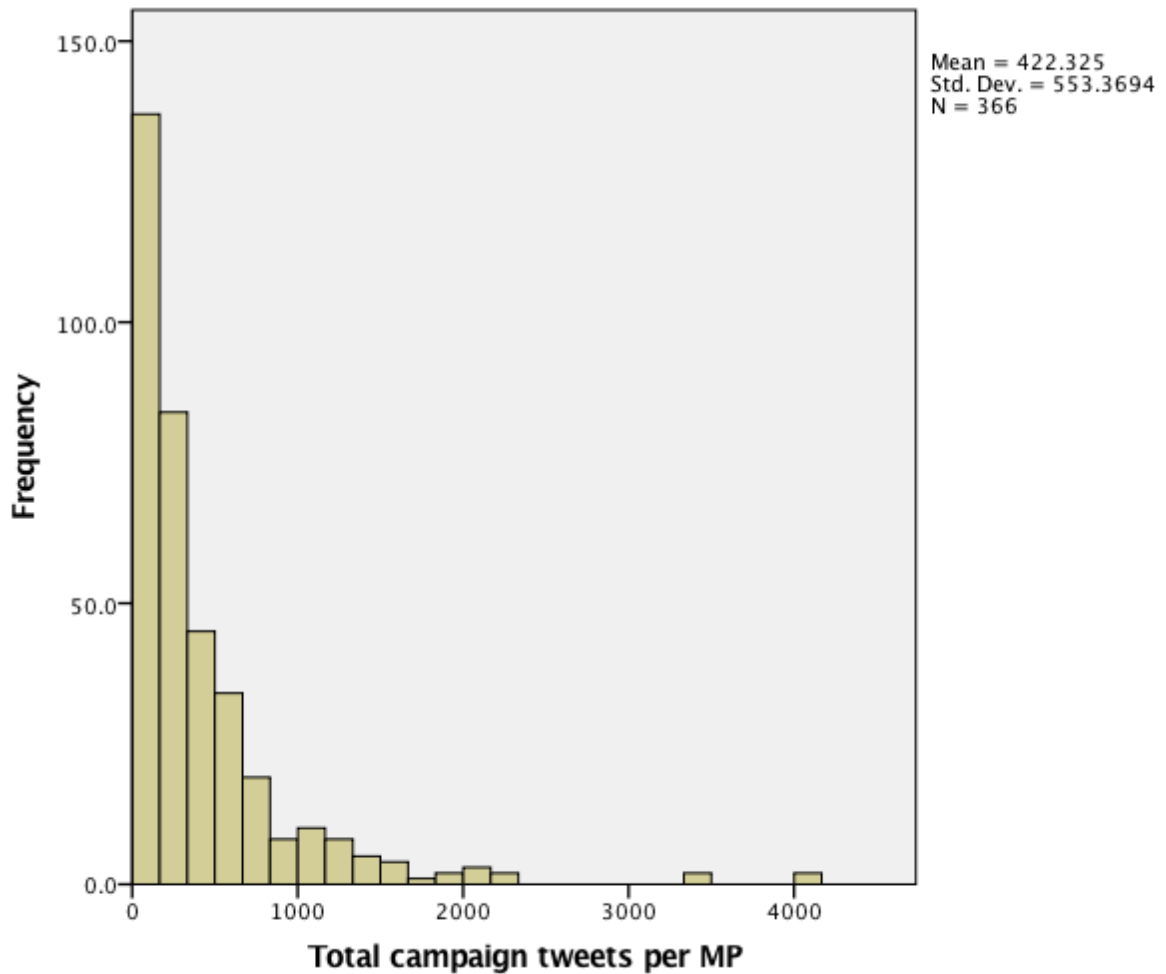
There are 13 MPs in the dataset who sent fewer than ten tweets during the entire campaign, of whom seven sent only one (thus one is the mode for this distribution). In comparison, at the top end of the scale the most active MPs sent thousands of tweets, as can be seen in Table 7 showing the ten MPs who sent the most tweets during the campaign. With the exception of perhaps Douglas Carswell, these are not particularly well-known MPs, suggesting that the Twitter activity of an MP is not necessarily related to their prominence.

Table 7 - Most active MPs on Twitter during campaign

MP	Party	Campaign tweets
Karl Turner	Labour	4,108
Jonathan Edwards	Plaid Cymru	4,048
Greg Mulholland	Liberal Democrat	3,494
Andrew Gwynne	Labour	3,491
Julian Huppert	Liberal Democrat	2,297
Michael Dugher	Labour	2,248
Mike Gapes	Labour	2,139
Douglas Carswell	Ukip	2,031
Karl McCartney	Conservative	2,021

A new variable was calculated for each MP, giving the percentage of their campaign tweets that were retweeted at least once. The aim here was to provide some measure of each MP's success at being retweeted that did not consider the number of retweets but simply whether or not a tweet was retweeted. This percentage ranged from zero to 100% with a mean of 83.39 and a median of 87.5, showing that most MPs managed to get most of their tweets retweeted at least once. Once again this distribution is not normal ($D(366) = 0.152$ $p < .001$, skewness = -1.577, kurtosis = 3.698).

To test for relationships between volume of Twitter activity and the number of retweets each MP generated per tweet, a series of correlations were run. As the data are not normally distributed Spearman's correlation coefficient is the most appropriate test to use (Field, 2009). These correlations show that:

- There is no significant relationship between the total number of tweets sent by each MP during their time on Twitter and the number of retweets per campaign tweet they achieve ($r_s = .02$ (one-tailed), $p < .307$)
- There is no significant relationship between the number of tweets sent by each MP during the campaign and the number of retweets per campaign tweet they achieve ($r_s = .01$ (one-tailed), $p < .407$)

6.3.1.2. Number of followers and people followed

When looking at the MPs' follower numbers, similar long-tailed non-normal distributions are observed (number of followers $D(366) = 0.897$, $p < .001$, skewness 12.174, kurtosis 176.113). The mean number of followers is 18,949 with a range from 1,204 (John Howell, Conservative MP) to 1,085,183 (David Cameron), and a standard deviation of 67,818.380. Once again, the mean is pushed upwards by outliers at the top of the distribution with huge numbers of followers relative to typical MPs. The median number of followers is substantially lower at 7,919. These numbers show that MPs' Twitter use (and indeed Twitter use generally) has changed very substantially since Jackson and Lilleker's (2011) examination of Twitter use during 2009 when Tom Watson was the MP with the most followers with 4,441, a number that would put an MP firmly in the bottom half of the distribution today.

The number of people each MP follows ranges from 13 (Mark Field, Conservative MP) to 18,078 (Steve Reed, Labour MP), with a mean of 1,496, a median of 766 and a standard deviation of 2,077.018. This distribution is also heavily skewed and non-normal ($D(366)=0.238$, $p < .001$, skewness 3.752, kurtosis 19.461). Again, MPs are more active now than they were in 2009 when the median number of people they followed was 133 (Jackson and Lilleker, 2011). There is a wider distribution of follower numbers than following numbers because of the substantial number of individuals within the dataset with huge numbers of Twitter followers. It is generally the case that the higher the number of Twitter followers someone has, the less likely they are to follow them all back, so at the top end of the scale follower numbers will be higher than following numbers (Cha *et al.*, 2012).

A measure of how interactive MPs are on Twitter can be gained by looking at their ratio of followers to followees (Jackson and Lilleker, 2011). A ratio of 1 would indicate that an MP follows exactly as many people as they are followed by. In this dataset the ratio ranges from 0.827 (Conservative MP Ann McIntosh, who follows 6,432 people whilst being followed by 5,320) to 2,825.997 (David Cameron, who follows 384 people whilst being followed by 1,085,183).

Further Spearman's rank correlations show that:

- There is a significant relationship between the number of followers an MP has and the number of retweets per campaign tweet they achieve (r_s = .47, p (one-tailed) <.001).

- There is a significant relationship between the number of people an MP follows and the number of retweets per campaign tweet they achieve (r_s = .14 (one-tailed), p <.003)

- There is a significant relationship between the MPs' ratio of followers to followees and the number of retweets per campaign tweet they achieve (r_s = .21 (one-tailed), p <.001)

This suggests that number of followers one has and the number of people whom one follows both play a greater role in determining retweets than does the volume of tweets that one sends.

6.3.1.3.   Account verification

Just over two thirds of the MPs in this dataset have had their accounts officially verified by Twitter (67%). Whether or not a tweet is retweeted is related to whether or not the sender's account is verified (Table 8). Just over 86% of tweets from verified accounts are retweeted, compared to just over 80% for those from non-verified accounts, and this difference is significant (p < .001). This is likely to be because verification is an indication of an MP's status. It is not the fact that it comes from a verified account itself that leads to a tweet being retweeted. More likely, it is because those MPs who have verified accounts are the highest profile MPs, thus they have more followers and hence more people to retweet their tweets, as well as being more likely to be retweeted by automated Twitter bots[20].

---

[20] Twitter bots (short for robots) are accounts that have been set up purely to automatically retweet particular groups of tweets (for example all tweets using a particular hashtag, mentioning a certain word or, as in this case, sent by a particular group of people). Once the bot is up and running there is no human intervention required – all tweets that meet the bot's criteria are automatically retweeted. The more well-known the MP is, the more likely their tweets are to be picked up by bots.

Table 8 - Influence of account status on retweets

| | | | Retweeted? | | |
			No	Yes	Total
Twitter Verified	No	Count	2237	9299	11536
		% within Twitter Verified	19.4%	80.6%	100.0%
		Standardized Residual	11.1	−4.7	
	Yes	Count	4273	26635	30908
		% within Twitter Verified	13.8%	86.2%	100.0%
		Standardized Residual	−6.8	2.9	
Total		Count	6510	35934	42444
		% within Twitter Verified	15.3%	84.7%	100.0%

Chi-Square Tests

	Value	df	Asymptotic Significance (2-sided)	Exact Sig. (2-sided)	Exact Sig. (1-sided)
Pearson Chi-Square	200.462[a]	1	.000		
Continuity Correction[b]	200.033	1	.000		
Likelihood Ratio	192.732	1	.000		
Fisher's Exact Test				.000	.000
N of Valid Cases	42444				

a. 0 cells (0.0%) have expected count less than 5. The minimum expected count is 1769.38.

b. Computed only for a 2x2 table

However, the relationship does not extend to the number of retweets. Whilst there is a large difference between the mean number of retweets per tweet between verified accounts (9.75) and non-verified accounts (3.07), the difference between the medians is much less pronounced (1.9 for verified accounts, 1.7 for non-verified). An independent samples median test shows that this difference is not significant (Figure 30).

Figure 30 - Comparison of retweets per tweet for verified/non verified accounts
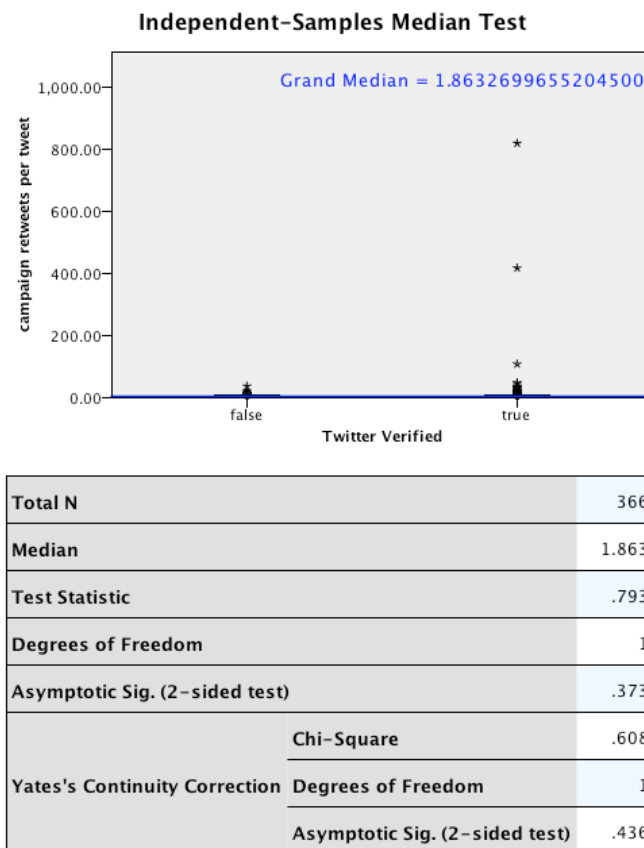
Independent-Samples Median Test

Grand Median = 1.8632699655204500

Total N	366
Median	1.863
Test Statistic	.793
Degrees of Freedom	1
Asymptotic Sig. (2-sided test)	.373
Yates's Continuity Correction Chi-Square	.608
Degrees of Freedom	1
Asymptotic Sig. (2-sided test)	.436

1. Multiple comparisons are not performed because the overall test does not show significant differences across samples.

6.3.2.   Senders' personal characteristics

6.3.2.1.   Gender

The majority of the MPs in the dataset are male – 265 (72%) compared to 101 (28%) female. In 2015 female MPs made up 22% of the total in parliament, so these numbers suggest female MPs were slightly more likely to be active on Twitter during the 2015 election campaign than their numbers in parliament would indicate. This is probably a reflection of the number of Labour MPs in the dataset as Labour accounts for the majority of female MPs.

The sender's gender has a small but significant influence on the chances of their tweets being retweeted. As Table 9 shows, 85% of male MPs' tweets were retweeted

compared to 83% of women's tweets. The chi-square test shows that this difference is significant, with the standardised residuals showing that tweets sent by women are particularly overrepresented in the 'not retweeted' category.

Table 9 - Influence of gender on retweets

			Tweet retweeted?		Total
			No	Yes	
gender	female	Count	1880	9144	11024
		% within gender	17.1%	82.9%	100.0%
		Standardized Residual	4.6	−2.0	
	male	Count	4630	26790	31420
		% within gender	14.7%	85.3%	100.0%
		Standardized Residual	−2.7	1.2	
Total		Count	6510	35934	42444
		% within gender	15.3%	84.7%	100.0%

Chi-Square Tests

	Value	df	Asymptotic Significance (2-sided)	Exact Sig. (2-sided)	Exact Sig. (1-sided)
Pearson Chi-Square	33.764[a]	1	.000		
Continuity Correction[b]	33.586	1	.000		
Likelihood Ratio	33.138	1	.000		
Fisher's Exact Test				.000	.000
N of Valid Cases	42444				

a. 0 cells (0.0%) have expected count less than 5. The minimum expected count is 1690.85.

b. Computed only for a 2x2 table

A comparison of the number of retweets per tweet that the MPs achieve initially suggests a substantial difference by gender, with men generating a mean of 8.92 retweets per tweet compared to just 3.99 for women. However, the skewed nature of this data means it is more appropriate to compare the medians rather than the means. Here the difference is negligible. Men have a median of 1.85 retweets per tweet compared to 1.95 for the women. An independent samples median test (Figure 31) shows that this difference is not significant.

Figure 31 - Comparing median retweets per tweet by gender
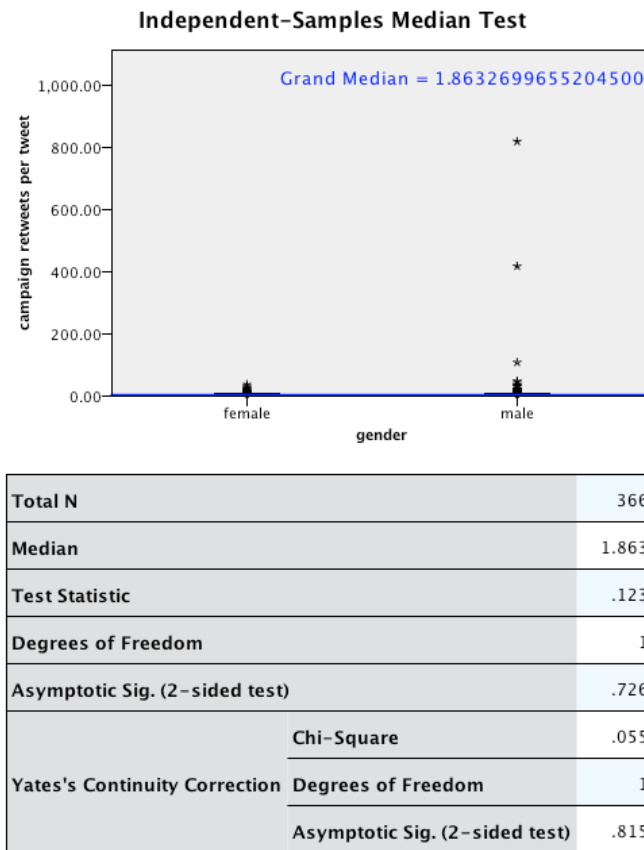
Independent-Samples Median Test

Grand Median = 1.8632699655204500

Total N		366
Median		1.863
Test Statistic		.123
Degrees of Freedom		1
Asymptotic Sig. (2-sided test)		.726
	Chi-Square	.055
Yates's Continuity Correction	Degrees of Freedom	1
	Asymptotic Sig. (2-sided test)	.815

1. Multiple comparisons are not performed because the overall test does not show significant differences across samples.

So, there appears to be a small relationship between gender and whether one's tweets get retweeted or not, but no significant relationship between gender and the median number of retweets achieved per tweet. This is most likely because the highest profile politicians are the most retweeted, and they tend to be men. Here gender is probably a proxy for some combination of seniority and public profile rather than a direct influence on retweets itself.

6.3.2.2.   Age

The average age of the MPs in the dataset is 52 (mean 52.07, median 52), with a range from 30 to 81. The average age of all the elected MPs in 2010 was 50 (www.parliament.uk, n.d.), suggesting that there is no significant relationship

between the age of the MPs and how likely they are to be on Twitter. For ease of analysis, a new variable of 'age group' was created and the MPs re-categorised according to which group they fell into. Table 10 shows that the largest group of MPs is those aged 45-54, followed by 55-64.

Table 10 - MPs by age group

		Frequency	Percent
Age group	25-34	9	2.5
	35-44	80	21.9
	45-54	129	35.2
	55-64	104	28.4
	65+	44	12.0
	Total	366	100.0

Figure 32 shows the median number of campaign tweets sent by each age group, to determine whether there are any significant differences in Twitter activity according to age. An independent samples median test shows these differences to be significant (test statistic 24.332, df 4, $p < .001$), suggesting that MPs in the group 35-44 were significantly more active on Twitter than those in other groups. However, the same test was run to assess whether the median number of retweets per tweet generated was influenced by age group. This test was not significant (test statistic 7.804, df 4, $p = .099$). This shows that whilst how active MPs are on Twitter may be influenced by their age, how successful they are at getting their tweets retweeted is not.

Figure 32 - Median campaign tweets by age group
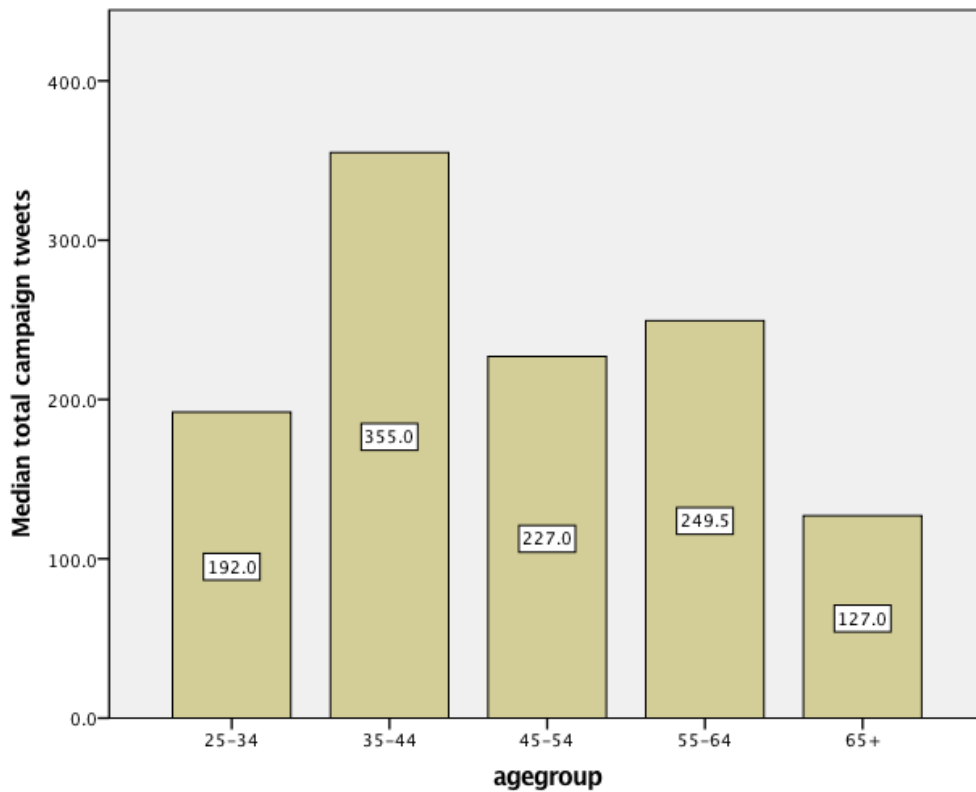
6.3.3. Senders' political characteristics

6.3.3.1. Party affiliation

Table 11 shows the number of MPs from each party active on Twitter during the campaign and the percentage of the total tweeting MPs that came from each party. This shows that Labour politicians punched above their weight on Twitter in the run up to the 2015 General Election. Their 157 active tweeters represent 43% of the tweeting politicians whilst only comprising 39% of MPs in parliament at the time (of the remainder, 46% of MPs were Conservatives and 9% were Liberal Democrats). Just over half of Conservative MPs were active on Twitter compared to 61% of Labour MPs. Smaller parties did better still. Almost 70% of Liberal Democrats tweeted and several small parties had 100% of their MPs represented on Twitter (although in most cases this represents a single person).

Table 11 - Tweeting MPs by party

	Number of tweeting MPs from each party	% of tweeting MPs from each party	% of parliamentary party tweeting
Labour	157	42.9%	61%
Conservative	155	42.35%	51%
Liberal Democrat	38	10.38%	69%
Sinn Fein	3	0.82%	60%
SNP	3	0.82%	50%
Ukip	2	0.54%	100%
DUP	2	0.54%	25%
SDLP	2	0.54%	66%
Green	1	0.27%	100%
Respect	1	0.27%	100%
Plaid Cymru	1	0.27%	33%
Alliance	1	0.27%	100%

Of course, this table only shows how many MPs from each party were active on Twitter – it does not reveal anything about the extent of that activity. 'Active' means sending at least one tweet, so a party could have many MPs active on Twitter whilst still only generating a very small number of tweets. Figure 33 takes this analysis further and shows the numbers of actual tweets sent during the campaign by party. Once again Labour perform well, generating 49% of all the tweets sent during the campaign (compared to 32% for their nearest rivals, the Conservatives). In total Labour generated 75,270 tweets, over 50% more than the Conservatives, but from only two more tweeting MPs.

Figure 33 - Total number of campaign tweets sent by party
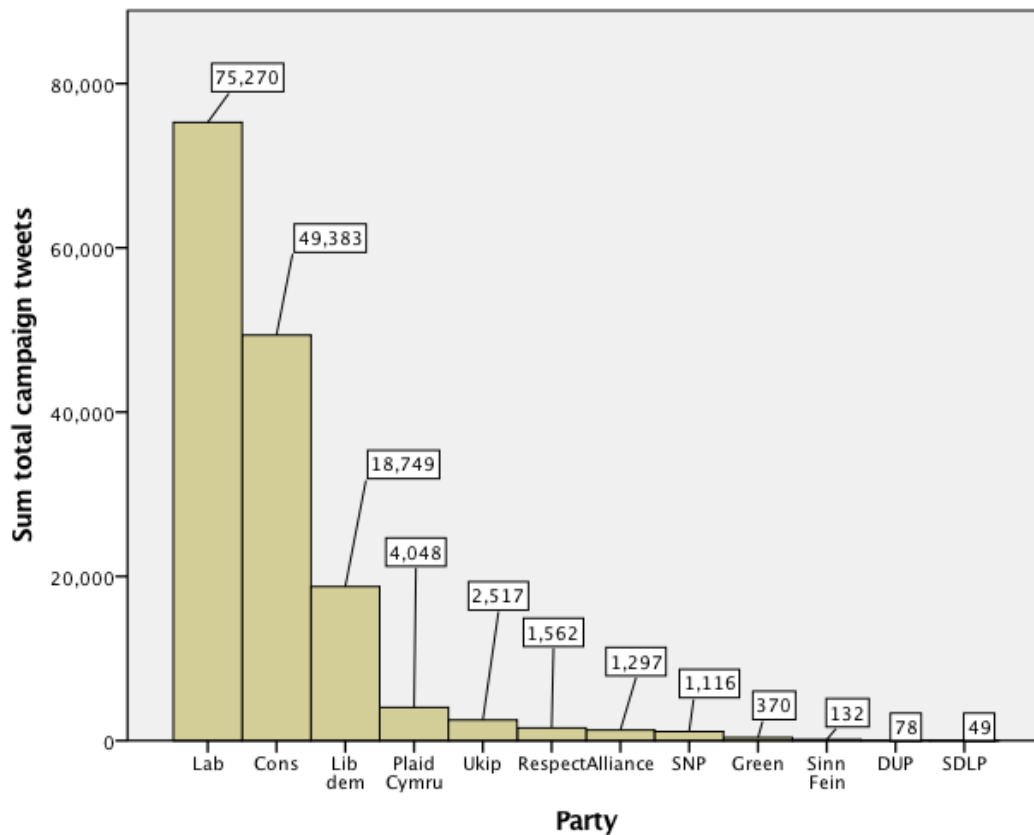
However, Labour do not claim the top spot in terms of number of campaign tweets per MP (Figure 34). Here the smaller parties do much better, and in most cases this is down to one or two individuals being extremely active. For example, Plaid Cymru tops the table with 4,048 tweets during the campaign, all sent by one MP, whereas the Conservatives have a median of 186 tweets per MP spread across 155 tweeting MPs

Figure 34 - Median campaign tweets per tweeting MP by party
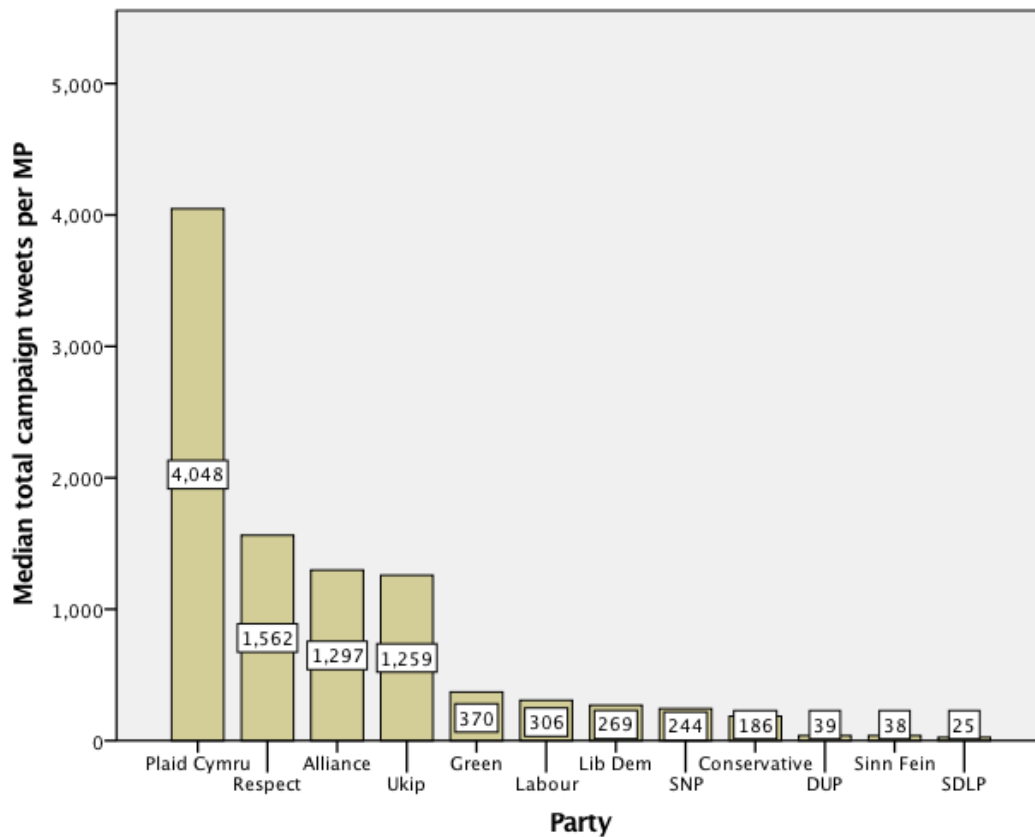
As discussed in chapter two, extant research suggests that MPs tend to operate in broadcast mode on Twitter, making limited use of interactive functions such as retweets and replies. Breaking down these numbers by party (Table 12) reveals substantial variation, with some parties generating relatively few original posts and relying heavily on retweets whilst others focused more on posts.

Table 12 - Types of tweets posted per party

	Frequency and percentage of tweets per party			
	Normal post	@ replies	Retweets	Total
Labour	20,298 26.7%	14,503 19.3%	40,669 54%	75,270
Conservative	16,843 34.1%	9,417 19.1%	23,117 46.8%	49,377
Liberal Democrat	3,647 19.5%	6,212 33.1%	8,890 47.4%	18,749
Sinn Fein	36 27.3%	3 2.3%	93 70.5%	132
SNP	577 51.7%	247 22.1%	292 26.2%	1,116
Ukip	555 22.1%	750 29.8%	1,212 48.2%	2,517
DUP	15 19.2%	14 17.9%	49 62.8%	78
SDLP	14 28.6%	4 8.2%	31 63.3%	49
Green	113 30.5%	55 15%	202 54.6%	370
Respect	213 13.7%	261 16.7%	1,088 69.6%	1,562
Plaid Cymru	282 7%	199 4.9%	3,567 88.1%	4,048
Alliance	51 3.9%	932 72%	314 24.2%	1,297
Total	42,444 27.5%	32,597 21.1%	79,524 51.5%	

Chi-square 8101.558, df 22, *p* <.001

As Figure 35 and Figure 36 show, these differences are particularly pronounced when comparing the tweeting patterns of the smaller parties to those of the three largest parties. The three major parties (Figure 35) all hover around the 50% mark for number of retweets, but show more variation in terms of the percentage of original posts and replies that they generate. The Liberal Democrats appear to be the most interactive of the three major parties, with 33% of their tweets being replies to other people. This falls to 19% for Labour and the Conservatives.

Figure 35 - Types of tweets by party (three largest parties)
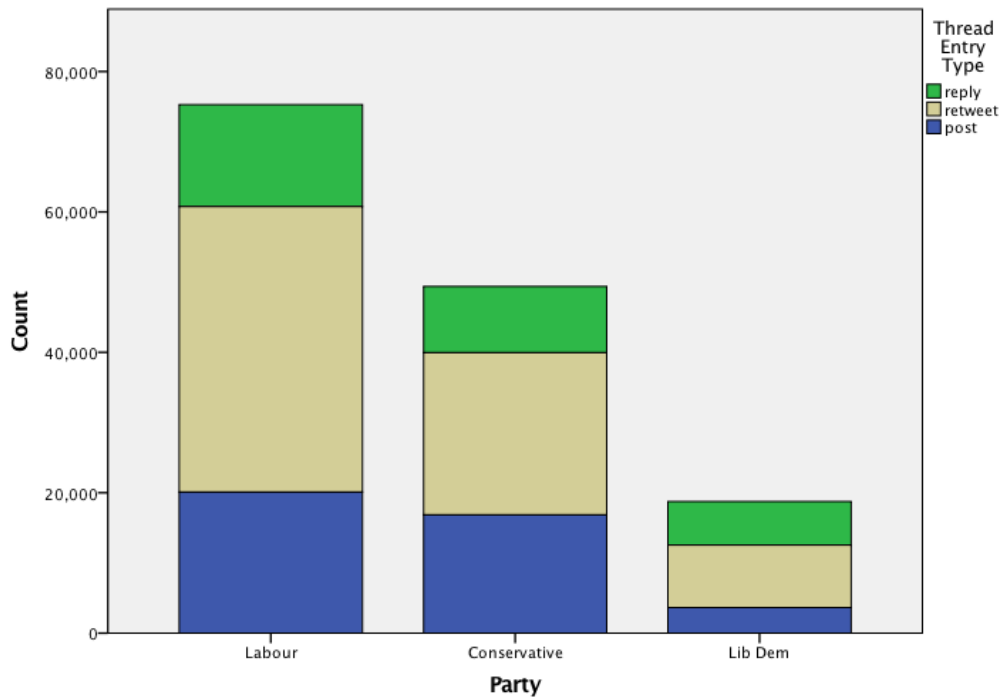
When one looks at the pattern of the smaller parties the differences are more pronounced (Figure 36) and again, tend to come down to differences in the Twitter habits of a few individuals. For example, the Respect Party (consisting of one person – George Galloway) generated only 213 original tweets out of 1,562 total tweets. Almost 70% of Respect's tweets were retweets. The only party with a higher percentage of retweets is Sinn Fein, with 70% retweets out of a total of 132 tweets. The highest percentage of original posts was generated by the SNP with 52%, and the highest absolute number is the Labour Party with 20,298 original posts.

Figure 36 - Types of tweet by party (smaller parties)
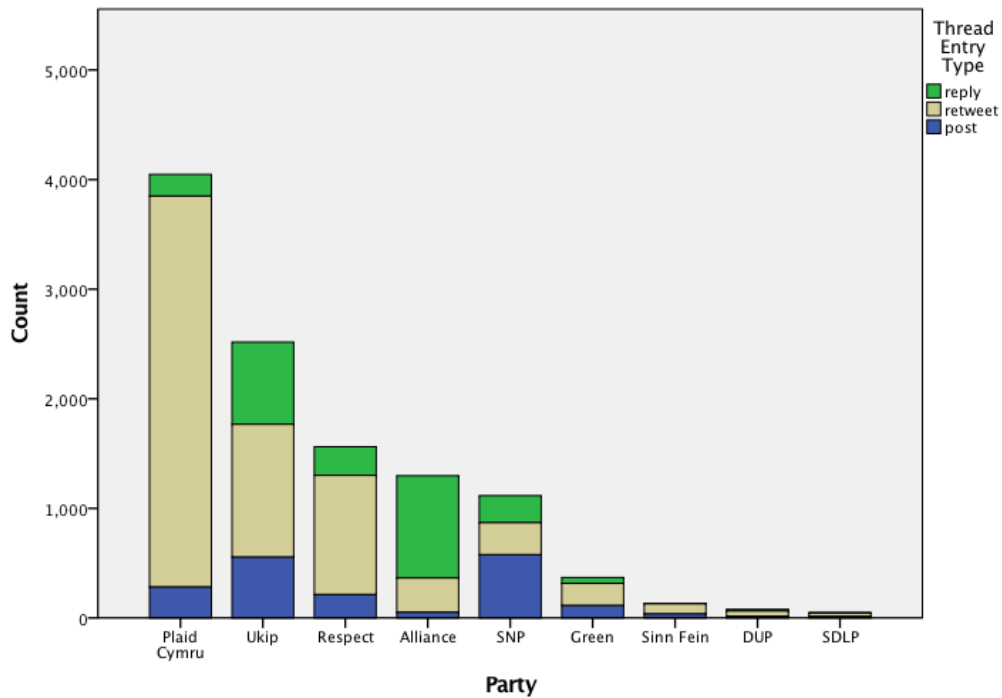
The sender's party affiliation also has a significant impact on the chances of their original tweets being retweeted. The percentage of tweets getting retweeted varies substantially, from Respect, Ukip and the Green Party, all of whom manage to get almost 100% of their tweets retweeted at least once, to the Alliance and the DUP who manage to get only 47% and 40% retweeted respectively (Table 13). The three largest parties do reasonably well here too. Again, Labour performs the best with just over 87% of its tweets being retweeted, compared to 82% for the Conservatives and 80% for the Liberal Democrats.

Table 13 - Percentage of tweets that are retweeted or not according to party

			Retweeted yes/no		
			No	Yes	Total
Party	Alliance	Count	27	24	51
		% within party	52.9%	47.1%	100.0%
	Conservative	Count	3077	13766	16843
		% within party	18.3%	81.7%	100.0%
	DUP	Count	9	6	15
		% within party	60.0%	40.0%	100.0%
	Green	Count	5	108	113
		% within party	4.4%	95.6%	100.0%
	Labour	Count	2567	17531	20098
		% within party	12.8%	87.2%	100.0%
	Liberal Democrat	Count	712	2935	3647
		% within party	19.5%	80.5%	100.0%
	Plaid Cymru	Count	61	221	282
		% within party	21.6%	78.4%	100.0%
	Respect	Count	3	210	213
		% within party	1.4%	98.6%	100.0%
	SDLP	Count	2	12	14
		% within party	14.3%	85.7%	100.0%
	Sinn Fein	Count	13	23	36
		% within party	36.1%	63.9%	100.0%
	SNP	Count	26	551	577
		% within party	4.5%	95.5%	100.0%
	Ukip	Count	8	547	555
		% within party	1.4%	98.6%	100.0%
Total		Count	6510	35934	42444
		% within party	15.3%	84.7%	100.0%

Chi-square Tests

	Value	df	Asymptotic significance (2-sided)
Pearson chi-square	538.488[a]	11	.000
Likelihood ratio	591.277	11	.000
N of valid cases	42444		

a. 2 cells (8.3%) have expected count less than 5. The minimum expected count is 2.15.

This is likely to be a function of how organised the party's campaigning infrastructure is as well as of how many followers its politicians have. For example, all Green Party tweets are sent by one person – Caroline Lucas – who is well-known and has 138,455 Twitter followers (at the time of the election), particularly when compared to Naomi Long, the single representative of the Alliance, who has only 14,102 followers. The difference in retweet performance between large parties and smaller parties can more clearly be seen in Figure 37 and Figure 38.

Figure 37 - Proportion of tweets that are retweeted for larger parties
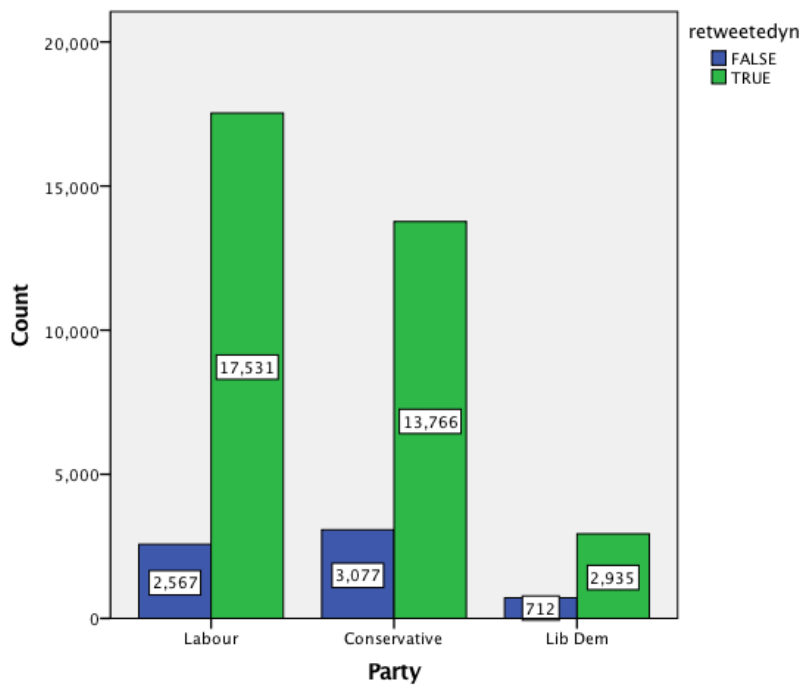
Figure 38 - Proportion of tweets that are retweeted for smaller parties
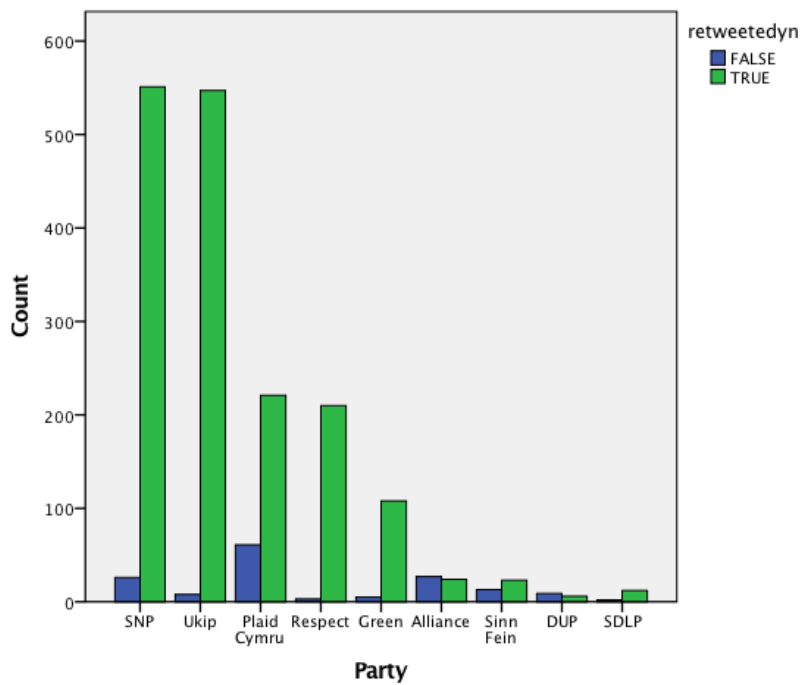
Some parties are also more effective than others, not just at getting their tweets retweeted but also at generating a substantial volume of retweets. For example, the SNP manage to get a median of almost 30 retweets per tweet compared to less than one for the Alliance and Plaid Cymru (Figure 39). Here the three main parties perform poorly when compared to the smaller parties, the Liberal Democrats in particular managing a median of only just over one retweet per tweet. Again, it is likely that these differences come down to differences in the social media strategies of the various parties and the focus that some put on the importance of getting retweets and on politicians retweeting each other's tweets.

Figure 39 - Median retweets per campaign tweet by party
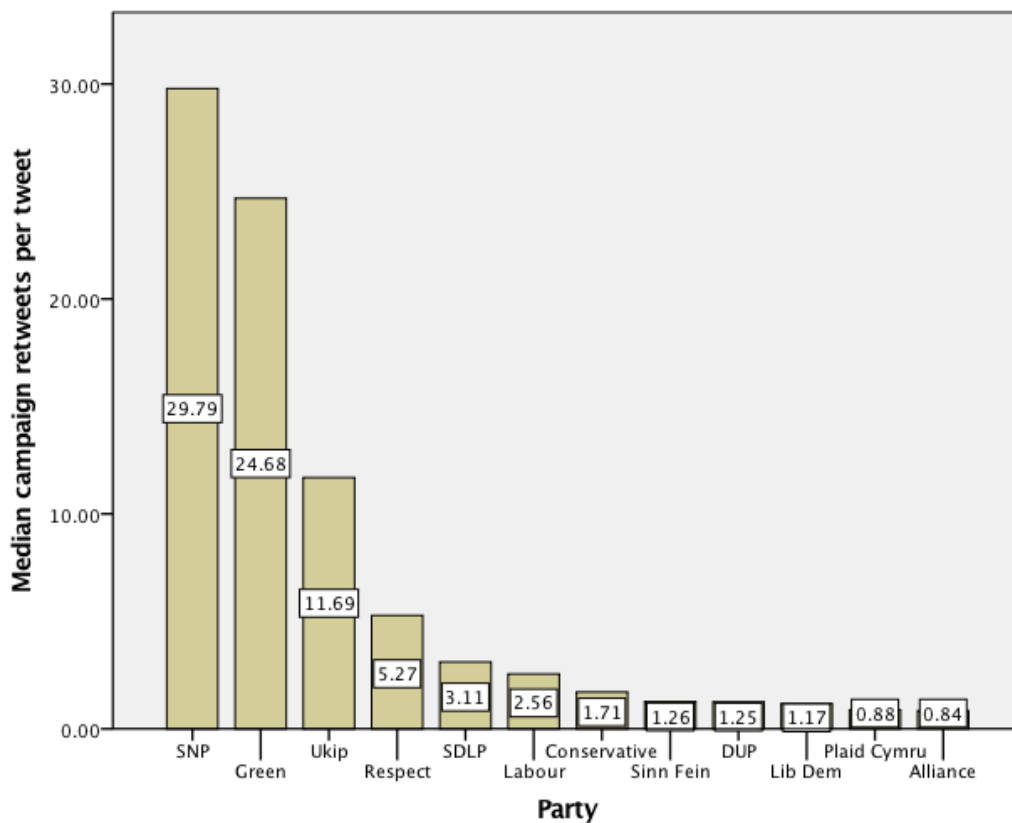
Of the tweeting MPs, 81% held their seats, 13% lost and 6% stood down. Overall, across all MPs in the election 72% won, 14% lost and 14% stood down. These figures suggest that there is no significant relationship between being on Twitter and winning or losing but that, perhaps unsurprisingly, MPs who knew that they were standing down were much less likely to tweet than others. Table 14 supports this contention, showing that just over 82% of the tweets were sent by people who held their seats compared to 15% sent by people who went on to lose and less than 3% sent by people who stood down.

Table 14 - Tweeting behaviour compared to election outcome

	% of tweets	% of MPs in dataset	% of total MPs
Won	82.38%	81.15%	72.01%
Lost	15.02%	12.84%	14.15%
Stood down	2.6%	6.01%	13.84%

Whilst simply being on Twitter is not significantly related to winning or losing, the volume of one's Twitter activity is. Those MPs who held their seats sent a median of 245 tweets during the campaign, compared to 279 for those who lost their seats and 123.5 for those who stood down. An independent samples median test reveals that these differences are significant (test statistic 6.261, df 2, p = .044).

MPs who held their seats achieved a median of 2.04 retweets per campaign tweet, compared to 1.4 for those who lost their seats and 1.11 for those who stood down. An independent samples median test shows these differences to be significant (test statistic 9.922, df 2, p = .007). Of course, these figures could be skewed by a small number of tweets that achieved a massive number of retweets, all of which were sent by very well known politicians who held their seats.

Each MP is flagged according to whether they held their seat, lost it or stood down at the election. The closeness of each contest is determined by the size of their majority and divided into three possible options – safe seats, near-marginal seats and marginal

seats[21]. Regarding closeness of the election, one might expect that MPs in marginal seats would be more motivated to campaign than would MPs in safe seats.

In terms of marginality of seat, 62% (226) of the tweeting MPs were from safe seats, 22% (81) from near-marginal seats and 16% (59) from marginal seats. One might expect that the MPs in marginal seats would be more active on Twitter than those in safe seats as they would have more incentive to campaign actively. However, that does not appear to be the case. Those in marginal seats generated a median of 240 campaign tweets, compared to 352 for near-marginal MPs and 213 for those in the safe seats. An independent samples median test shows these differences are significant (test statistic 10.714, df 2, p = .005). However, this could be because some of the most active tweeting MPs happen to be in near-marginal seats. Whilst the safeness of the seat is related to how active the MPs are on Twitter, there is no significant relationship between the safeness of the seat and the median number of retweets generated per campaign tweet (test statistic 2.514, df 2, p = .285).

6.3.3.3.   Parliamentary cohort

Nearly half of the MPs in the dataset, 42% (154), entered parliament for the first time in 2010 (Figure 40). In contrast, only 35% (227) of all the MPs in parliament at the time were first elected in 2010 (Devlin *et al.*, 2015). This suggests that MPs from the 2010 cohort are more likely to be active on Twitter than their absolute numbers in parliament would indicate.

---

[21] Based on the approach taken by Jackson and Lilleker (2011) who categorised safe seat as those with a margin of 11% or more over the next nearest candidate, near-marginal seats as those with a majority of 5.1% to 10.9% and marginal seats any with a majority of 5% or less.

Figure 40 - Year in which the tweeting MPs entered parliament
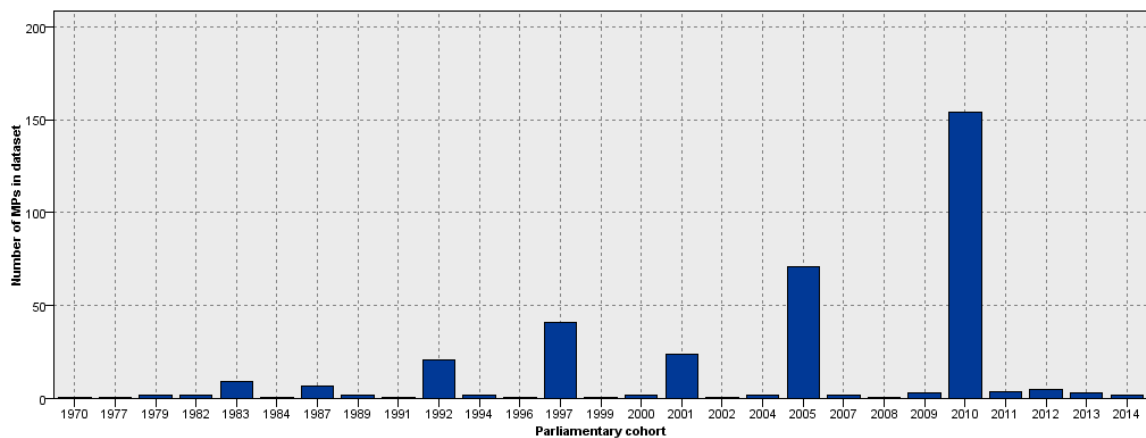
It could be posited that MPs elected more recently in a post-Twitter age might be more active on Twitter than other more established MPs. Thus, the MPs were divided into two groups – those elected in 2010 or later (46% of the total[22]) and those elected pre-2010 (54%). However, an independent samples comparison of medians showed no significant difference in the median number of campaign tweets sent by these two groups (test statistic 3.149, df 1, *p* = .095). Whilst the 2010 cohort is over-represented on Twitter, they are no more active than other MPs elected earlier.

6.4.   Descriptive analysis of tweet characteristics

Twitter's own research (Rogers, 2014) suggests that including elements such as links, images, videos and hashtags increases the chances of tweets being retweeted. Its analysis of tweets in the area of government and politics indicates that the most powerful way of boosting the chances of a retweet is to include an image. Note that this research is based on a sample of verified users only rather than a more general sample, and so cannot be directly compared with the findings of this research which considers both verified and non-verified accounts. This section presents findings from the descriptive analysis of the tweets themselves, focusing on how elements such as

---

[22] This number is higher than just the number of MPs elected in 2010 because it also includes the small number of MPs elected at by-elections since 2010.

hashtags, mentions and links influence retweeting, as shown in the conceptual model presented again in Figure 41. The primary dependent variable on which this research is focused is whether or not the tweet is retweeted – a yes/no categorical flag that makes no distinction between tweets based on how many times they have been retweeted. However, the descriptive analysis does also include some consideration of a secondary dependent variable – the number of times that the tweet is retweeted.

Figure 41 - Conceptual model showing this stage of the descriptive analysis
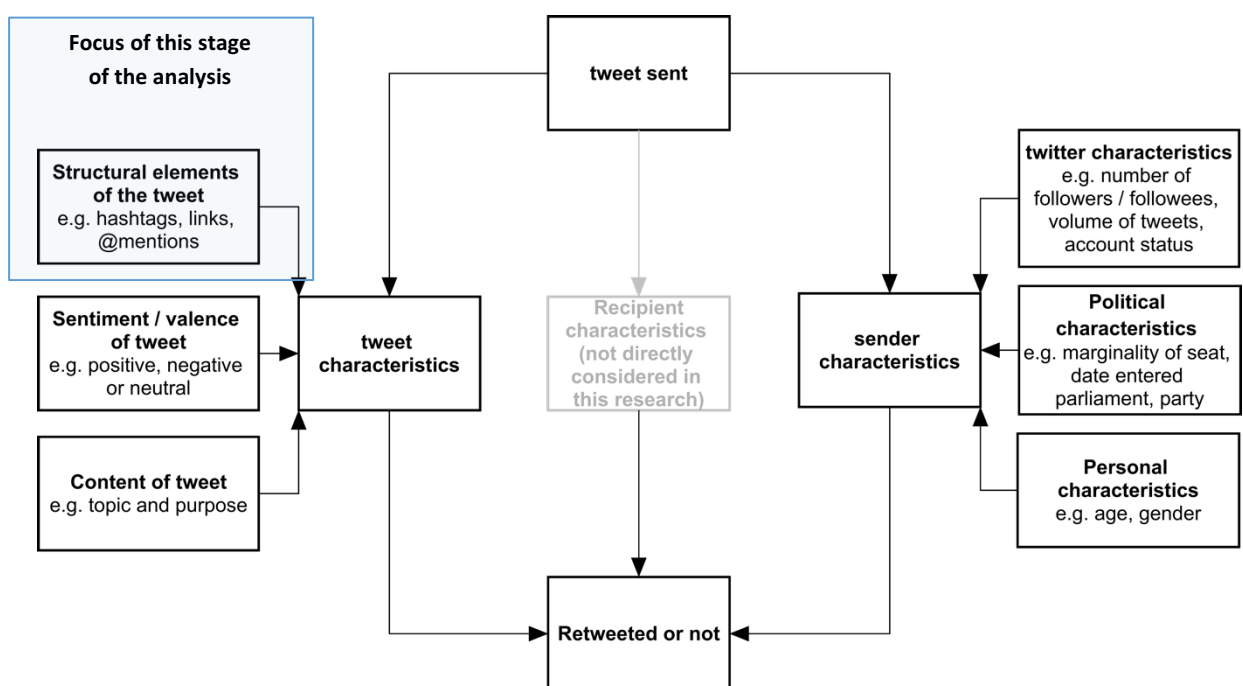
6.4.1. Influence of hashtags on retweeting

Most of the tweets (59%) did not include any hashtags, and so both the median and mode number of hashtags is zero, with a mean of 0.6. For comparison, Golbeck *et al*. (2010) found that only 0.08% of the tweets sent by US Congresspeople included hashtags. However, their research is considerably older than that presented here, and hashtag use has become much more formalised on Twitter in the years since. There does not appear to be any more recent research specifically looking at politicians' use of hashtags but we do know that somewhere around 15% of all tweets contained

hashtags in 2014 (Liu, Kliman-Silver and Mislove, 2014), compared to 41% of the MPs' tweets, suggesting that perhaps MPs are making more use of hashtags than is typical.

Unsurprisingly, including a hashtag improves the chances of a tweet being retweeted. Using hashtags means that your tweet is seen by more than just your immediate followers, and the more people who see your tweet, the more chance there is of it being retweeted. In the population of MPs' tweets, 82% of those without hashtags are retweeted rising to 89% of those with at least one hashtag. A chi-square test (Table 15) shows that this difference is significant ($p < .001$).

Table 15 - Influence of hashtags on retweeting

	Retweeted?		
	No	**Yes**	
No hashtags in tweet	4,587	20,508	Chi square 408.858
	18.3%	81.7%	Sig. 0.00
At least one hashtag in tweet	1,923	15,426	
	11.1%	88.9%	
Total	6,510	35,934	
	15.3%	84.7%	

The number of hashtags used ranges from zero to nine, but with very few tweets including more than three (Figure 42).

Figure 42 - Histogram of hashtags per tweet
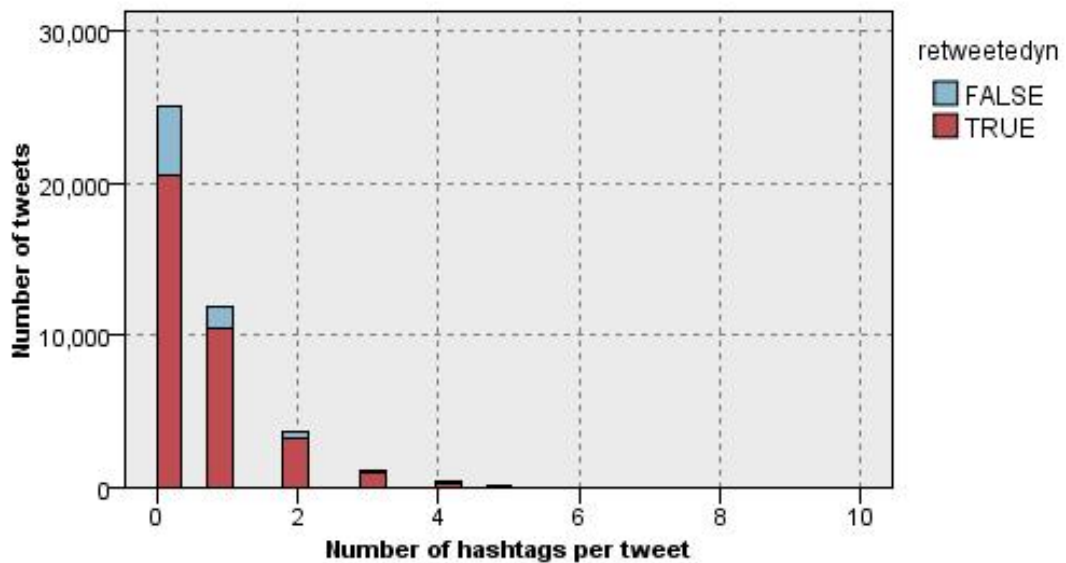
Having determined that including a hashtag significantly improves the chances of getting a tweet retweeted at least once, the next step was to consider the extent to which including a hashtag influences how many times a tweet is likely to be retweeted. Those tweets without any hashtags had a median of three retweets compared to a median of five for tweets that included at least one hashtag. The distribution of retweet numbers in the two groups is non-normal (without hashtags $D(25,079)=0.454$, $p < .001$ and with hashtags ($D(17,349)=0.416$, $p < .001$). An independent samples median test showed that the median retweet numbers for tweets with and without hashtags do differ significantly (Figure 43) – including at least one hashtag in your tweet means that it is likely to get more retweets than a tweet without any hashtags. However, Figure 43 also shows that those tweets with the very highest numbers of retweets did not contain any hashtags. Whilst the tweets with hashtags did better at attracting low numbers of retweets, all the tweets with more than 5,000 retweets were hashtag-free, suggesting that other factors influence very high retweet volume.

Figure 43 - The effect of whether a tweet contains a hashtag on retweeting

Independent–Samples Median Test
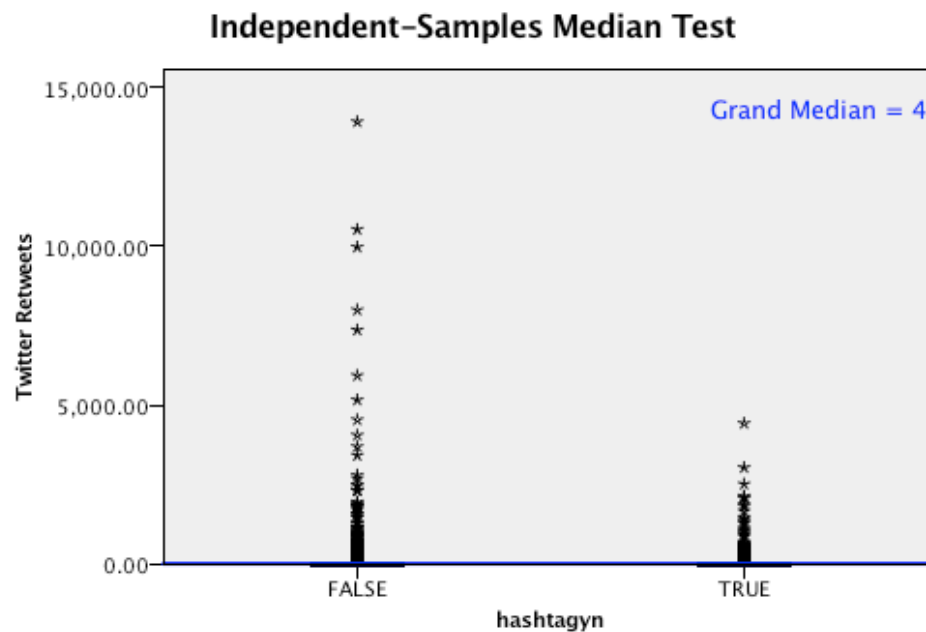
Total N	42,444
Median	4.000
Test Statistic	435.066
Degrees of Freedom	1
Asymptotic Sig. (2–sided test)	.000
Yates's Continuity Correction	Chi–Square	434.653
	Degrees of Freedom	1
	Asymptotic Sig. (2–sided test)	.000

1. Multiple comparisons are not performed because there are less than three test fields.

Which hashtag you include can also make a significant difference to the chances of your tweet being retweeted. For example, 92% of tweets that included #leadersdebate, 96% that included #bbcdebate and 96% mentioning BBC's Question

Time programme (#bbcqt) were retweeted (compared to a retweet rate for the sample as a whole of 85%). Similarly, 90% of #VoteConservative tweets were retweeted, 94% of #LabourDoorstep tweets, 97% of #VoteLabour tweets and 91% of #GE2015 tweets[23]. These are all mainstream hashtags used widely by both parties and the public to signal that their tweets related to the General Election or to the political conversations taking place on television during the campaign.

6.4.2. Influence of @mentions

Another common feature in tweets is use of @mentions. In the case of the MPs, only 38.5% of the tweets contained at least one mention leading to a mode and median of zero once again, with a mean of 0.57. This compares to over 50% across a sample of all tweets in 2014 (Liu, Kliman-Silver and Mislove, 2014) suggesting that MPs are less likely to make use of the @mention feature than average Twitter users. Including at least one mention in the tweet improves its chances of getting retweeted (Table 16), largely because the tweet can then be seen by the followers of the person who is mentioned thus expanding the possible audience and the more people to see the tweet, the more likely it is to get retweeted.

Table 16 - Influence of mentions on retweeting

	Retweeted?		
	No	**Yes**	
No mentions in tweet	4,246	22,046	Chi square 35.041
	16.1%	83.9%	Sig. 0.00
At least one mention	2,264	13,888	
	14%	86%	
Total	6,510	35,934	
	15.3%	84.7%	

---

[23] All significant at $p < .001$

6.4.3. Influence of links

Across the MPs' dataset, 57% of the tweets contain a link (Figure 44), and 34% contain a link specifically to a picture or a video (Figure 45). Golbeck *et al.* (2010) found that 45% of Congresspeople's tweets contained a link. Research in 2014 (Liu, Kliman-Silver and Mislove, 2014) suggests that the figure for the population at large is only 12%. This difference indicates that perhaps MPs are making more use of tweets as a medium for passing on information to others or directing people to websites than do average users.

Figure 44 - Number of tweets that do / do not contain links
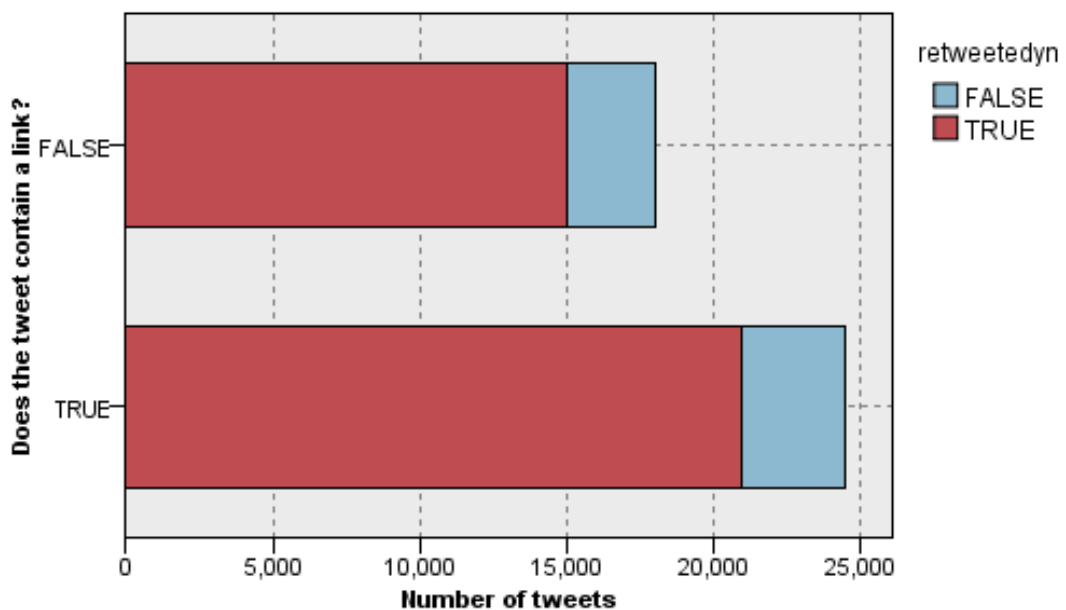
Figure 45 - Number of tweets containing links to videos or pictures
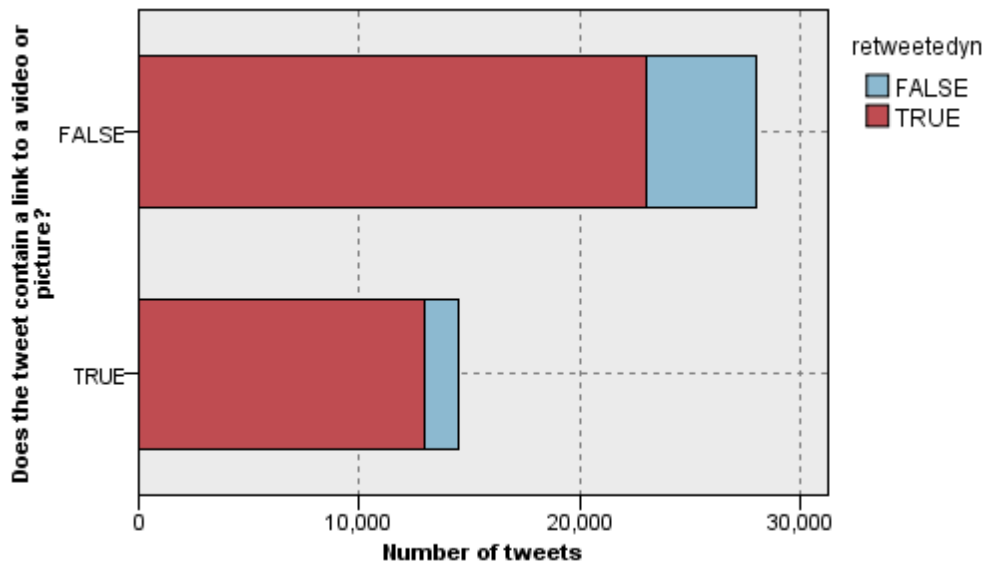
Whether or not a tweet includes a link does affect whether or not it gets retweeted (Figure 46) – tweets with links are more likely to get retweeted than those without.

Figure 46 - How including a link influences whether a tweet is retweeted

| | | | Retweeted yes / no | | |
			No	Yes	Total
Contains URL yes/ no	No	Count	2998	15028	18026
		% within urlyn	16.6%	83.4%	100.0%
		Standardized Residual	4.4	−1.9	
	Yes	Count	3512	20906	24418
		% within urlyn	14.4%	85.6%	100.0%
		Standardized Residual	−3.8	1.6	
Total		Count	6510	35934	42444
		% within urlyn	15.3%	84.7%	100.0%

Chi-Square Tests

	Value	df	Asymptotic Significance (2–sided)	Exact Sig. (2–sided)	Exact Sig. (1–sided)
Pearson Chi–Square	40.383[a]	1	.000		
Continuity Correction[b]	40.210	1	.000		
Likelihood Ratio	40.172	1	.000		
Fisher's Exact Test				.000	.000
N of Valid Cases	42444				

a. 0 cells (0.0%) have expected count less than 5. The minimum expected count is 2764.80.

b. Computed only for a 2x2 table

This effect is more pronounced if the link is to a picture or video rather than, for example, to a webpage (Figure 47). Almost 90% of the tweets that include pictures or videos are retweeted compared to 82% of those that do not.

Figure 47 - How including a picture or video influences retweeting

| | | | Retweeted yes / no | | |
			No	Yes	Total
Tweet contains link to picture or video	No	Count	4968	23009	27977
		% within medialinkyn	17.8%	82.2%	100.0%
		Standardized Residual	10.3	−4.4	
	Yes	Count	1542	12925	14467
		% within medialinkyn	10.7%	89.3%	100.0%
		Standardized Residual	−14.4	6.1	
Total		Count	6510	35934	42444
		% within medialinkyn	15.3%	84.7%	100.0%

Chi–Square Tests

	Value	df	Asymptotic Significance (2–sided)	Exact Sig. (2–sided)	Exact Sig. (1–sided)
Pearson Chi–Square	370.056[a]	1	.000		
Continuity Correction[b]	369.509	1	.000		
Likelihood Ratio	389.304	1	.000		
Fisher's Exact Test				.000	.000
N of Valid Cases	42444				

a. 0 cells (0.0%) have expected count less than 5. The minimum expected count is 2218.93.

b. Computed only for a 2x2 table

The descriptive data analysis presented in this section supports the contention of the conceptual model that both characteristics of a tweet's sender and of the tweet itself influence whether or not that tweet is retweeted. Considered one by one, several variables have been shown to be significantly related to both whether a tweet is retweeted and also to how many times it is retweeted. The next phase of the research is to bring these variables together into a series of more sophisticated predictive models that move beyond describing the relationships between the variables one by one to understanding how they interact with each other, and how they can be used to predict whether or not tweets are retweeted.

6.5.   Phase two - CHAID models using author and tweet variables

As explained in section 5.5, the predictive modelling was conducted using a sample of 13,020 tweets in which half were retweeted and the other half were not. Before modelling commenced the tweets were further partitioned into a training set and a testing set. Each model is built using the training data and then tested on the testing data. Results are reported for both sets. The conceptual model (shown again for convenience in Figure 48) suggests that the characteristics of the sender and of the tweet play a role in determining retweets, and the findings of the descriptive analysis support this contention. When building the CHAID models these two elements were considered separately before being brought together into a single 'master model' (in common with Dang-Xuan *et al*. (2013), Lawless (2012) and others). Additionally, new content and sentiment variables were created during this phase of the research, modelled separately and then included in the master model. The results of each stage of the modelling process are presented in the sections that follow.

Figure 48 - Conceptual model to be tested by CHAID analysis
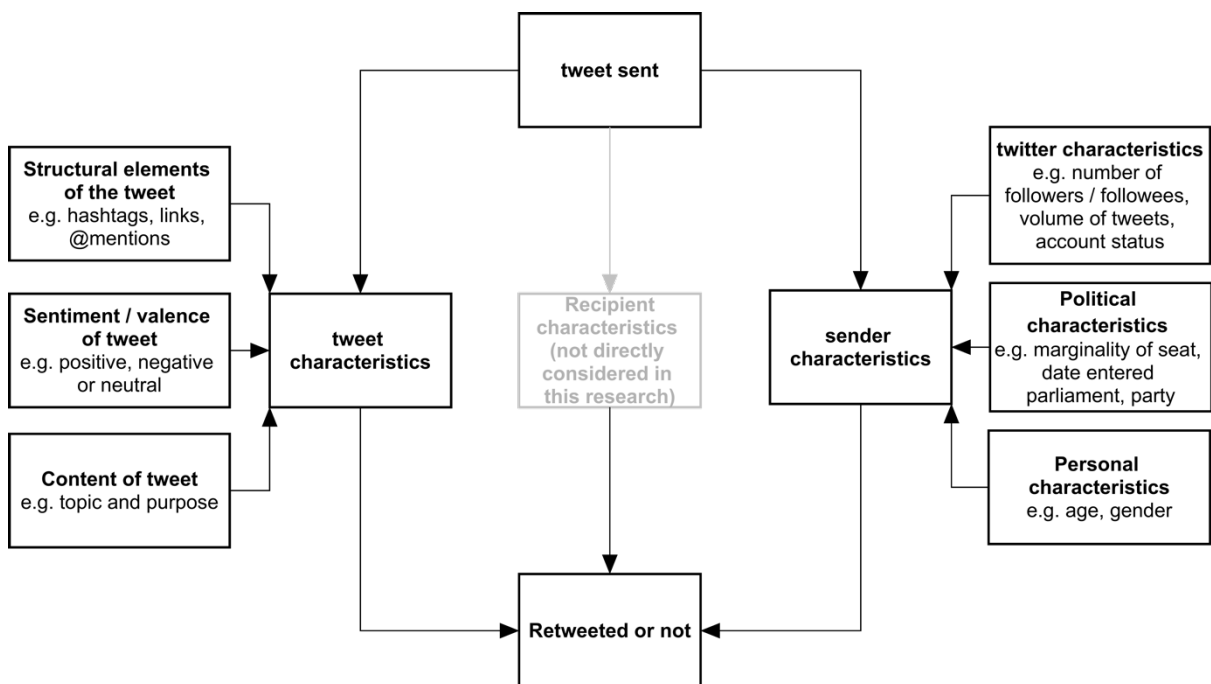
6.5.1. Model one: CHAID model based on structural elements of the tweet

The CHAID algorithm begins by splitting the training data in two groups based on the values of the variable it is trying to predict (in this case whether the tweet is retweeted, yes or no). The algorithm looks for the two purest subgroups possible in relation to the outcome of interest. In model one the first split is performed according to whether the tweets contain a hashtag or not. Node one (not shown) contains all the tweets that do not have hashtags, and node two contains all of those that do. The algorithm then looks at each of the two groups and splits them according to the same principles. In the CHAID decision tree a chi-square statistic is given for each split. This indicates how independent the target field (in this case whether or not a tweet is retweeted) is from the predictor under consideration. The higher the chi-square statistic, the lower the chance that the two variables under consideration are independent, meaning that the split the model has suggested is a good one. This continues until no further subgroups can be usefully generated (or until the maximum number of splits that the analyst specified has been reached). Figure 49 provides a guide to interpreting a CHAID decision tree, using a section of model one as an example.

Figure 49 - Guide to interpreting a CHAID decision tree

At the top of each split the model shows the variable on which the split has been performed. In this case this split is formed on the basis of whether the tweets contain a mention (TRUE) or not (FALSE). The chi-square test tells us that the difference between the two groups (mentionyn TRUE and FALSE) is significant.

The numbers inside each node relate to the variable being predicted. So in this case node 2 contains a total of 1,632 cases of which 39.216% were not retweeted (FALSE) and 60.784% were retweeted (TRUE).

Node 6 tells us that of those tweets that contained a mention (mentionyn=TRUE) 66.27% were retweeted.

Within the group of tweets that do not contain a mention, the model makes a further split between those that contain a URL and those which do not (urlyn)

TRUE

Node 2

Category	%	n
FALSE	39.216	640
TRUE	60.784	992
Total	36.380	1632

mentionyn
Adj. P-value=0.000, Chi-square=14.461, df=1

FALSE

Node 5

Category	%	n
FALSE	43.066	413
TRUE	56.934	546
Total	21.378	959

urlyn
Adj. P-value=0.007, Chi-square=7.285, df=1

TRUE

Node 6

Category	%	n
FALSE	33.730	227
TRUE	66.270	446
Total	15.002	673

medialinkyn
Adj. P-value=0.000, Chi-square=14.111, df=1

The first CHAID model was built using only variables that relate to the structural characteristics of the tweets themselves, namely the time at which the tweet was sent, whether it included a hashtag (y/n), the number of hashtags included, whether it included any mentions (y/n), the number of @mentions in the message, whether it included a link, and whether it included a link to a picture or video (i.e. media link). Model one is shown in Figure 50 and will be discussed in more detail on the following pages.

Figure 50 - Model one: variables relating to tweets' structural elements

List of variables in this model

hashtagyn – does the tweet contain a hashtag yes/no?

medialinkyn – does the tweet contain a link to an image or video yes/no?

urlyn – does the tweet contain a link of any kind yes/no?

mentionnumber – how many mentions are there in the tweet?

mentionyn – does the tweet contain a mention?
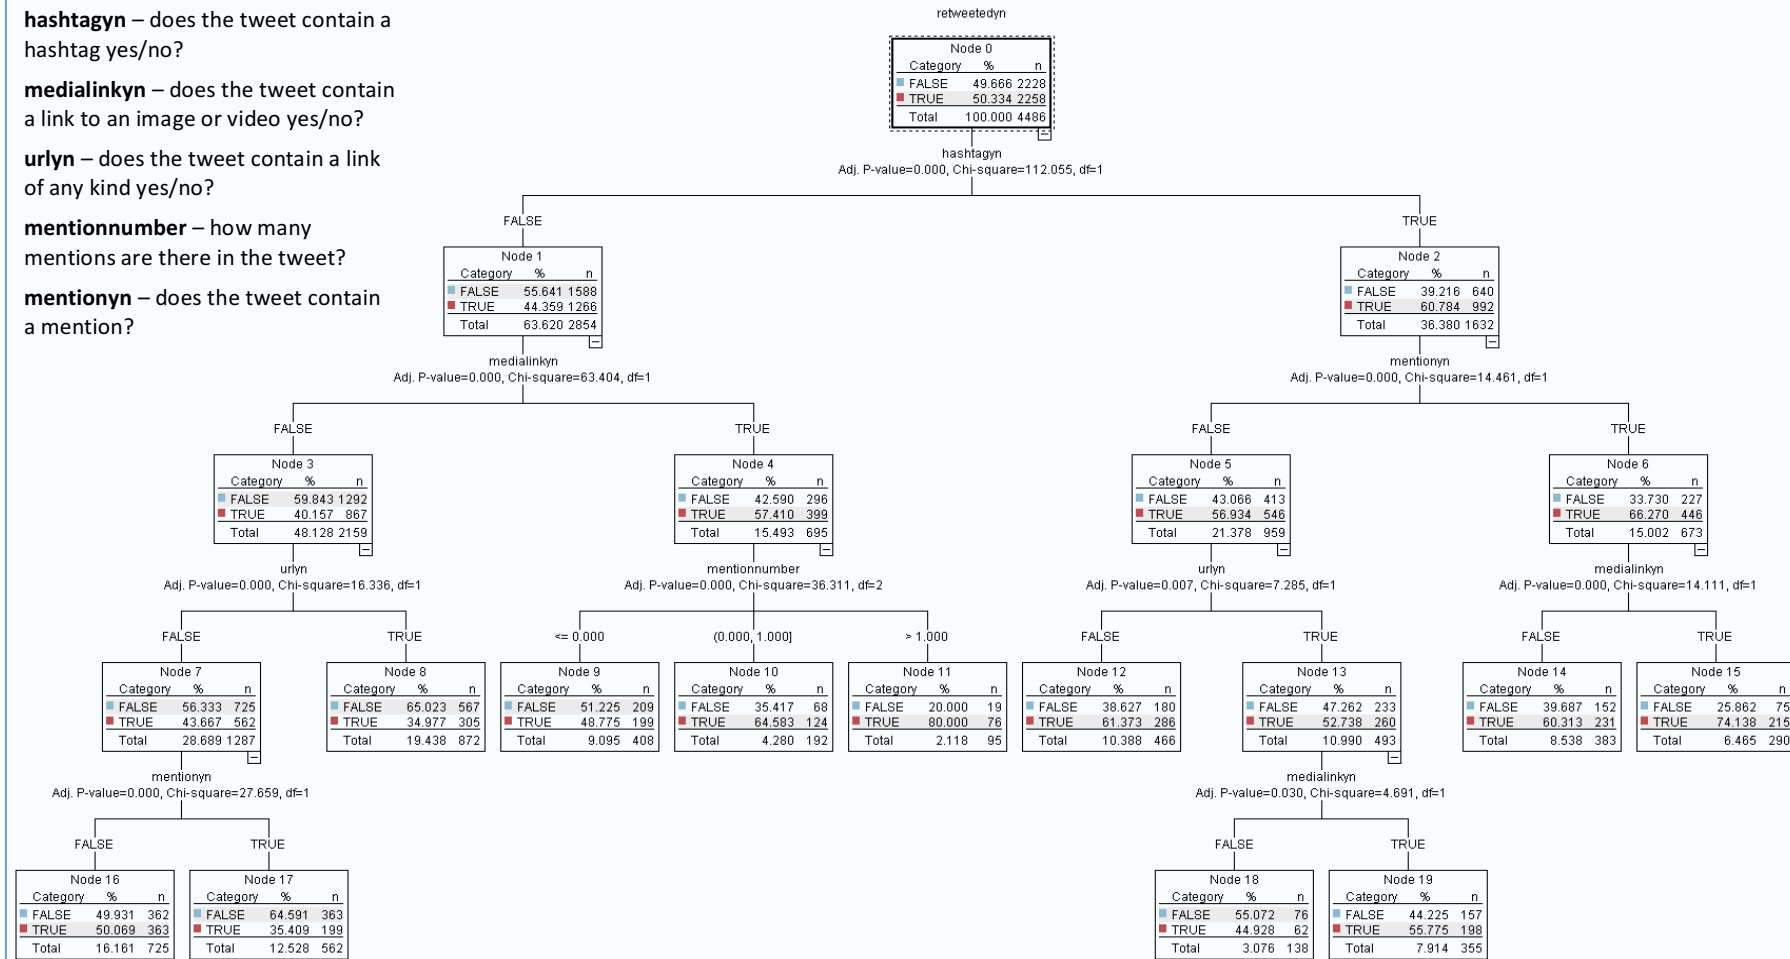
SPSS Modeler provides a predictor importance chart for each CHAID model (Figure 51) which gives an indication of how important each of the predictors was in estimating the model. The sum of the values for all the predictors is 1. In the case of model one, the CHAID algorithm has determined that whether or not a tweet contains a link to a video or picture is the most important predictor of retweeting, followed by whether it has a hashtag. Below that comes whether the tweet includes a mention, whether it has a URL link and the number of mentions that it includes. Unlike other modelling approaches such as regression analysis, the CHAID model automatically excludes variables that do not significantly add to the predictive power of the model. In this case, the time at which the tweet was sent was not deemed to be predictive of retweeting and so was excluded from the model.

Figure 51 - Relative importance of variables in model one
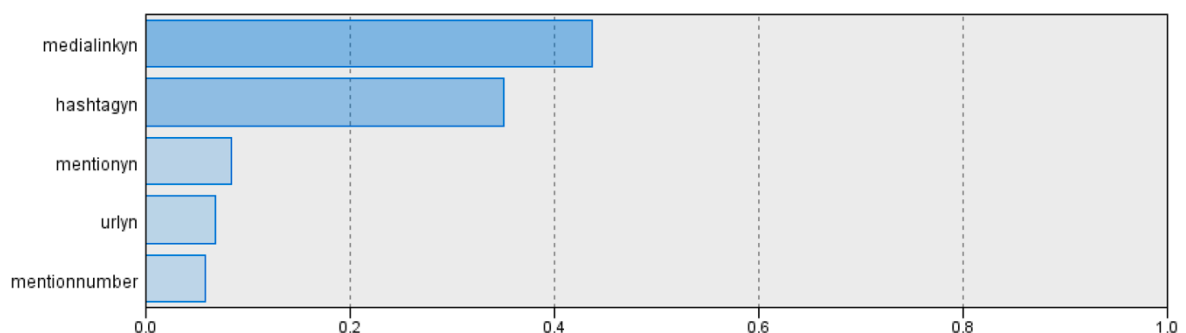
For ease of interpretation, the CHAID model can also be expressed in terms of a set of rules that it uses to determine which group a particular tweet is most likely to fall into. The full set of rules governing model one is shown in appendix six. In this case, the variable of interest is retweetyn which shows as FALSE if a tweet is not retweeted and TRUE if it is retweeted. For example:

- If a tweet does not have a media link, nor a URL, nor a mention, nor a hashtag then retweetyn will equal FALSE

- If a tweet contains a media link, a hashtag and a mention then retweetyn will equal TRUE

Overall, model one was correct in its predictions for 60.27% of the cases in the training data and 59.29% of the cases in the testing data (Table 17). However, within the model were some individual nodes with considerably higher success rates. Looking at the decision tree, it can be seen that Node 15 contains the one of the highest percentage of retweets – this node shows that if a tweet contains at least one hashtag, at least one mention and a link to a picture or a video then it will be retweeted 74% of the time. In contrast, Node 17 shows that tweets that do not have hashtags, nor media links, nor other kinds of URLs, but do include at least one mention will be retweeted only 33% of the time.

Table 17 - Comparison of performance on training and testing data for model one

	Training		Testing	
Correct	3,878	60.27%	3,905	59.29%
Wrong	2,556	39.73%	2,681	40.71%
Total	6,434		6,586	

The model generates a prediction for each individual tweet and a confidence score indicating how confident it is that its prediction is correct in each case. It then creates two new variables which are written back into the dataset. $R-retweetyn gives the model's prediction for each tweet – retweeted or not retweeted. $RC-retweetedyn gives a confidence score of between 0 and 1 indicating how confident the model is that its prediction for each tweet is correct. The confidence values report for model one shows that the model ranges in confidence in its predictions for each tweet from 0.5 to 0.779 (as shown in Table 18). A confidence score of 0.5 would indicate that the model is no more confident than random chance. A confidence score of 1 would indicate that the model was completely certain in its prediction for that particular tweet. The mean correct score is the average confidence score for all those tweets where the model correctly predicted the outcome. The *always correct above* figure

shows the confidence score above which the model was always correct in its predictions. In this model no cases fell into this category. Similarly, *always incorrect below* shows the confidence level below which the model was always wrong. Again, this did not apply to any cases for this model. *90% accuracy* shows the confidence score above which the model was correct in its predictions 90% of the time. That level was not reached for this model.

Table 18 - Evaluation of confidence scores for model one

	Training data	Testing data
Range	0.5-0.779	0.5-0.779
Mean correct	0.614	0.611
Mean incorrect	0.587	0.588
Always correct above	0.779 (0% of cases)	0.779 (0% of cases)
Always incorrect below	0.5 (0% of cases)	0.5 (0% of cases)
90% accuracy above	Never reached	Never reached

The predictive power of the model can also be evaluated by looking at the area under the curve (AUC) of the model. In this case, the AUC is 0.643 for the training data and 0.635 for the testing data. If the model were correct 100% of the time, then the AUC would be 1. If it were no better than random chance then the AUC would be 0.5, so an AUC of 0.635 indicates that the model performs better than random chance.

The predictive power of each CHAID model can also be evaluated visually in a gains chart (Figure 52). The central diagonal line on the gains chart shows the expected rate of retweeting for the whole sample if the CHAID model were not used and the tweets were sent out in a random order. The tweets are split 50/50 between those retweeted and not retweeted, so the line is a constant diagonal. The curved line above the diagonal shows how much better the retweet performance would be if the tweets were scored by the model before being sent and then sent out in order of likelihood to get retweeted, most likely through to least likely. In this case, one can see that sending out about 60% of the tweets, reordered in this way, would generate about 70% of the potential retweeted tweets.

Figure 52 - Gains chart for model one
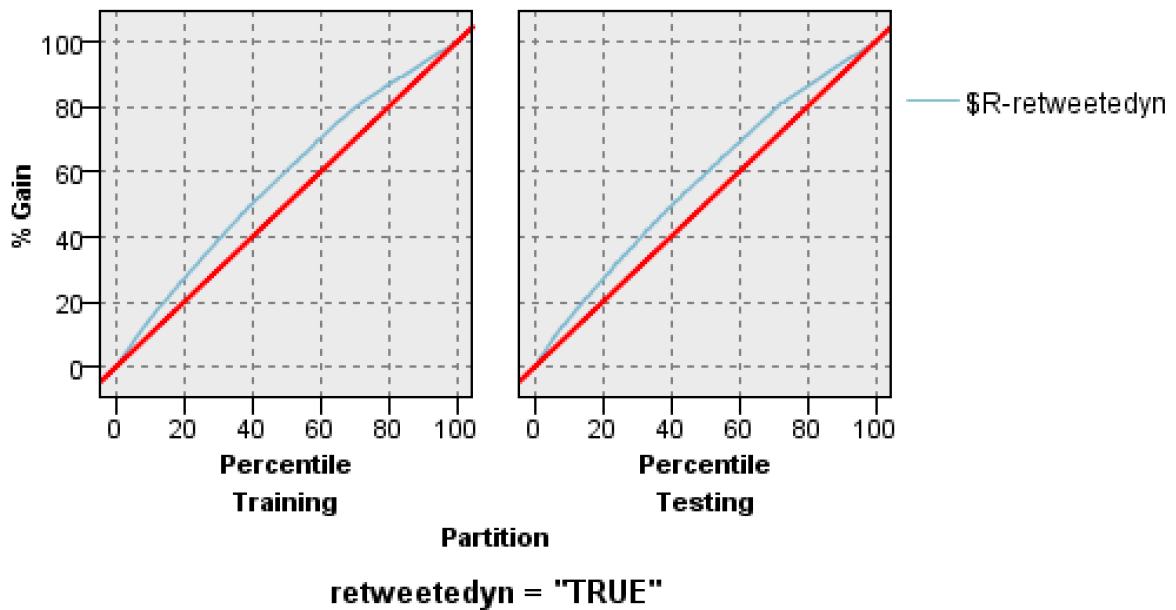
A gains chart would more typically be used for marketing applications such as scoring a customer database according to how likely each customer is to respond to a particular campaign. The gains chart would show what percentage of the database would need to be mailed to generate a particular percentage of responses, comparing the results if the model is not used with the results using the model. A gains chart has limited practical value for this research as there is no possibility that an MP would write a large number of tweets in advance, score them according to a CHAID model and then send them out in the order suggested by the model. However, the gains chart can still provide a useful visual indication of how much better than chance the developed model performs.

Finally, a gains table can be generated for each model, showing statistics for each of the terminal nodes in the CHAID tree (Table 19). The gain is an indication of how far the proportion of retweeted tweets in each node differs from the proportion in the data as a whole. The greater the difference between the proportion in the node and the proportion in the data, the more powerful the model. For example, node 15 has

an index value of almost exactly 150%. This means that the records in node 15 are about 1.5 times more likely to be retweeted than is the case in the dataset as a whole.

Table 19 - Gains table for model one

Nodes	Node: n	Node (%)	Gain: n	Gain (%)	Response (%)	Index (%)
11	95.00	2.12	76.00	3.37	80.00	158.94
15	290.00	6.46	215.00	9.52	74.14	147.29
10	192.00	4.28	124.00	5.49	64.58	128.31
12	466.00	10.39	286.00	12.67	61.37	121.93
14	383.00	8.54	231.00	10.23	60.31	119.83
19	355.00	7.91	198.00	8.77	55.77	110.81
16	725.00	16.16	363.00	16.08	50.07	99.47
9	408.00	9.09	199.00	8.81	48.77	96.90
18	138.00	3.08	62.00	2.75	44.93	89.26
17	562.00	12.53	199.00	8.81	35.41	70.35
8	872.00	19.44	305.00	13.51	34.98	69.49

In summary, model one tells us that overall the best way to get a tweet retweeted is to include a hashtag. The chances are further improved if someone else is mentioned in the tweet. If you do both of these things and also include a link to an image or a video then your tweet will be retweeted just under 75% of the time (node 15 Figure 50). If you do not include a hashtag you can still have a high chance of being retweeted if you include a link to an image or video and mention more than one other person. These tweets are retweeted 78% of the time (node 11).

6.5.2. Model two: CHAID model using author characteristics to predict retweets

The second CHAID model used only variables that relate to the characteristics of tweets' senders. The first iteration of this model included gender, party, age, percent majority before the election, marginality of seat, year of birth, year entered parliament, Kred outreach, Kred influence, number of Twitter followers, number of people following, follower/followee ratio, total number of tweets sent, total number of campaign tweets, mean number of tweets per day of the campaign and account status (verified or not). The first iteration (not shown here) was 72.4% accurate in its predictions (on the testing data), so considerably better than random chance and also better than the previous model which used just tweet characteristics.

The most important predictor was the number of followers that someone has. It is not surprising that people with more followers get more retweets. If you have a lot of followers, then more people see your tweets and can potentially retweet them[24]. This then is not particularly helpful information when trying to advise an MP on how to compose a tweet with a high chance of being retweeted, as the number of followers one has is not something that can be directly controlled in the same way as the content of a tweet can be. Thus, number of followers was removed from the model. Using the same logic, age, gender, year of birth, percent majority, marginality of seat, year entered parliament and account status were also removed as they cannot be influenced or changed by the MP either. This second iteration of the model was accurate 71.54% of the time (on the testing data), so removing these variables had a relatively small effect on the model's predictive power.

This second iteration showed Kred influence to be the most important predictor of retweeting. An MP accumulates Kred points whenever someone replies to one of their posts, retweets it, follows them or adds them to a list. In effect then a high Kred score is a measure of the number of retweets that someone generates, so keeping it in the model is saying that to get a tweet retweeted one needs to generate a lot of retweets – not particularly helpful advice. Therefore, Kred influence and outreach scores were removed from the model too.

The final version of the author variables model was based just on the number of people the MP followed, their ratio of followers to followees, the total number of tweets they sent, the number of campaign tweets they sent and the mean number of campaign tweets per day they sent. Of those variables, the model determined that only four were predictive – campaign tweets per day, number of people followed, total tweets sent during the person's entire time on Twitter and their ratio of followers to followees, with tweets per day being the most important (Figure 53). The

---

[24] Note, it is not necessary to follow someone in order to retweet their tweets, so we cannot assume that all an MP's retweets come from their direct followers. Retweets could also come from the followers of their followers, hence why follower numbers matter as a larger number of followers exposes your tweets to more people.

total number of campaign tweets was excluded from the model. This version of the model was correct in its predictions 69.36% of the time which shows that it is possible to build a model using only those author-related variables that the author has control over which is only slightly less predictive than a model including all author-related variables.

Figure 53 - Relative importance of predictors in author variables model
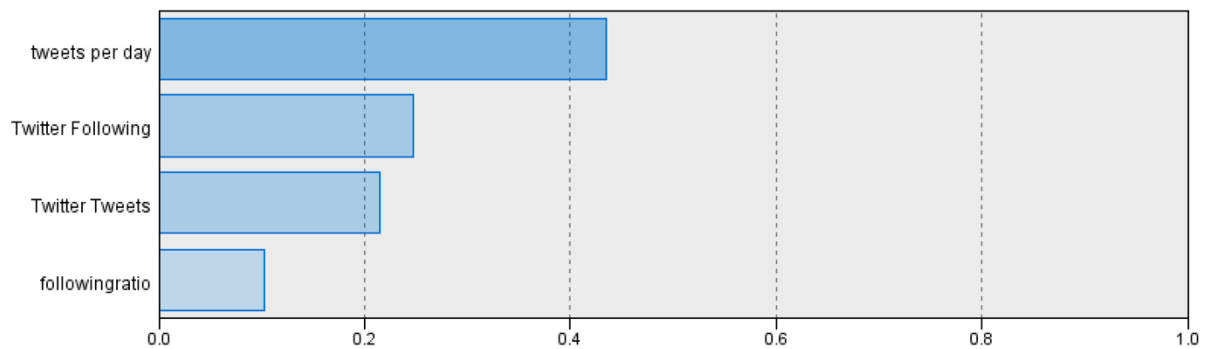
Tweets per day is a continuous field and CHAID is not restricted to binary splits, so has divided the number of tweets per day into several groups that it determines to be predictive. Thus, model two is very wide and has been split for ease of viewing across Figure 56-61[25]. The relationship between tweets per day and retweeting is not linear. Looking at the first level nodes in the CHAID tree one can see that the model has split the tweets per day into nine groups, and that the percentage of tweets that are retweeted does not simply go up as the number of tweets goes up.

Figure 54 shows the relationship between number of tweets per day of the campaign and the percentage of tweets that are retweeted. A Spearman's rank correlation shows that number of tweets an MP sends per day and percentage of their tweets that get retweeted are not significantly correlated (r_s = -0.28, p (one-tailed) = .297).

---

[25] A guide to interpreting CHAID output is given in Figure 49 on page 190

Figure 54 - Relationship between mean tweets per day and % of posts retweeted
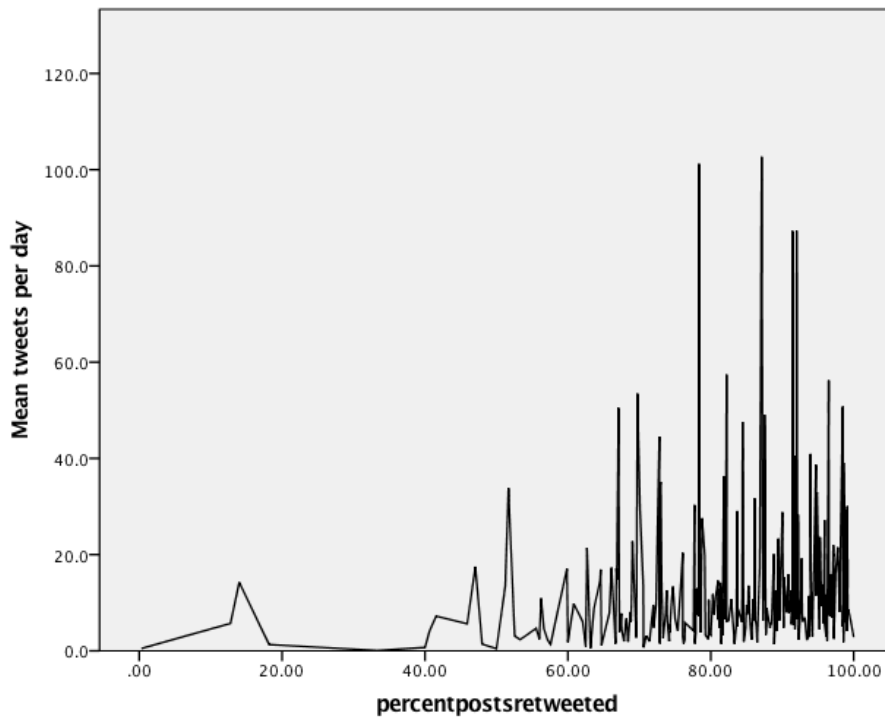
The CHAID model is not looking for linear correlations so will split the data in any way that will enable it to segment into retweeted / not retweeted groups. Thus, it can on occasion produce very wide models when working with continuous variables, as is the case here, which can be difficult to interpret meaningfully, acknowledged as a limitation of this method. The fact that number of tweets per day shows as predictive but that its relationship with retweet chance is not linear suggests that perhaps number of tweets per day is acting as a proxy in this model for some other variable which is not included such as, for example, the number of followers that one has. The full set of rules governing this model can be found in appendix six.

Figure 55 - Full model two using author characteristics
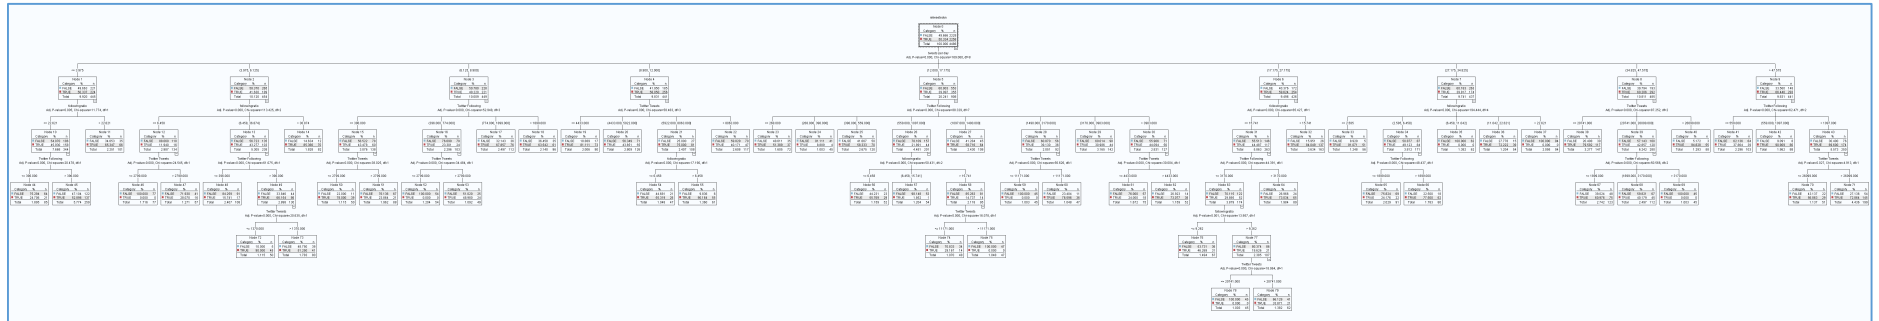
Figure 56 - Model two using author characteristics, part one
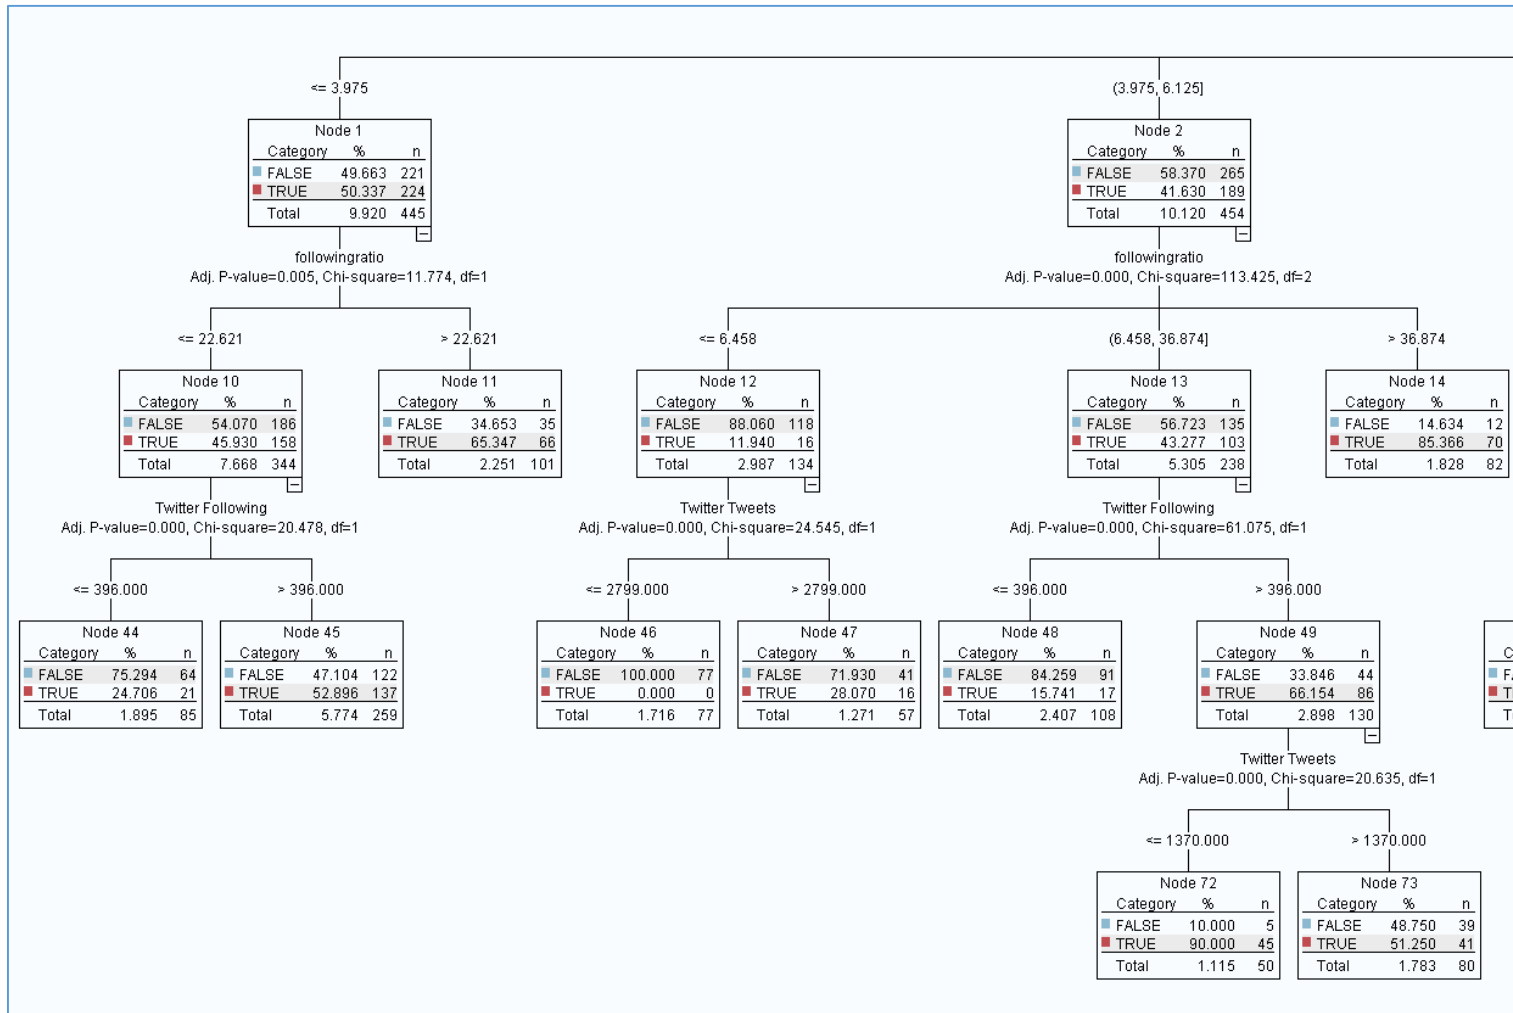
Variables included in this model

Tweets per day – the number of tweets that the MP sent per day of the campaign

Followingratio – number of followers divided by number of followees

Twitter following – number of people the MP follows

Twitter tweets – total number of tweets sent by the MP in their time on Twitter

Figure 57 - Model two using author characteristics, part two

Variables included in this model

Tweets per day – the number of tweets that the MP sent per day of the campaign

Followingratio – number of followers divided by number of followees

Twitter following – number of people the MP follows

Twitter tweets – total number of tweets sent by the MP in their time on Twitter

Figure 58 - Model two using author characteristics part three

Variables included in this model

Tweets per day – the number of tweets that the MP sent per day of the campaign

Followingratio – number of followers divided by number of followees

Twitter following – number of people the MP follows

Twitter tweets – total number of tweets sent by the MP in their time on Twitter

Figure 59 - Model two using author characteristics part four

Variables included in this model

Tweets per day – the number of tweets that the MP sent per day of the campaign

Followingratio – number of followers divided by number of followees

Twitter following – number of people the MP follows

Twitter tweets – total number of tweets sent by the MP in their time on Twitter

Figure 60 - Model two using author characteristics part five
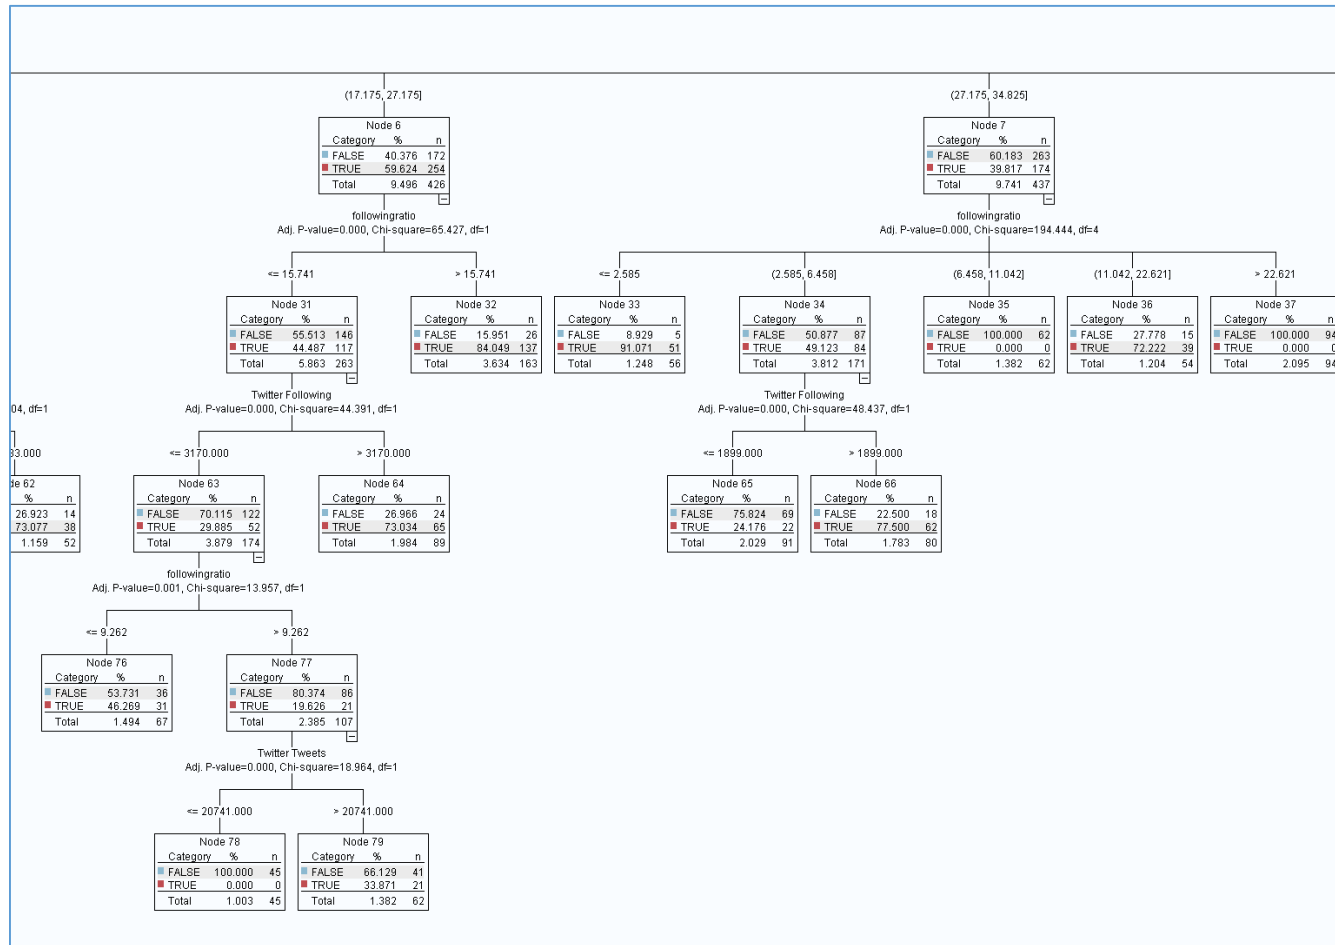
Variables included in this model

Tweets per day – the number of tweets that the MP sent per day of the campaign

Followingratio – number of followers divided by number of followees

Twitter following – number of people the MP follows

Twitter tweets – total number of tweets sent by the MP in their time on Twitter

Figure 61 - Model two using author characteristics part six
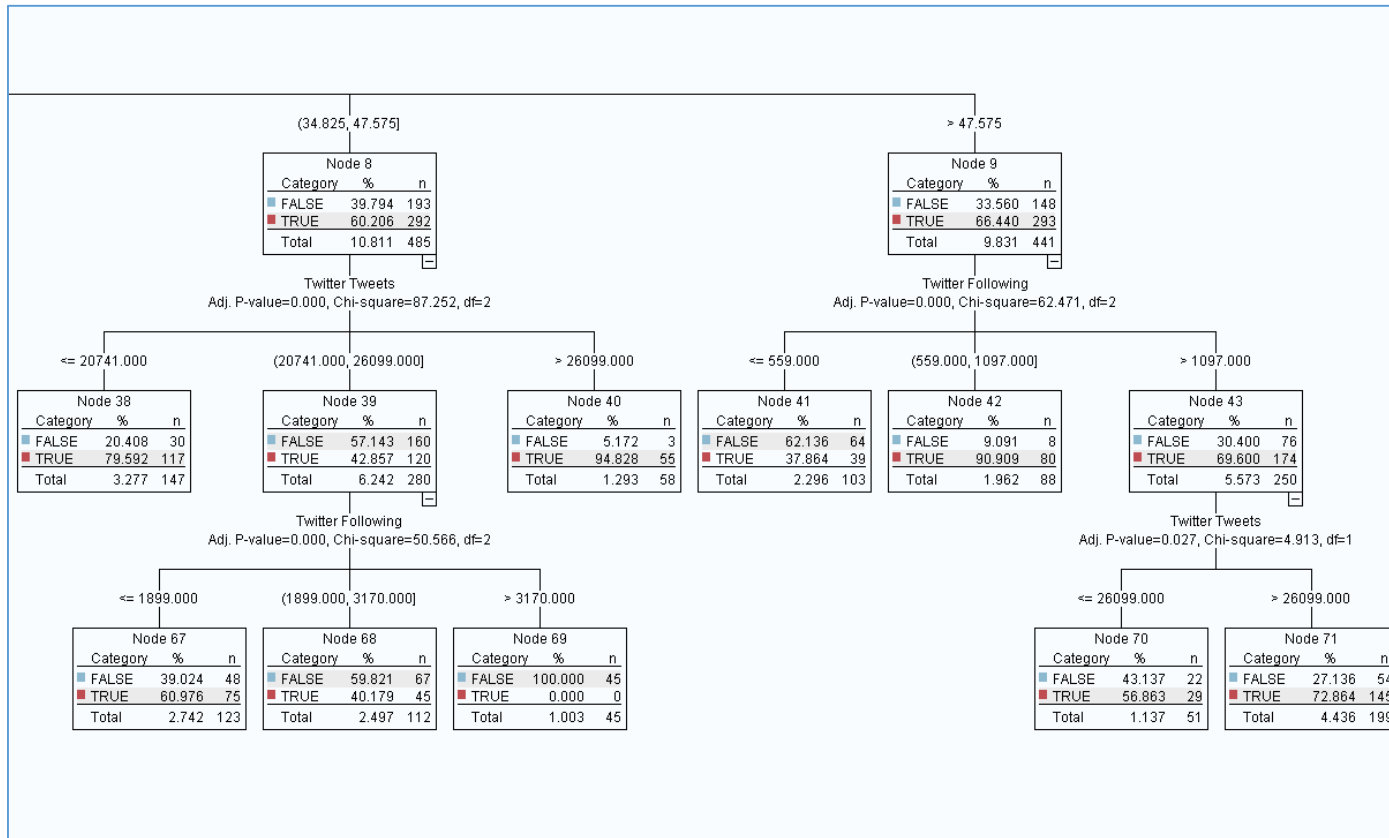
Variables included in this model

Tweets per day – the number of tweets that the MP sent per day of the campaign

Followingratio – number of followers divided by number of followees

Twitter following – number of people the MP follows

Twitter tweets – total number of tweets sent by the MP in their time on Twitter

As can be seen in Table 20, model two is accurate in its predictions just under 70% of the time. Thus, model two, which focuses on the authors' characteristics, performs considerably better than model one which only included data relating to the structural content of the tweet.

Table 20 - Comparison of performance on training and testing data for model two

	Training		Testing	
Correct	4,549	70.7%	4,568	69.36%
Wrong	1,885	29.3%	2,018	30.64%
Total	6,434		6,586	

Table 21 - Evaluation of confidence scores for model two

	Training data	Testing data
Range	0.508-0.994	0.508-0.994
Mean correct	0.735	0.737
Mean incorrect	0.626	0.631
Always correct above	0.97 (10.32% of cases)	0.97 (10.32% of cases)
Always incorrect below	0.508 (0% of cases)	0.508 (0% of cases)
90.74% accuracy above	0.775	0.778

Figure 62 - Gains chart for model two
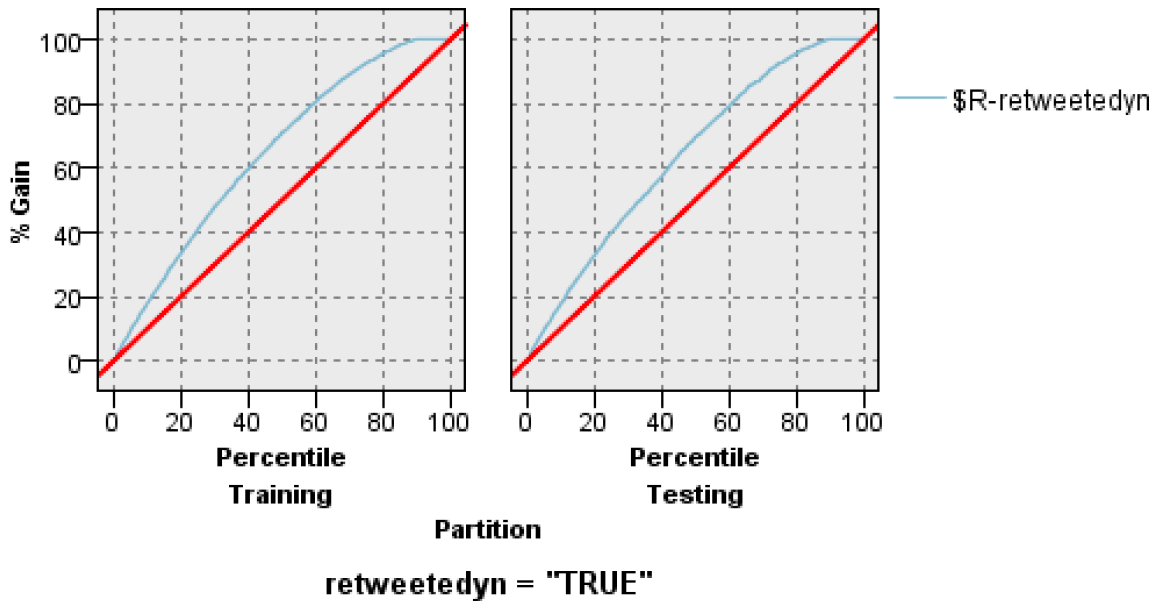
The area under the curve (shown visually in Figure 62) is 0.791 for the training data and 0.775 for the testing data. As model two is extremely large, the gains table is shown by deciles rather than listing the gains for each individual node. As Table 22 shows, there are many nodes in this model that perform better than random chance. For example, the nodes in the tenth percentile are nearly twice as likely to contain tweets that were retweeted than is the dataset as a whole.

Table 22 - Gains table for model two

Nodes	Percentile	Percentile: n	Gain: n	Gain (%)	Response (%)	Index (%)
40,33,42,55,72...	10.00	449.00	401.00	17.78	89.40	177.60
32,19,38,50,66	20.00	897.00	762.00	33.76	84.98	168.84
66,60,62,64,71...	30.00	1346.00	1093.00	48.39	81.17	161.26
36,17,11,18,67	40.00	1794.00	1383.00	61.23	77.07	153.11
67,27,25,70,56...	50.00	2243.00	1638.00	72.54	73.02	145.08
45,23,73,53,76	60.00	2692.00	1867.00	82.70	69.37	137.82
76,20,68,22,41	70.00	3140.00	2051.00	90.82	65.31	129.75
41,79,29,74,47...	80.00	3589.00	2183.00	96.68	60.82	120.84
65,61,51,48,24...	90.00	4037.00	2258.00	100.00	55.93	111.12
37,46,35,52,75...	100.00	4486.00	2258.00	100.00	50.33	100.00

6.5.3. Model three: CHAID model including both author and structural data

For the third model tweet structural data and author data were blended in a single model (model three shown in Figure 63) to see if this would improve the model's predictive power or make the results easier to interpret.

Figure 63 - Model three: combined author and tweet characteristics[26]
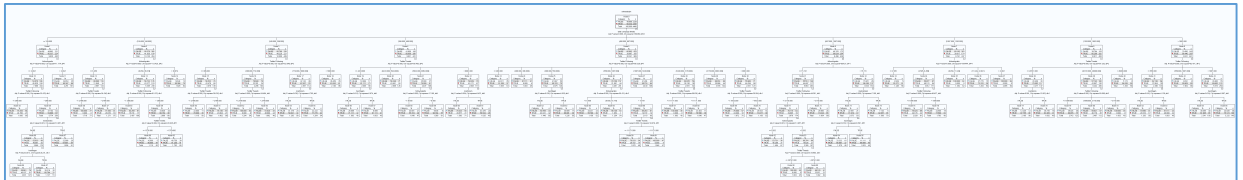
This combined model was accurate 71.7% of the time on the training data and 69.92% on the testing data (Table 23). Thus combining the two types of variable together does not significantly increase the predictive power of the model, and the characteristics of the person sending the tweet are more important than the characteristics of the tweet itself when it comes to determining how likely the tweet is to get retweeted. Model three is shown in Figure 63 and split for ease of reading across Figure 64, Figure 65, Figure 66 and Figure 67.

Table 23 - Comparison of performance on training and testing data for model three

	Training		Testing	
Correct	4,613	71.7%	4,605	69.92%
Wrong	1,821	28.3%	1,981	30.08%
Total	6,434		6,586	

---

[26] A guide to interpreting CHAID output is given in Figure 49 on page 190

Figure 64 - Model three: combined author and tweet characteristics part one

Variables included in this model

Total campaign tweets – number of tweets MP sent during the campaign

Tweets per day – the number of tweets sent per day of the campaign

Followingratio – number of followers divided by number of followees

Twitter following – number of people the MP follows

Twitter tweets – total number of tweets sent by the MP in their time on Twitter

Medialinkyn – whether the tweet contains a link to an image or video yes/no

Hashtagyn – whether the tweet contains any hashtags yes/no

Mentionyn – whether the tweet contains any mentions yes/no

Figure 65 - Model three: combined author and tweet characteristics part two
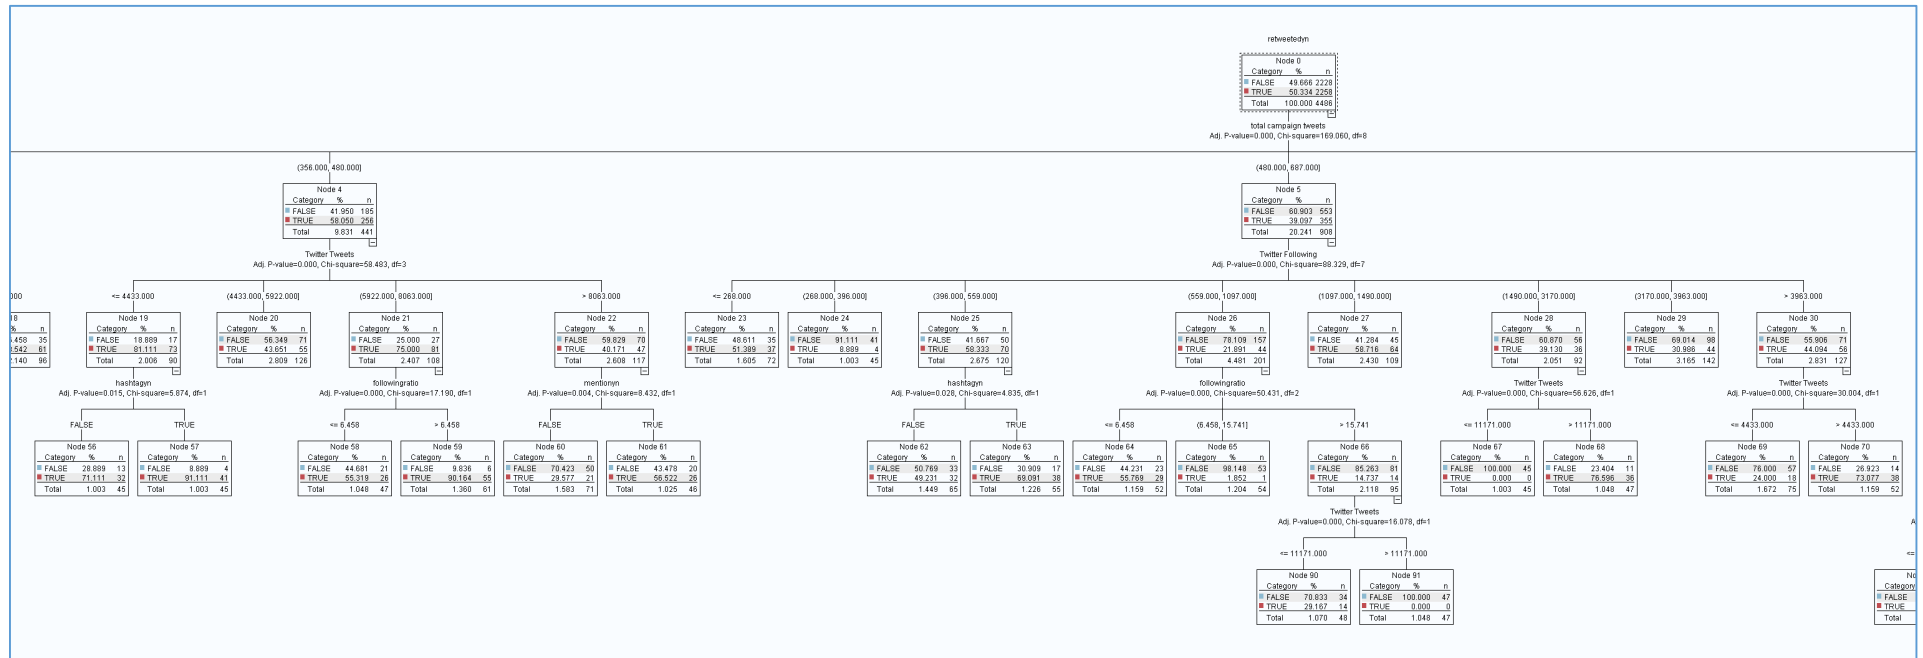
Variables included in this model

Total campaign tweets – number of tweets MP sent during the campaign

Tweets per day – the number of tweets sent per day of the campaign

Followingratio – number of followers divided by number of followees

Twitter following – number of people the MP follows

Twitter tweets – total number of tweets sent by the MP in their time on Twitter

Medialinkyn – whether the tweet contains a link to an image or video yes/no

Hashtagyn – whether the tweet contains any hashtags yes/no

Mentionyn – whether the tweet contains any mentions yes/no

Figure 66 - Model three: combined author and tweet characteristics part three
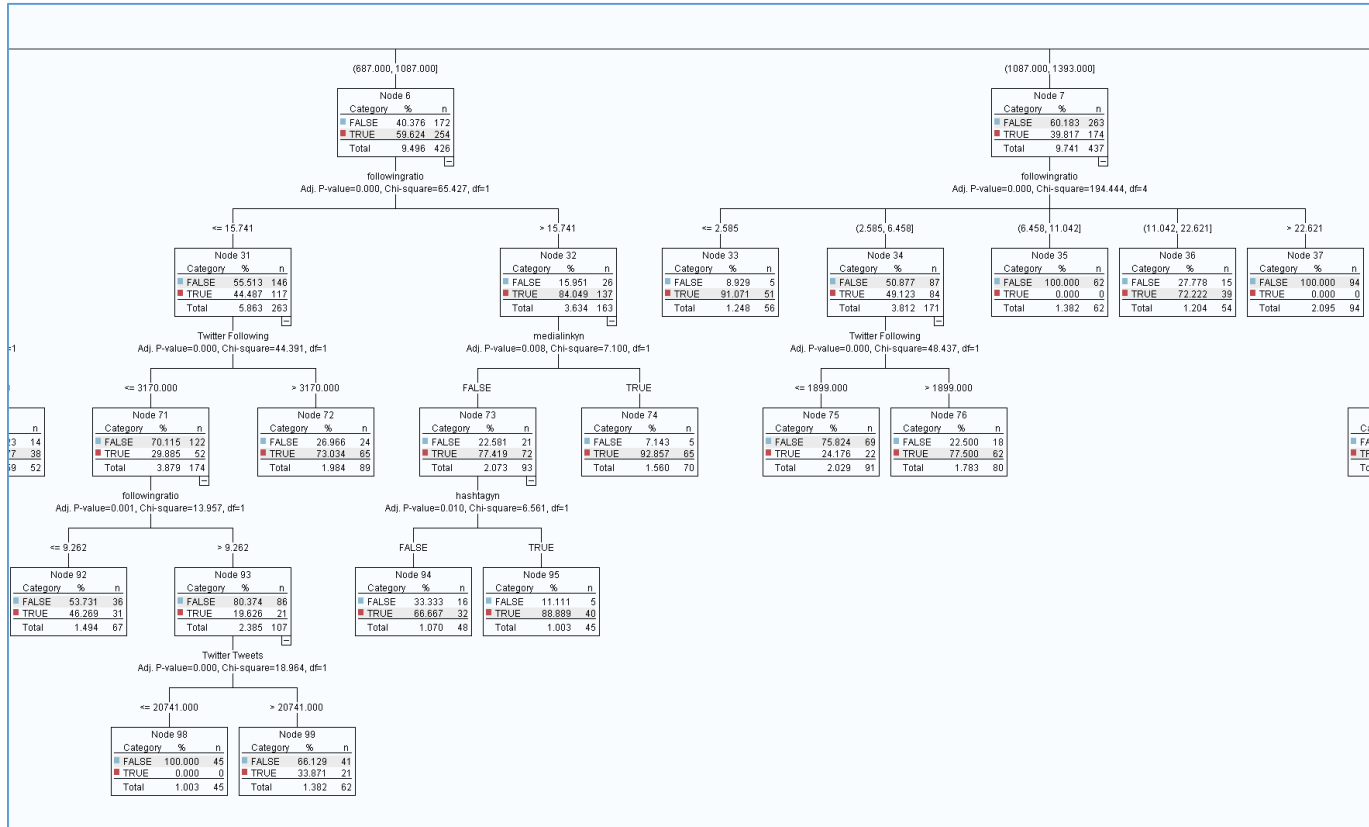
Variables included in this model

Total campaign tweets – number of tweets MP sent during the campaign

Tweets per day – the number of tweets sent per day of the campaign

Followingratio – number of followers divided by number of followees

Twitter following – number of people the MP follows

Twitter tweets – total number of tweets sent by the MP in their time on Twitter

Medialinkyn – whether the tweet contains a link to an image or video yes/no

Hashtagyn – whether the tweet contains any hashtags yes/no

Mentionyn – whether the tweet contains any mentions yes/no

Figure 67 - Model three: combined author and tweet characteristics part four

(1393.000, 1903.000] > 1903.000

Node 8

Category	%	n
FALSE	39.794	193
TRUE	60.206	292
Total	10.811	485

Node 9

Category	%	n
FALSE	33.560	148
TRUE	66.440	293
Total	9.831	441

Twitter Tweets
Adj. P-value=0.000, Chi-square=87.252, df=2

Twitter Following
Adj. P-value=0.000, Chi-square=62.471, df=2

<= 20741.000 (20741.000, 26099.000] > 26099.000 <= 559.000 (559.000, 1097.000] > 1097.000

n
94
0
94

Node 38

Category	%	n
FALSE	20.408	30
TRUE	79.592	117
Total	3.277	147

Node 39

Category	%	n
FALSE	57.143	160
TRUE	42.857	120
Total	6.242	280

Node 40

Category	%	n
FALSE	5.172	3
TRUE	94.828	55
Total	1.293	58

Node 41

Category	%	n
FALSE	62.136	64
TRUE	37.864	39
Total	2.296	103

Node 42

Category	%	n
FALSE	9.091	8
TRUE	90.909	80
Total	1.962	88

Node 43

Category	%	n
FALSE	30.400	76
TRUE	69.600	174
Total	5.573	250

mentionyn
Adj. P-value=0.001, Chi-square=12.062, df=1

Twitter Following
Adj. P-value=0.000, Chi-square=50.566, df=2

hashtagyn
Adj. P-value=0.042, Chi-square=4.129, df=1

hashtagyn
Adj. P-value=0.024, Chi-square=5.067, df=1

FALSE TRUE <= 1899.000 (1899.000, 3170.000] > 3170.000 FALSE TRUE FALSE TRUE

Node 77

Category	%	n
FALSE	12.245	12
TRUE	87.755	86
Total	2.185	98

Node 78

Category	%	n
FALSE	36.735	18
TRUE	63.265	31
Total	1.092	49

Node 79

Category	%	n
FALSE	39.024	48
TRUE	60.976	75
Total	2.742	123

Node 80

Category	%	n
FALSE	59.821	67
TRUE	40.179	45
Total	2.497	112

Node 81

Category	%	n
FALSE	100.000	45
TRUE	0.000	0
Total	1.003	45

Node 82

Category	%	n
FALSE	70.690	41
TRUE	29.310	17
Total	1.293	58

Node 83

Category	%	n
FALSE	51.111	23
TRUE	48.889	22
Total	1.003	45

Node 84

Category	%	n
FALSE	38.095	40
TRUE	61.905	65
Total	2.341	105

Node 85

Category	%	n
FALSE	24.828	36
TRUE	75.172	109
Total	3.232	145

Variables included in this model

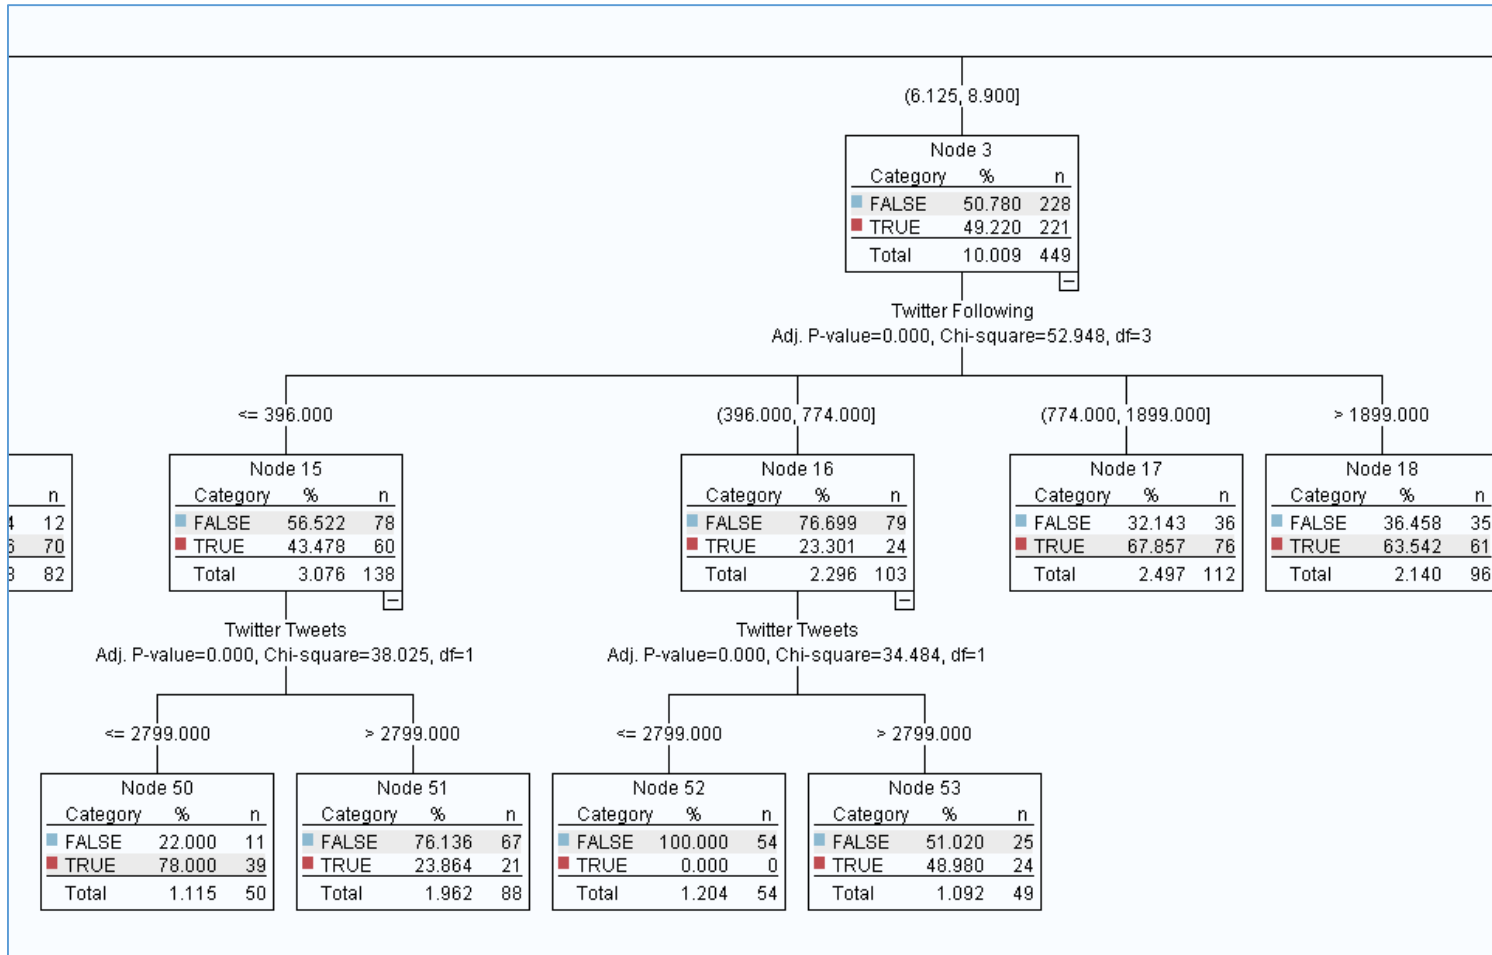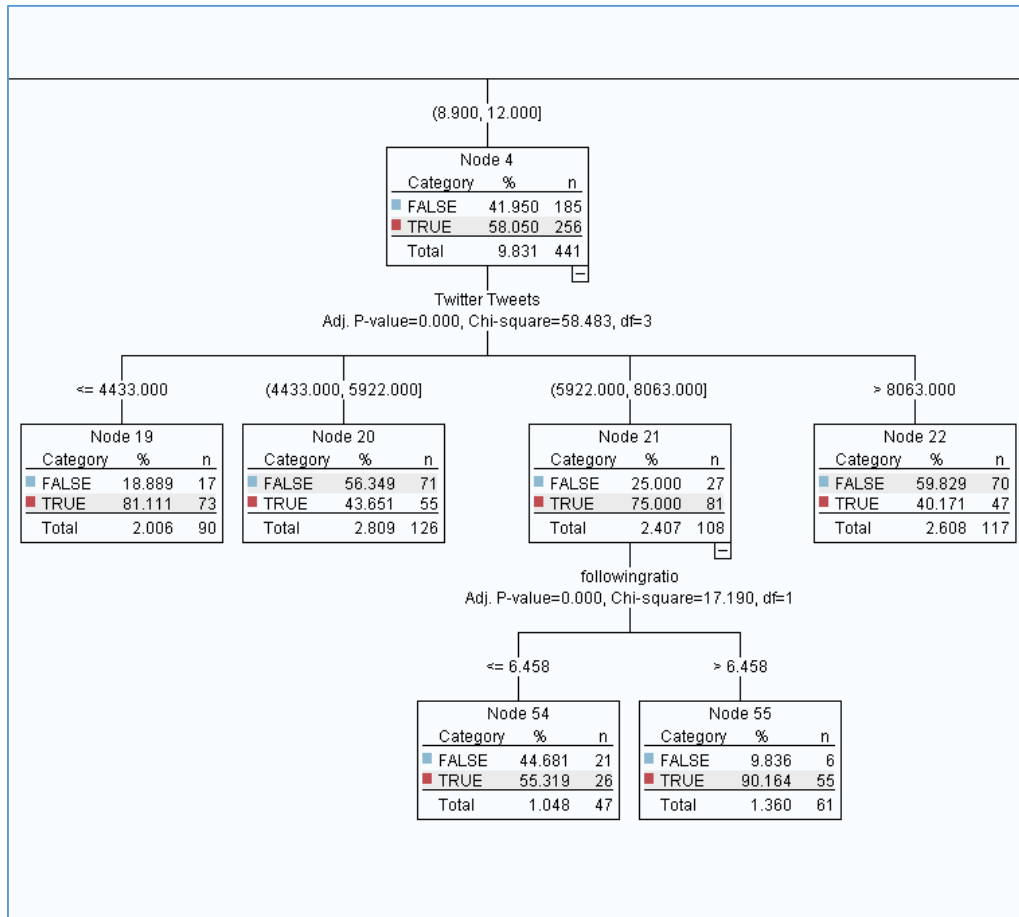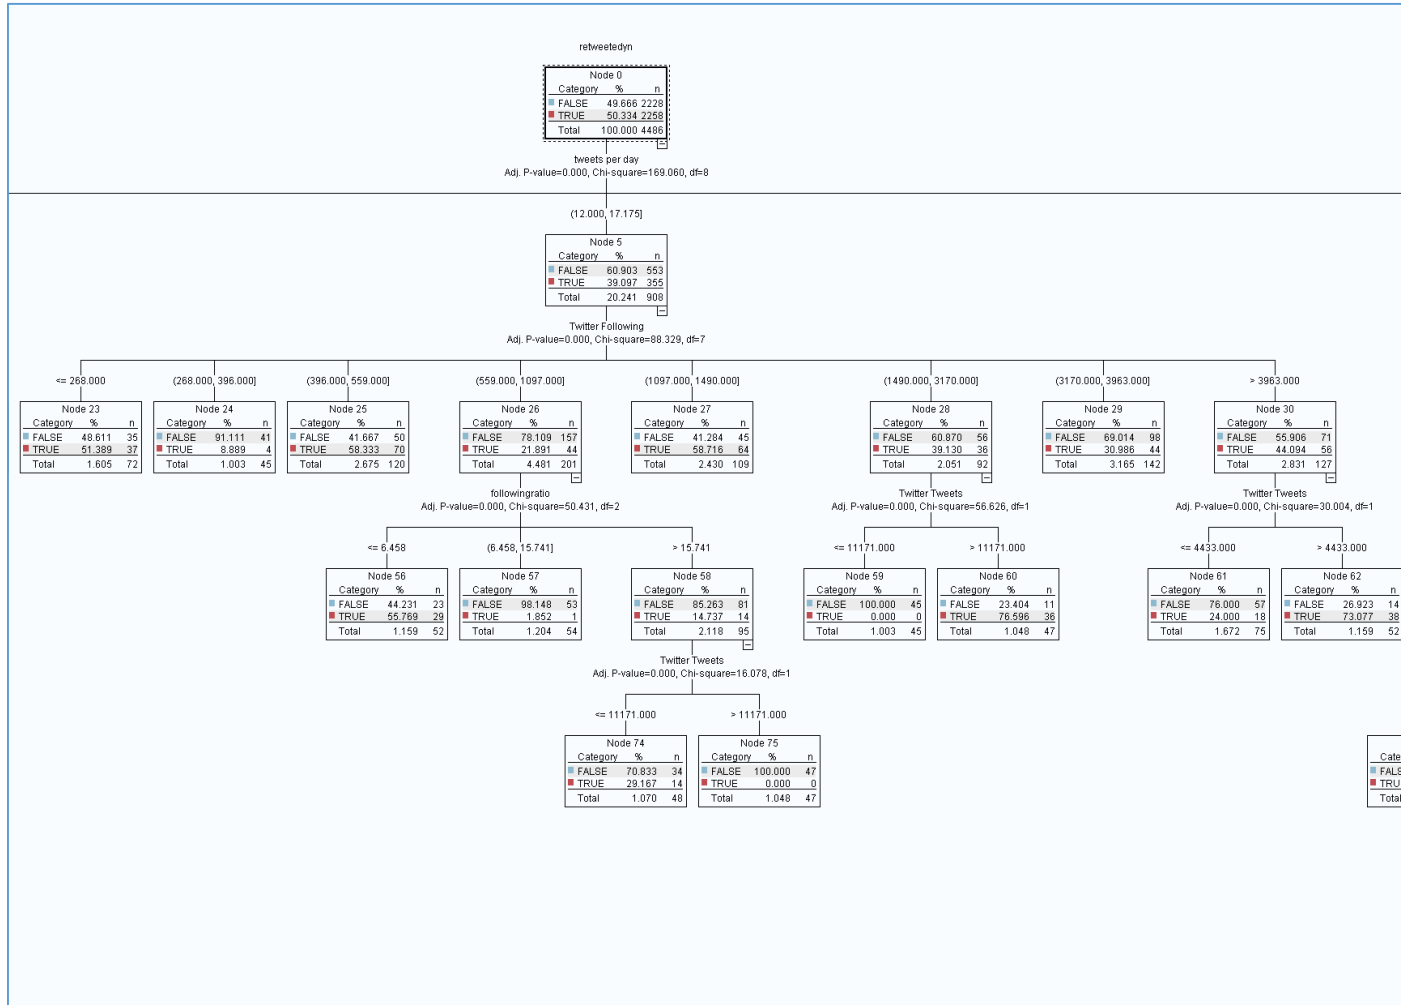**Total campaign tweets** – number of tweets MP sent during the campaign

Tweets per day – the number of tweets sent per day of the campaign

Followingratio – number of followers divided by number of followees

Twitter following – number of people the MP follows

Twitter tweets – total number of tweets sent by the MP in their time on Twitter

Medialinkyn – whether the tweet contains a link to an image or video yes/no

Hashtagyn – whether the tweet contains any hashtags yes/no

Mentionyn – whether the tweet contains any mentions or not

As Figure 68 shows, the single most important factor in this model is the number of tweets that one sends during the campaign, followed by the number of people that one follows, suggesting that to a large extent how one behaves on Twitter is more important than what one says. Of course, whilst the number of followers has been removed from this model, some of these variables are likely to be very closely related to the number of followers. For example, people who have sent a very large number of tweets probably have more followers than people who have only sent a very small number of tweets, so Twitter tweets may be acting as a proxy for follower numbers in this model. The full set of rules governing this model can be seen in appendix six.

Figure 68 - Model three: relative importance of author and tweet structural variables
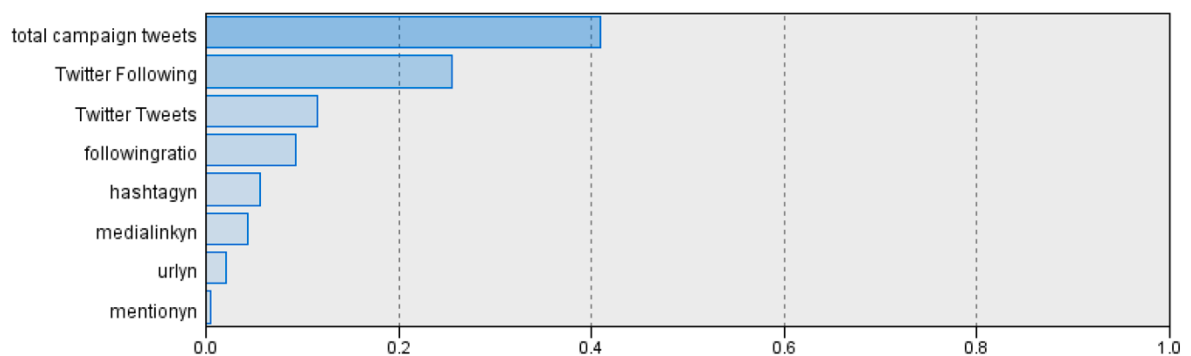
Table 24 - Evaluation of confidence scores for model three

	Training data	**Testing data**
Range	0.512-0.994	0.512-0.994
Mean correct	0.743	0.746
Mean incorrect	0.637	0.642
Always correct above	0.97 (10.32% of cases)	0.97 (10.52% of cases)
Always incorrect below	0.512 (0% of cases)	0.512 (0% of cases)
90.74% accuracy above	0.778	0.796

In this case, the area under the curve (as shown in Figure 69) is 0.799 for the training data and 0.783 for the testing data.

Figure 69 - Gains chart for model three combining author and tweet variables
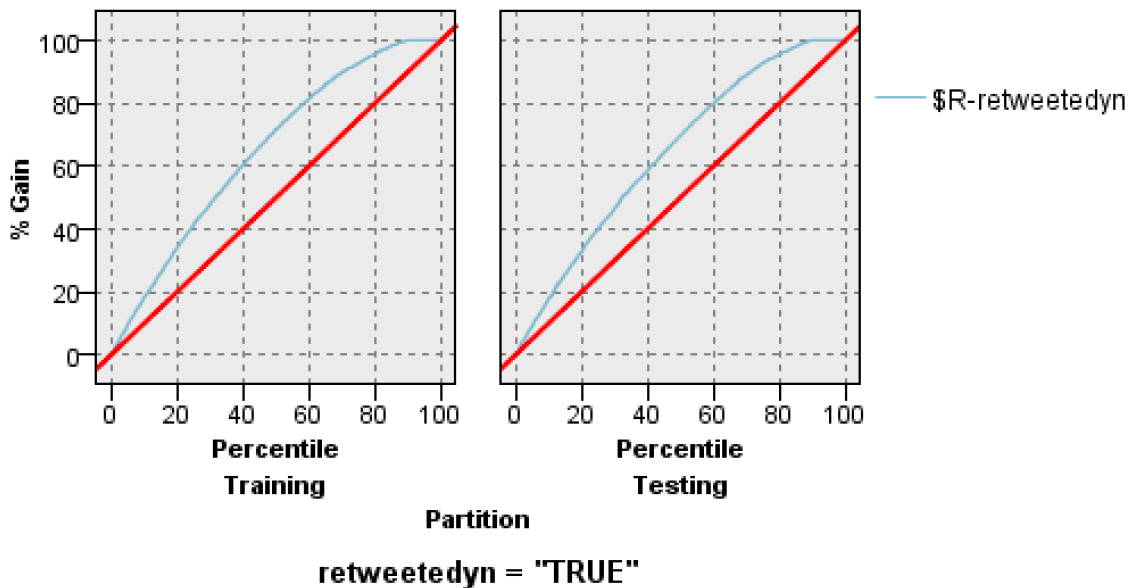
As with model two, there are too many nodes to meaningfully present each one in a gains table so Table 25 shows the nodes divided into percentiles. The nodes in the tenth percentile are 80% more likely to include retweets than in the dataset as a whole.

Table 25 - Gains table for model three combining author and tweet variables

Nodes	Percentile	Percentile: n	Gain: n	Gain (%)	Response (%)	Index (%)
40,74,57,33,42...	10.00	449.00	411.00	18.19	91.46	181.71
95,77,14,50,55...	20.00	897.00	777.00	34.40	86.60	172.04
85,70,72,36,56...	30.00	1346.00	1103.00	48.85	81.94	162.80
87,94,11,18,78...	40.00	1794.00	1392.00	61.64	77.58	154.13
84,79,97,27,54...	50.00	2243.00	1657.00	73.38	73.87	146.77
64,58,23,89,62...	60.00	2692.00	1880.00	83.28	69.85	138.78
20,80,96,99,29	70.00	3140.00	2056.00	91.05	65.48	130.08
29,60,82,90,47...	80.00	3589.00	2183.00	96.68	60.82	120.84
75,69,51,48,24...	90.00	4037.00	2258.00	100.00	55.93	111.12
37,46,35,52,91...	100.00	4486.00	2258.00	100.00	50.33	100.00

6.5.4. Model four: CHAID model using hashtags

The next stage of the analysis was to determine whether the content of the tweet influences whether or not the tweet will be retweeted. The sample of 13,020 tweets includes some basic content-related information. For example, variables were created for the most popular hashtags (#votelabour, #labourdoorstep, #voteconservative, #GE2015, #BBCQT, #leadersdebate and #bbcdebate). Additionally, Brandwatch provides a sentiment score for each tweet (positive, negative or neutral). These variables were fed into a CHAID model, along with the two other hashtag variables – whether the tweet contains hashtags or not, and the number of hashtags it contains. As can be seen in Figure 70, the fact of including a hashtag at all is substantially more predictive than the specific hashtag one chooses.

Figure 70 - Most predictive variables in model four
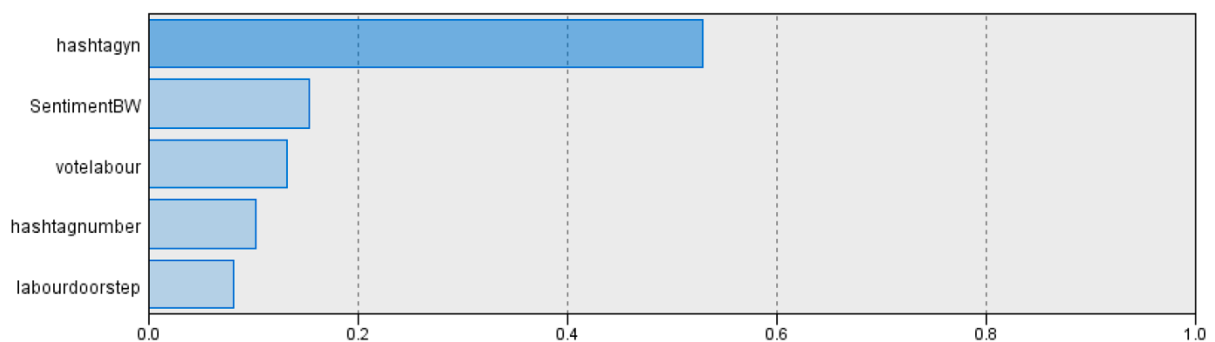
These variables produce a relatively straightforward model (Figure 71) which is simple to interpret[27]. The most importance decision is whether to include a hashtag or not. In this election Labour hashtags got more retweets than Conservative hashtags. If one does not include a hashtag then the most important predictor from this selection is sentiment, with tweets that Brandwatch deems to be positive less likely to be retweeted than those it considers to be negative or neutral. The impact of sentiment on retweeting will be considered in more detail later in this chapter.

---

[27] A guide to interpreting CHAID output is provided in Figure 49 on page 190

Figure 71 - Model four using hashtags and sentiment

Variables included in this model

Hashtagyn – does the tweet contain any hashtags yes/no?

SentimentBW – the Brandwatch sentiment score for the tweet (positive, negative or neutral)

Votelabour – does the tweet contain the hashtag #votelabour or not?

Labourdoorstep – does the tweet contain the hashtag #labourdoorstep or not?

Hashtagnumber – the absolute number of hashtags in the tweet

Despite extremely high retweeting levels for some of the hashtags the model itself was only accurate 57.53% of the time (Table 26).

Table 26 - Comparison of performance on training and testing data for model four

	Training		Testing	
Correct	3,687	57.3%	3,789	57.53%
Wrong	2,747	42.7%	2,797	42.47%
Total	6,434		6,586	

The area under the curve for model four is 0.592 on the training data and 0.594 on the testing data (Figure 68).

Figure 72 - Gains chart for model four
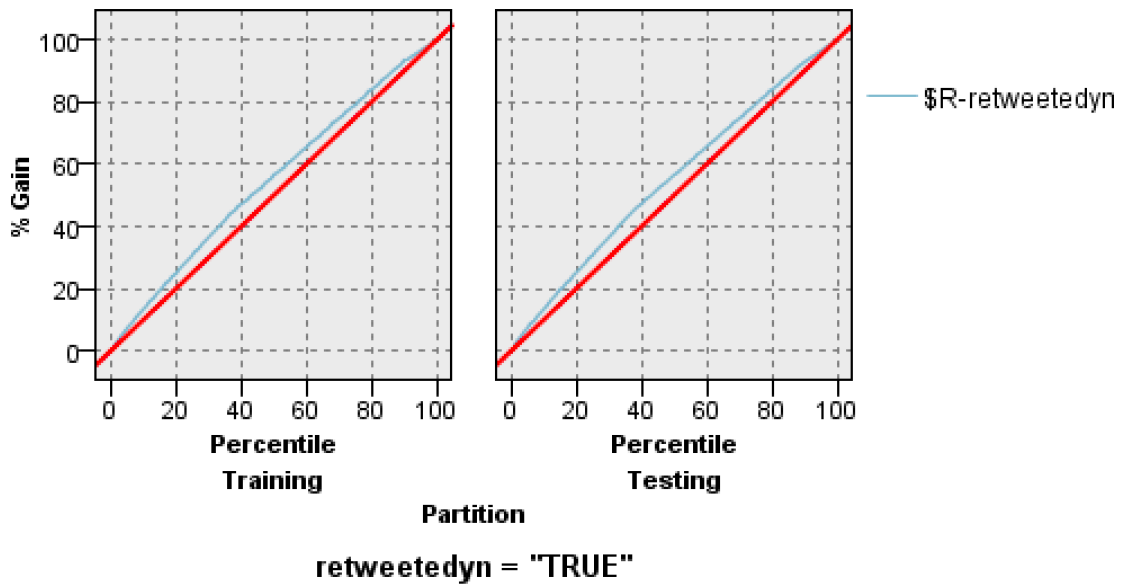
Table 27 - Evaluation of confidence scores for model four

	Training data	Testing data
Range	0.539-0.802	0.539-0.802
Mean correct	0.578	0.578
Mean incorrect	0.566	0.567
Always correct above	0.802(0% of cases)	0.802(0% of cases)
Always incorrect below	0.539 (0% of cases)	0.539 (0% of cases)
90% accuracy above	Never reached	Never reached

Table 28 - Gains table for model four

Nodes	Node: n	Node (%)	Gain: n	Gain (%)	Response (%)	Index (%)
6	67.00	1.49	51.00	2.26	76.12	151.23
10	86.00	1.92	61.00	2.70	70.93	140.92
8	176.00	3.92	122.00	5.40	69.32	137.72
9	1303.00	29.05	758.00	33.57	58.17	115.57
3	2397.00	53.43	1109.00	49.11	46.27	91.92
4	457.00	10.19	157.00	6.95	34.35	68.25

This model is little better than random chance because of the high number of tweets not containing any of these hashtags, as the model had no basis on which to make a prediction in these cases. Additionally, this model only considers a small selection of the possible hashtags that MPs included in their tweets. There are many other hashtags that could also significantly impact retweeting which further analysis could uncover. An examination of the percentage of tweets containing these hashtags getting retweeted clearly shows that including an appropriate hashtag in a tweet can almost guarantee that it will get retweeted (as previously discussed in section 6.4.1). For example, 80% of the tweets containing the #votelabour hashtag got retweeted (node 6) (The full rule set for model four can be seen in appendix six). Thus, it is clear that content plays an important role in determining whether tweets get retweeted but a more sophisticated set of content-related variables is needed in order to better understand it. The process of generating these variables and using them for modelling is explained in the sections that follow.

6.5.5.   Model five: CHAID model of machine-generated content concepts

The first stage of the content analysis was a machine-based generation of content and sentiment concepts and categories. The sample of 13,020 tweets was loaded into SPSS Text Analytics, which used its pre-set dictionaries to identify concepts within the data. The only intervention from the researcher was to specify that concepts should only be created if they occurred more than 10 times in the data, done to prevent an unmanageable number of concepts being generated. Concepts are themes that the software identifies in the data based on the occurrences of similar words in the content. The automated concept creation settings come with a default limit of 100 concepts. The researcher looked through these and manually removed ten which were not relevant (e.g. 'Twitter', 'fb', 'pic', 'bit' and so on, all of which are parts of URLs rather than meaningful concepts in their own right). A CHAID model (model five shown in Figure 73) was then run using the remaining 90 concepts as predictors[28]. This

---

[28] A guide to interpreting CHAID output is provided in Figure 49 on page 190

was correct in its predictions 55.32% of the time on the training data and 55.5% on the testing data so just slightly better than random chance (Table 29).

Table 29 - Comparison of performance on training and testing data for model five

	Training		Testing	
Correct	3,559	55.32%	3,655	55.5%
Wrong	2,875	44.68%	2,931	44.5%
Total	6,434		6,586	

Figure 73 - Model five: machine-generated content-related concept variables

Variables included in this model

Concept_labour – tweets flagged as being in the SPSS concept 'labour'

Concept_check – tweets flagged as being in the SPSS concept 'check'

Concept_tories – tweets flagged as being in the SPSS concept 'tories'

Concept_@uklabour – tweets flagged as being in the SPSS concept '@uklabour'

Concept_thanks – tweets flagged as being in the SPSS concept 'thanks'

All the variables in this model are based on the concepts determined by SPSS Text Analytics' automated content analysis.

retweetedyn

Node 0		
Category	%	n
■ N	50.047	3220
■ Y	49.953	3214
Total	100.000	6434

Concept_labor
Adj. P-value=0.000, Chi-square=99.822, df=1

F / T

Node 1		
Category	%	n
■ N	52.049	3035
■ Y	47.951	2796
Total	90.628	5831

Node 2		
Category	%	n
■ N	30.680	185
■ Y	69.320	418
Total	9.372	603

Concept_check
Adj. P-value=0.000, Chi-square=71.968, df=1

F / T

Node 3		
Category	%	n
■ N	51.358	2949
■ Y	48.642	2793
Total	89.245	5742

Node 4		
Category	%	n
■ N	96.629	86
■ Y	3.371	3
Total	1.383	89

Concept_tories
Adj. P-value=0.000, Chi-square=32.598, df=1

F / T

Node 5		
Category	%	n
■ N	51.892	2921
■ Y	48.108	2708
Total	87.488	5629

Node 6		
Category	%	n
■ N	24.779	28
■ Y	75.221	85
Total	1.756	113

Concept_@uklabour
Adj. P-value=0.000, Chi-square=30.764, df=1

F / T

Node 7		
Category	%	n
■ N	52.365	2900
■ Y	47.635	2638
Total	86.074	5538

Node 8		
Category	%	n
■ N	23.077	21
■ Y	76.923	70
Total	1.414	91

Concept_thanks
Adj. P-value=0.000, Chi-square=29.140, df=1

F / T

Node 9		
Category	%	n
■ N	51.180	2560
■ Y	48.820	2442
Total	77.743	5002

Node 10		
Category	%	n
■ N	63.433	340
■ Y	36.567	196
Total	8.331	536

The area under the curve for this model is 0.58 on the training data and 0.583 on the testing data (shown visually in Figure 70). A full set of rules used in model five is shown in appendix six.

Table 30 - Evaluation of confidence scores for model five

	Training data	Testing data
Range	0.521-0.956	0.521-0.956
Mean correct	0.566	0.567
Mean incorrect	0.536	0.537
Always correct above	0.956 (0% of cases)	0.956 (0% of cases)
Always incorrect below	0.512 (0% of cases)	0.512 (0% of cases)
93.48% accuracy above	0.763	0.763

Figure 74 - Gains chart for model five
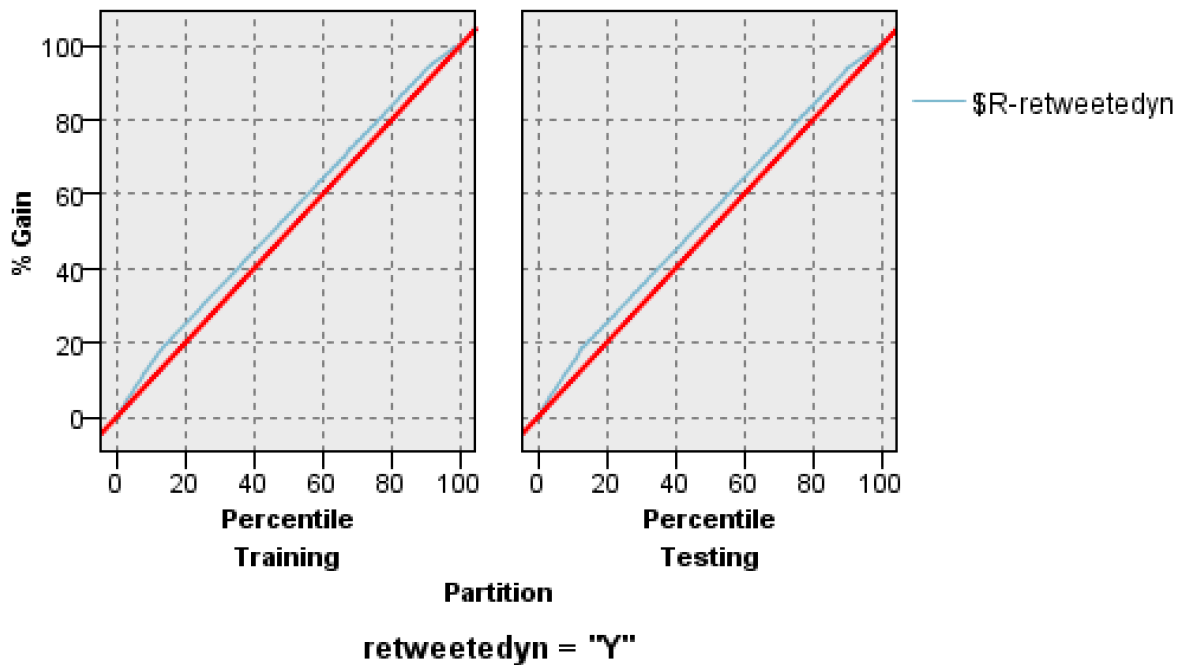
Whilst this model is not highly predictive, it does provide further evidence that there must be something in the content of the tweets that influences whether or not they are retweeted. The relatively unsophisticated way in which the computer has identified concepts in the data will limit the predictiveness of those concepts, as will the fact that many of the tweets will not contain any of the most predictive concepts thus giving the model nothing to go on when trying to predict whether those tweets will be retweeted. However, despite these limitations, the computer has identified some useful concepts that can be used for prediction. Examining CHAID model five in more depth shows five concepts are useful when predicting retweets (Figure 75) – concept_labour, concept_check, concept_thanks, concepts_tories and concept_@uklabour.

Figure 75 - Model five: most predictive machine-generated content-related variables
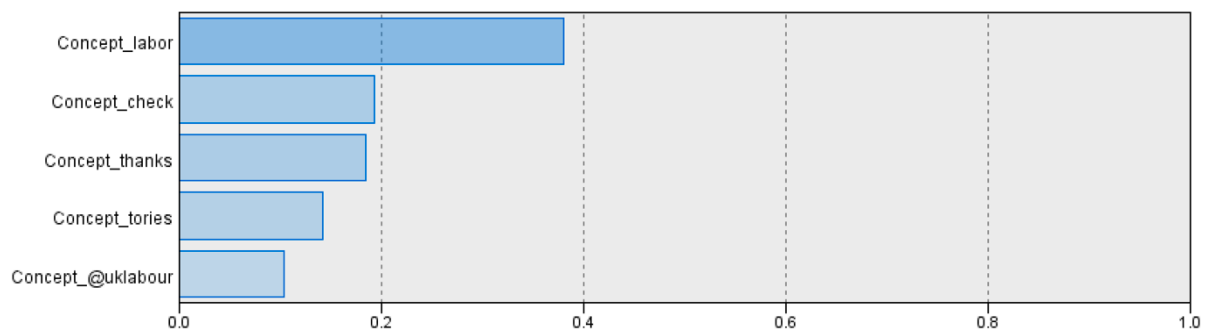
A closer examination of the CHAID model itself (Figure 73) shows the following:

- 69% of the tweets in category concept_labour were retweeted (node 2). This concept includes all tweets that mention either labor or labour.
- Only 11% of the tweets in concept_check were retweeted (node 4). Almost all the tweets in this category were automated tweets from MPs to constituents sent on polling day reminding them to vote.
- 75% of the tweets in concept_tories were retweeted (node 6). This concept includes tweets that mention the word 'tories'.

- 76% of tweets in concept @uklabour were retweeted (node 8). On the first iteration SPSS Text Analytics has treated @uklabour as if it were a word whereas actually it indicates that these are tweets in which the sender has mentioned @uklabour. The high percentage of retweets suggests that part of Labour's campaign strategy may be to have its @uklabour Twitter feed retweet its MPs' tweets whenever they mention it.

- Only 37% of tweets in concept_thanks were retweeted (node 10). There are a substantial number of underlying terms included in the concept_thanks (e.g. huge thanks, special thanks, ta, thnx, thx, thank you and many others). This result suggests that tweets thanking other people are unlikely to get retweeted.

6.5.6. Model six: CHAID model using machine-generated content categories

The next stage of the analysis was to run the concept extraction again but this time to allow SPSS Text Analytics to group the concepts into categories. Categories are more meaningful clusters of concepts that Text Analytics' dictionaries suggest are related to each other. Some additional manual cleaning was done to correct some idiosyncratic allocation of categories. For example, 'Richmond Hill' had been incorrectly classified as a person rather than a place. Such errors were corrected but no further qualitative input from the researcher was given. 30 broad categories (each including a number of sub-categories) were automatically generated based on tweet content. A CHAID model (model six - Figure 76) was built using these 30 categories as predictors[29].

---

[29] A guide to interpreting CHAID output is provided in Figure 49 on page 190

Figure 76 - Model six: machine-generated content categories

Variables included in this model

Category_labour – tweets flagged as being in the SPSS category 'labour' (T = true, F = false)

Category_vote – tweets flagged as being in the SPSS category 'vote'

Category_election – tweets flagged as being in the SPSS category 'election'

Category_people – tweets flagged as being in the SPSS category 'people'

Category_finance/tax – tweets flagged as being in the SPSS category 'finance/tax'

All the variables in this model are based on the categories determined by SPSS Text Analytics' automated content

retweetedyn

Node 0		
Category	%	n
N	49.666	2228
Y	50.334	2258
Total	100.000	4486

Category_labor
Adj. P-value=0.000, Chi-square=61.771, df=1

F ———— T

Node 1		
Category	%	n
N	51.564	2094
Y	48.436	1967
Total	90.526	4061

Node 2		
Category	%	n
N	31.529	134
Y	68.471	291
Total	9.474	425

Category_vote
Adj. P-value=0.000, Chi-square=16.850, df=1

F ———— T

Node 3		
Category	%	n
N	52.201	2040
Y	47.799	1868
Total	87.115	3908

Node 4		
Category	%	n
N	35.294	54
Y	64.706	99
Total	3.411	153

Category_election
Adj. P-value=0.002, Chi-square=9.884, df=1

F ———— T

Node 5		
Category	%	n
N	52.678	1987
Y	47.322	1785
Total	84.084	3772

Node 6		
Category	%	n
N	38.971	53
Y	61.029	83
Total	3.032	136

Category_people
Adj. P-value=0.001, Chi-square=10.460, df=1

F ———— T

Node 7		
Category	%	n
N	53.194	1932
Y	46.806	1700
Total	80.963	3632

Node 8		
Category	%	n
N	39.286	55
Y	60.714	85
Total	3.121	140

Category_finance/tax
Adj. P-value=0.001, Chi-square=11.596, df=1

F ———— T

Node 9		
Category	%	n
N	53.574	1911
Y	46.426	1656
Total	79.514	3567

Node 10		
Category	%	n
N	32.308	21
Y	67.692	44
Total	1.449	65

Model six was accurate in its predictions 56.14% of the time on the training data and 55.36% on the testing data (Table 31). Although this model is only slightly better than random chance, it does provide some further evidence that content influences retweeting. That the model is not more predictive is due to the limitations of the categories created by the machine analysis.

Table 31 - Comparison of performance on training and testing data for model six

	Training		Testing	
Correct	3,612	56.14%	3,646	55.36%
Wrong	2,822	43.86%	2,940	44.64%
Total	6,434		6,586	

Table 32 - Evaluation of confidence scores for model six

	Training data	Testing data
Range	0.536-0.684	0.536-0.684
Mean correct	0.564	0.564
Mean incorrect	0.554	0.554
Always correct above	0.684 (0% of cases)	0.684 (0% of cases)
Always incorrect below	0.536 (0% of cases)	0.536 (0% of cases)
90% accuracy above	Never reached	Never reached

The area under the curve for this model is 0.563 on the training data and 0.557 on the testing data, shown visually in Figure 77. A full set of rules for model six can be seen in appendix six.

Figure 77 - Gains chart for model six
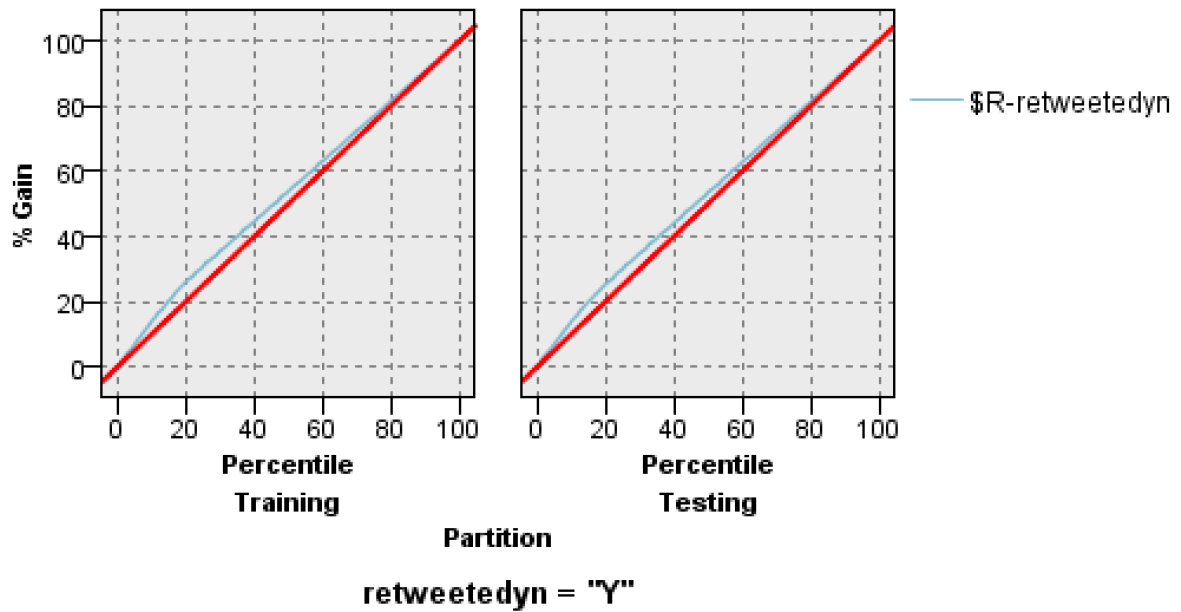
Table 33 - Gains table for model six

Nodes	Node: n	Node (%)	Gain: n	Gain (%)	Response (%)	Index (%)
2	425.00	9.47	291.00	12.89	68.47	136.03
10	65.00	1.45	44.00	1.95	67.69	134.49
4	153.00	3.41	99.00	4.38	64.71	128.55
6	136.00	3.03	83.00	3.68	61.03	121.25
8	140.00	3.12	85.00	3.76	60.71	120.62
9	3567.00	79.51	1656.00	73.34	46.43	92.23

As this model is based on computerised content analysis it has similar limitations to model five. The computer is not able to really understand the content of the tweets and so the categories are not as reliable as hand coded categories would be. Additionally, many tweets do not fall into any of the most predictive categories identified, meaning that in those cases the model has nothing to go on when making a prediction and hence the overall predictiveness of the model is low. However, individual nodes within it contain high numbers of retweeted tweets, providing further evidence that content does influence retweeting. Figure 78 shows categories that the model selected as most predictive.

Figure 78 - Model six: most important machine-generated category predictors
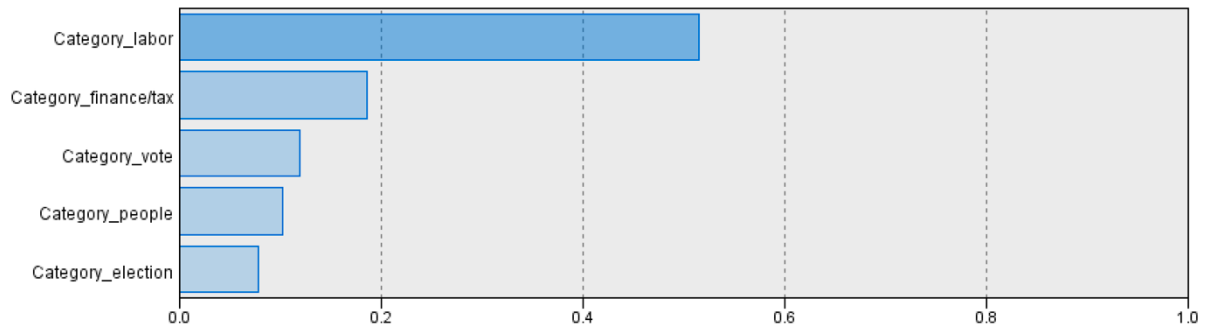
Once again category_labor (node 2) was highly predictive with 68% of these tweets being retweeted, providing further evidence suggesting that retweeting may have been part of the Labour Party's campaign strategy. The next most significant category was category_vote (node 4) with 65% of tweets being retweeted. Then came category_finance/tax (node 10) in which 68% of tweets were retweeted. In category_election (node 6) 61% of tweets were retweeted and in category_people (node 8) 61% were retweeted. All the CHAID models were set to automatically stop at five levels deep which is why only five categories were identified.

6.5.7. The influence of tweet valence

As well as machine-generated content categories, some analysis of machine-generated sentiment variables is also possible. The data collected from Brandwatch includes Brandwatch's own assessment of whether a tweet is positive, negative or neutral. How the Brandwatch sentiment algorithm works is not revealed so it is not possible to evaluate how reliable its scoring system is or determine the factors that it uses to make its sentiment assessment. However, we can see that the Brandwatch sentiment categories do have an influence on whether a tweet is retweeted or not. Table 34 shows that, whilst 58% of tweets that Brandwatch deems to be negative tweets are retweeted, only 42% of the tweets it considers positive tweets are retweeted. This suggests that further consideration of the role of sentiment in predicting retweets would be valuable.

Table 34 - Influence of Brandwatch sentiment score on retweets

		Not retweeted	Retweeted
Negative	Count	137	193
	Residual	-28.00	28.00
	Row %	41.51%	58.48%
Neutral	Count	5,258	5,513
	Residual	-127.500	127.500
	Row %	48.82%	51.18%
Positive	Count	1,115	804
	Residual	155.500	-155.500
	Row %	58.1%	41.9%

Chi-square 65.942, df 2, *p* < .001

SPSS Text Analytics also includes an automated sentiment analysis element which categorises content as either positive or negative (or neither, but there is no neutral category). Although this uses a different method to Brandwatch in order to categorise the valence of the tweets[30], the same pattern emerges with negative tweets being significantly more likely to get retweeted whilst positive tweets are less likely to be retweeted (Table 35).

Table 35 - Influence of SPSS Text Analytics' tweet valence score on retweets

		Retweeted?		Total	
		No	Yes		
Negative	Count	1,236	1,497	2,733	Chi-square 31.547
	Residual	-130.500	130.500		df 1
	Percentage	45.22%	54.77%		*p* < .001
Positive	Count	3,700	3,318	7,018	Chi-square 45.105
	Residual	191.000	-191.000		df 1
	Percentage	52.77%	47.27%		*p* < .001

There are substantial differences between Brandwatch and SPSS in terms of the volume of tweets that each considers to be negative. The Brandwatch algorithm flags

---

[30] Note, Brandwatch creates a single variable within which each tweet is coded as positive, negative or neutral. SPSS Text Analytics creates two separate flag variables – positive (yes / no) and negative (yes / no) which is why there are two chi-square results above for the SPSS sentiment variables and only one for the Brandwatch variable. Note too that the SPSS approach means that the same tweet can be categorised as both positive and negative if it includes elements of both. Using the Brandwatch approach a tweet can only be positive, negative or neutral.

only 330 tweets as being negative – just over 2.5% of the total. By contrast, SPSS Text Analytics flags just over 20% as being negative. This is a substantial difference and is another reason to be cautious of the results of machine-based sentiment analysis. If two machines are able to come up with such vastly different numbers then they cannot both be measuring the same thing, suggesting that machine-based measures of sentiment may not be valid. That said, both approaches agree that negative tweets, however they are determined, are more likely to be retweeted than positive tweets.

Overall, the machine-generated content concepts and categories provide an indication that content and sentiment play a role in the retweeting of politicians' tweets but they are not highly predictive themselves. The limitations of machine-based content analysis, as already discussed in the methods chapter, come into play here. The computer looks for patterns in the tweets and associated words but it cannot really understand what those words mean or what the intention or purpose of a particular tweet is. Only a human coder with knowledge of the relevant political context can make those judgements.

Particularly with regard to political tweets, it is very hard for a computer to accurately assess whether some of them are positive or negative due to high levels of sarcasm and the large extent to which the negative implication of a tweet is entirely contextual. For example, in the General Election campaign many Conservative MPs tweeted about the extent to which the SNP would be helping Labour to govern in the event of a Labour victory. A computer algorithm might assume that any tweet including the word 'helping' is positive or would, best case scenario, determine that such a tweet did not have a strong valence either way whereas a human with knowledge of the relevant political context would clearly determine that such a tweet was intended to be negative.

Thus the final step of the research was to manually code the data in order the determine whether manually-generated content and sentiment variables could play a role in predicting which tweets get retweeted.

6.6. Phase three - manual content and sentiment analysis

6.6.1. Manual sentiment analysis

A random sample of ten percent of the tweets being analysed was selected, comprising a subset of 1,212 tweets to be manually coded. The balance between retweets and non-retweets was maintained at 50/50. The researcher considered each tweet in turn and decided whether it was positive, negative or neutral. No information about whether each tweet was retweeted or not was available to the researcher, to avoid this knowledge biasing the coding.

Positive tweets were deemed to be any which included some positive sentiment such as a statement of achievement, thanking people, offering support to someone else, reporting good news and so on. Negative tweets were those which were clearly intended as criticism of the other side, planting fear or concern in people's minds regarding what might happen if the other side won, use of sarcasm to imply criticism. Neutral tweets were those which did not include any emotional element, for example simply passing on information or describing something that had happened without any additional commentary.

It was decided that each tweet would only have one code (in common with the Brandwatch approach but a variation on the approach taken by SPSS). Where a tweet contained both positive and negative sentiment the researcher decided which element was dominant within the tweet and coded accordingly.

Out of 1,212 tweets, the researcher coded 17% as negative, 47% as neutral and 36% as positive. In comparison, the Brandwatch algorithm identified just over 2.5% of the tweets it analysed as being negative and SPSS Text Analytics determined that 21% of the tweets were negative. As shown above both the Brandwatch and Text Analytics scores were slightly predictive of retweeting however the numbers suggest that Brandwatch is being too conservative in its assessment of negativity whilst SPSS Text Analytics may be over-counting.

The manually coded valence of the tweets has a substantial impact on how likely they are to be retweeted (Figure 79 and Table 36). Almost 80% of the negative tweets are

retweeted, compared to just over 50% of the positive ones. For neutral tweets the retweet rate falls to just under 40%. These differences are significant at the 0.05 level showing that negative tweets are significantly more likely to be retweeted than are either positive or neutral ones.

Figure 79 - Effect of manually coded sentiment variables on retweeting
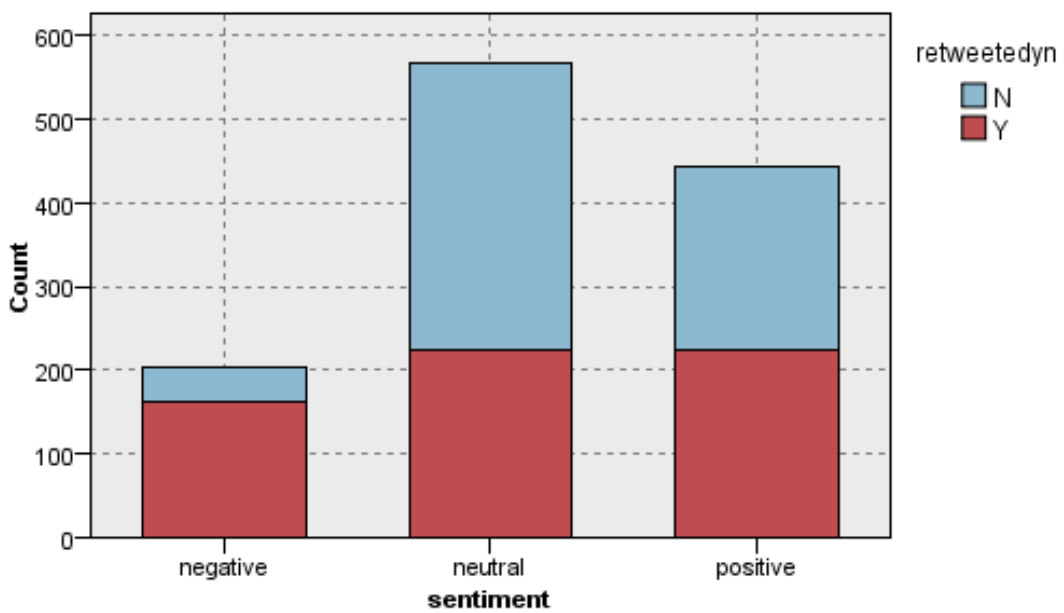
Table 36 - Effect of manually coded tweet valence on retweeting

	Not retweeted	Retweeted	Total
Negative	41 -59.822 20.098%	163 59.822 79.8%	204
Neutral	341 61.269 60.247%	225 -61.269 39.753	566
Positive	217 -1.447 49.095%	225 1.447 50.905%	442

Chi-square 96.731, df 2, $p < .000$

Although this research is not attempting to predict the volume of retweets, it is still interesting to see if the sentiment of the tweet influences how many times it gets retweeted. The median retweet numbers for each manually coded sentiment category were compared[31]. Positive tweets got a median of five retweets (mean 17), neutral tweets got a median of four retweets (mean 9.71) and negative tweets got a median of 11 retweets (mean 125.71). Retweet numbers follow a non-normal distribution and so a non-parametric comparison of independent medians was conducted which showed the differences between these three groups to be significant (Figure 80).

Figure 80 - Comparison of median retweets by sentiment including outlier

Total N	607
Median	5.000
Test Statistic	31.334
Degrees of Freedom	2
Asymptotic Sig. (2-sided test)	.000

---
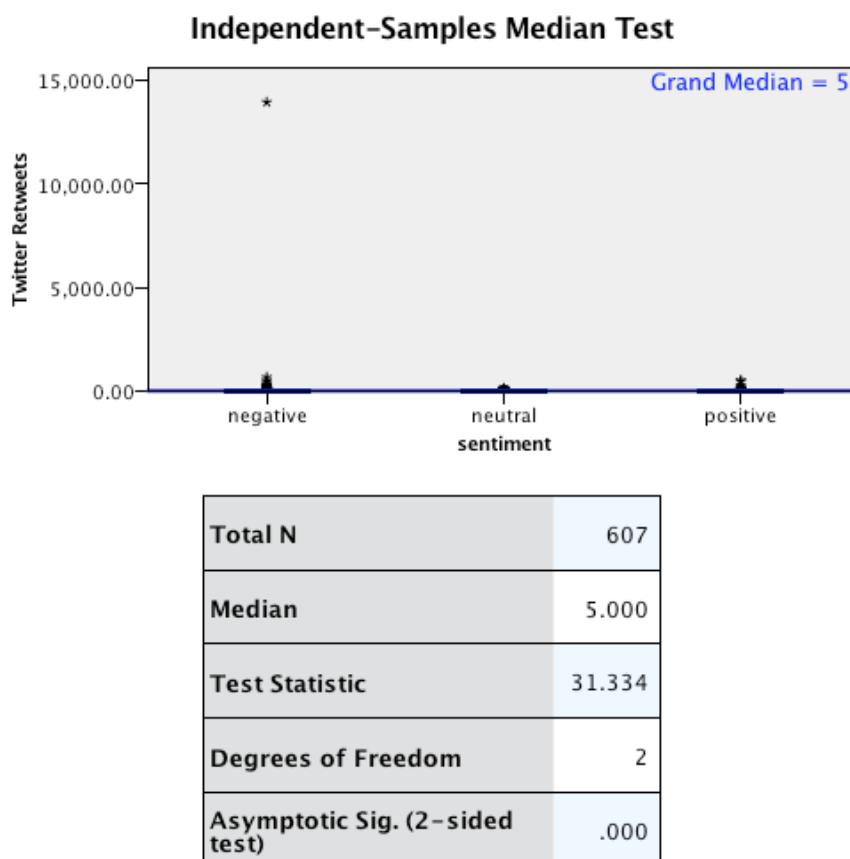
[31] This analysis was performed on only the 50% of the tweets in the dataset that got retweeted at least once, as leaving in the 50% that did not get retweeted pulls the mean and median numbers of retweets down artificially.

This test shows one tweet with vastly more retweets than any other (Figure 80). This is Ed Miliband's tweet about David Cameron's refusal to take part in the leader's debate, which was retweeted more than 13,000 times and is the single most retweeted tweet in the entire dataset. It was decided to remove this tweet from the sample and compare the medians once again to make sure that any differences were not because of this single tweet. The medians did not change at all and the test results were still significant (Figure 81), indicating that tweet sentiment has an impact not only on the chances of getting a retweet but on the volume of retweets one gets.

Figure 81 - Comparison of median retweets by tweet sentiment excluding outlier

Independent-Samples Median Test

Total N	606
Median	5.000
Test Statistic	30.774
Degrees of Freedom	2
Asymptotic Sig. (2-sided test)	.000
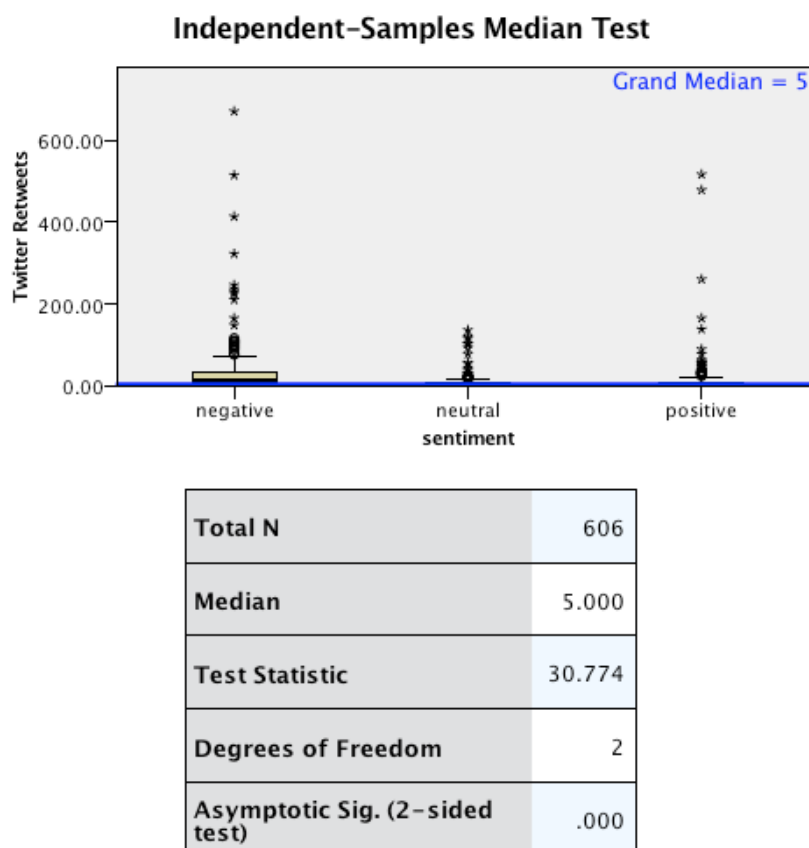
This analysis provides further evidence that the manually coded sentiment categories are robust. They are clearly providing a measure of some kind of meaningful difference between the various tweets.

6.6.2. Manual content analysis

The computer-based content categories were based on the appearance of certain words and phrases in tweets so the categories largely represented the topic of the tweet rather than its purpose. It is beyond the power of a text analytics computer to be able to determine tweet purpose so coding for that must be done manually. The manual codes were applied to the same sample of 1,212 tweets used for the manual sentiment analysis.

A detailed discussion of how the coding schema was developed was given in section 5.8.5 on page 125 and the coding schema itself can be seen in appendix five. The codes are neither mutually exclusive nor are they prioritised in any form. As many codes as were relevant were applied to each tweet. The frequency of each code is shown in Figure 82. Unsurprisingly, given that the majority of the MPs are standing for re-election, local mentions appear with by far the most frequency, followed by mentions of the campaign trail, attack tweets, personal tweets and thank you tweets.

Figure 82 - Frequency of manual codes
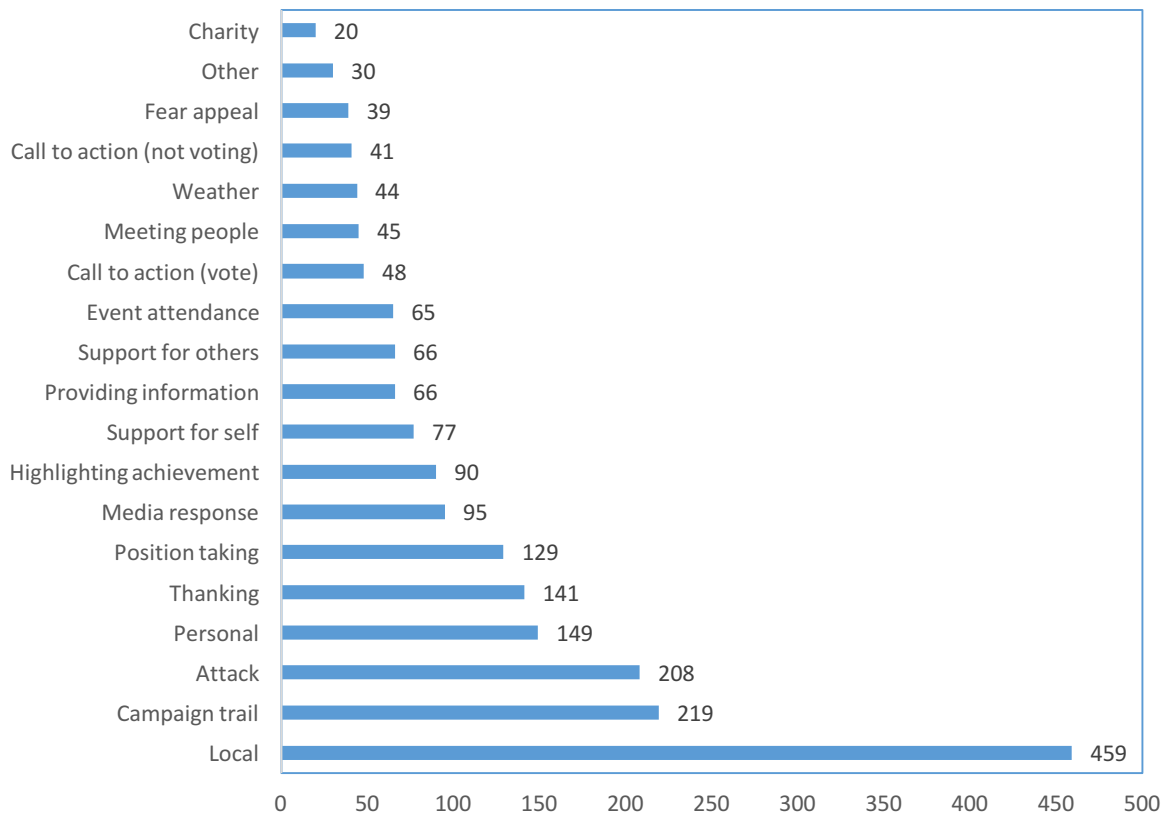
Chi-square tests were run on each of the coded categories to determine whether there were significant relationships between the number of tweets in each category that got retweeted. Eight out of the nineteen categories showed significant results, meaning that the content of the tweet influenced whether or not that tweet was retweeted. These eight categories are shown in Table 38. The results of the chi-square tests for the full list of variables are shown in appendix seven. Note, not all the relationships are positive – some content categories are negatively associated with retweeting meaning tweets of that kind are significantly less likely to get retweeted.

Table 37 - Significance of content categories in determining retweets

Variables positively associated with retweeting

		Retweeted?		
		No	**Yes**	
Fear appeal	Count	4	35	Chi Square 24.729
	Residuals	-3.5	3.5	df 1
	Percentage	10.3%	89.7%	Sig (two tailed) < .001
Support for others	Count	15	51	Chi Square 19.900
	Residuals	-3.1	3.1	df 1
	Percentage	22.7%	77.3%	Sig (two tailed) <.001
Attack	Count	55	153	Chi Square 53.046
	Residuals	-4.7	4.7	df 1
	Percentage	26.4%	73.6%	Sig (two tailed) < .001
Media response	Count	29	66	Chi Square 14.742
	Residuals	-2.6	2.6	df 1
	Percentage	30.5%	69.5%	Sig (two tailed) <.001
Position taking	Count	50	79	Chi Square 6.566
	Residuals	-1.7	1.7	df 1
	Percentage	38.8%	61.2%	Sig (two tailed) .010
Campaign trail	Count	86	133	Chi Square 11.023
	Residuals	-2.1	2.1	df 1
	Percentage	39.3%	60.7%	Sig (two tailed) .001

Variables negatively associated with retweeting

		Retweeted?		
		No	**Yes**	
Event	Count	40	25	Chi Square 4.034
	Residuals	1.4	-1.4	df 1
	Percentage	61.5%	38.5%	Sig (two tailed) .045
Personal	Count	103	46	Chi Square 26.389
	Residuals	3.4	-3.4	df 1
	Percentage	69.1%	30.9%	Sig (two tailed) <.001

As with the manual sentiment categories, so with the significant content categories an additional analysis was run to determine whether the median number of retweets differed significantly. This analysis was run excluding Ed Miliband's outlier tweet and those tweets that were not retweeted at all. The results can be seen in Table 38.

Table 38 - Comparison of median retweets by type of tweet content

	Median retweets by content category			
	In category	Not in category	*p*	Sig?
Attack	10	5	< .001	Y
Media response	12.5	5	.004	Y
Support for others	5	5	.992	N
Fear appeals	11	5	.012	Y
Position taking	8	5	.022	Y
Campaign trail	4	6	.073	N
Event attendance	3	6	< .001	Y
Personal	3	6	.038	Y

As Table 38 shows, the tweet's content significantly impacts both whether it gets retweeted and also how many times it gets retweeted. Media response tweets get the highest number of median retweets, followed by the fear appeals and the attack tweets. The categories event attendance and personal both have a negative impact on the number of retweets, achieving fewer retweets than those tweets which don't mention event attendance or include a personal element.

There are 129 tweets in the dataset with more than 15 retweets (deemed by SPSS to be extreme values). An analysis was run to see how many times the different content categories were represented in these 129 tweets. This showed that 50% of the most highly retweeted tweets were negative, compared with 17% in the dataset as a whole (Table 39). Additionally, 41.9% of the most retweeted tweets are attack tweets (Table 40), indicating that sending attack tweets is both a good way of getting retweeted in the first place and also of getting a substantial volume of retweets.

Table 39 - Sentiment of the most retweeted tweets

| | | | Volume of retweets | | |
			0–15	16+	Total
sentiment	negative	Count	138	65	203
		% within mostretweeted	12.9%	50.4%	17.0%
	neutral	Count	528	25	553
		% within mostretweeted	49.5%	19.4%	46.3%
	positive	Count	400	39	439
		% within mostretweeted	37.5%	30.2%	36.7%
Total		Count	1066	129	1195
		% within mostretweeted	100.0%	100.0%	100.0%

Chi-Square Tests

	Value	df	Asymptotic Significance (2-sided)
Pearson Chi-Square	119.238[a]	2	.000
Likelihood Ratio	96.369	2	.000
N of Valid Cases	1195		

a. 0 cells (0.0%) have expected count less than 5. The minimum expected count is 21.91.

Table 40 - Content categories of most retweeted tweets

Category	% of most retweeted tweets	% in whole sample	p	Significant?
Media response	21.7%	7.9%	<.001	Y
Fear appeal	11.6%	3.3%	<.001	Y
Achievement	12.4%	7.5%	.026	Y
Information	1.6%	5.4%	.042	Y
Position taking	20.9%	10.8%	<.001	Y
Thanking	5.4%	11.8%	.018	Y
Campaign trail	7%	18.2%	<.001	Y
Event	0%	5.4%	.004	Y
Personal	4.7%	12.1%	.006	Y
Local	17.1%	38.1%	<.001	Y
Attack	41.9%	17.2%	<.001	Y

6.6.3. Model seven: CHAID model using manually coded content variables

Having determined that the manual content and sentiment variables definitely influenced retweeting, the next stage of the analysis was to use variables based on the content codes (excluding sentiment) to build another CHAID model (model seven, Figure 83)[32].

---

[32] A guide to interpreting CHAID output is provided in Figure 49 on page 190

Figure 83 - Model seven: manual content variables

Variables included in this model

Attack – is the tweet an attack tweet yes/no?

Supportforothers – does the tweet express support for others yes/no?

Campaigntrail – is the tweet an update from the campaign trail yes/no?

Fearappeal – does the tweet contain a fear appeal yes/no?

Positiontaking – does the tweet take a position on a policy issue yes/no?

Thanking – does the tweet express thanks to someone yes/no?

Charity – does the tweet mention charity yes/no?

The content variables attack, support for others, position-taking, campaign trail, charity, fear appeal and thanking were shown to be influential in model seven (Figure 84).

Figure 84 - Model seven: most predictive manual content variables
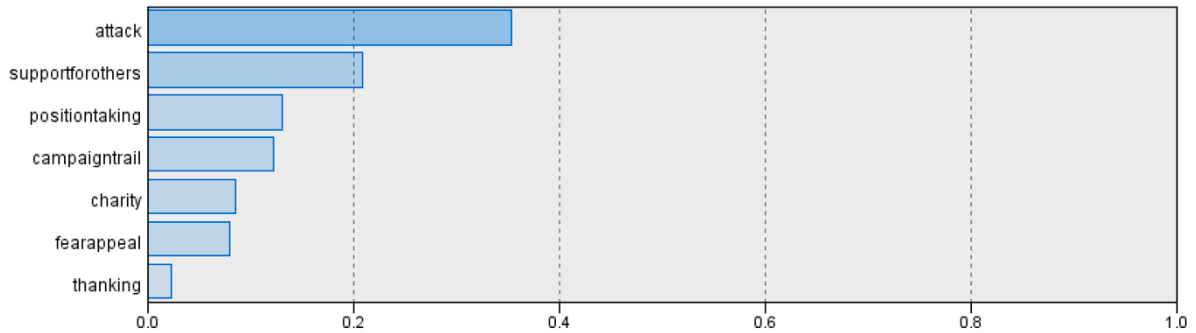
CHAID model seven (Figure 83) shows that 69.34% of attack tweets were retweeted (node 2)[33]. Within that category, 89.47% of fear appeals were retweeted (node 6). Of non-attacking tweets, those most likely to be retweeted either expressed support for others (80.85% retweeted – node 4), and more specifically those support tweets that also included some mention of the campaign trail (100% - node 10), or mentioned some aspect of the campaign and thanked other people (86.36% retweeted – node 14). Overall, this model was accurate in its predictions 64.79% of the time on the testing data and 65% on the training data. In this case, the area under the curve is 0.676 on the training data and 0.679 on the testing data (shown visually in Figure 85). The full set of CHAID decision rules for this tree can be seen in appendix six.

Table 41 - Comparison of performance on training and testing data for model seven

	Training		Testing	
Correct	541	64.79%	234	65%
Wrong	294	35.21%	126	35%
Total	835		360	

---

[33] Note, these numbers differ slightly from the numbers shown in Table 38 because the numbers in the table are based on the whole sample of manually coded tweets whereas the numbers in the CHAID model are based on a partitioned testing sample. For this phase of the analysis the cases were split 70/30 into training and testing samples.

Table 42 - Evaluation of confidence scores for model seven

	Training data	**Testing data**
Range	0.544-0.947	0.544-0.947
Mean correct	0.653	0.655
Mean incorrect	0.626	0.638
Always correct above	0.857 (2.04% of cases)	0.947 (0% of cases)
Always incorrect below	0.544 (0% of cases)	0.544 (0% of cases)
91.38% accuracy above	0.688	Never reached

Figure 85 - Gains chart for model seven
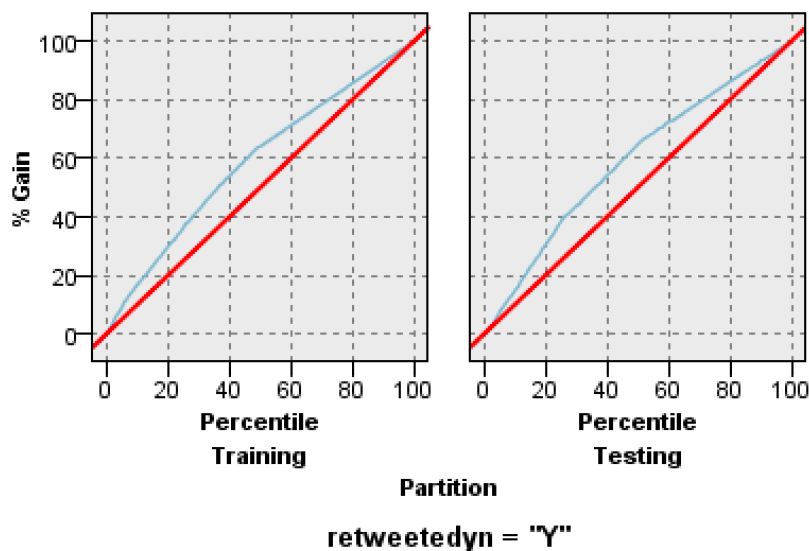
6.6.4. Model eight: CHAID model using manual content and sentiment
variables

Adding sentiment into the model (model eight – Figure 86) changed things slightly,
with sentiment becoming the most important predictor (Figure 87) as 77.44% of
negative tweets were retweeted. Both the attack and fear appeal categories
disappeared from the model, most likely because they would all be included within
the negative sentiment category and negative sentiment is predictive enough on its

own. For positive tweets to be retweeted the best thing an MP can do is call to vote (100% retweeted – node 7 Figure 86) or express support for others including a mention of the campaign trail (100% retweeted – node 15). Mentions of the campaign trail or support for others also boost the chances of neutral tweets being retweeted. For negative tweets it is enough just to be negative – no additional content is required as 77.4% of all negative tweets are retweeted (node 1).

Figure 86 - Model eight: manual sentiment and content variables
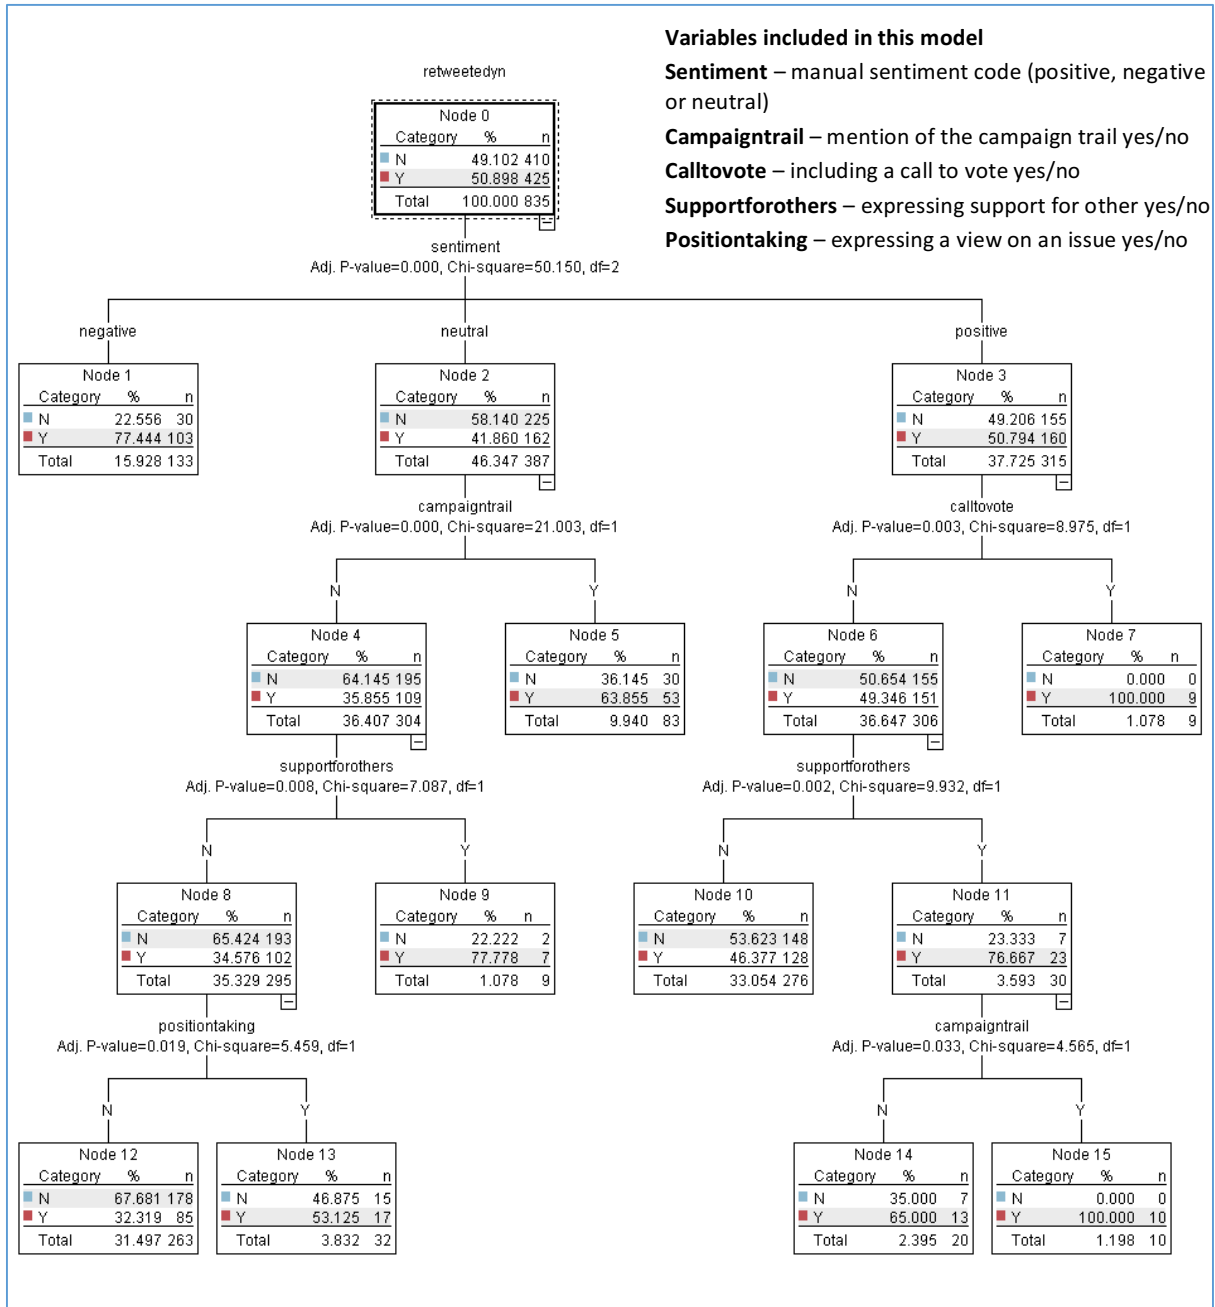
Figure 87 - Model eight: most predictive manual content and sentiment variables
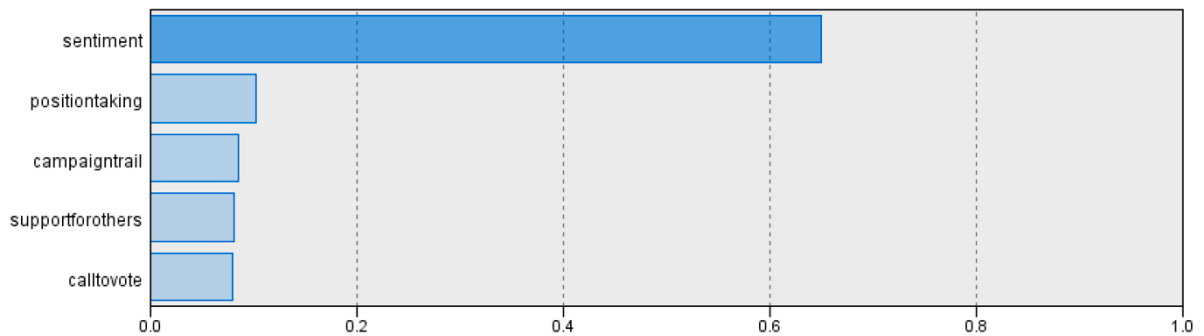
Including sentiment in the model did not improve the predictive power of the model which in fact dropped slightly to 64.43% on the training data and 63.31% on the testing data (Table 44). A full set of decision rules for this model can be seen in appendix six. In this case, the area under the curve was 0.691 for the training data and 0.695 for the testing data (Figure 88).

Table 43 - Comparison of performance on training and testing data for model eight

	Training		Testing	
Correct	538	64.43%	229	63.31%
Wrong	297	35.57%	131	36.39%
Total	835		360	

Table 44 - Evaluation of confidence scores for model eight

	Training data	Testing data
Range	0.529-0.917	0.529-0.917
Mean correct	0.655	0.669
Mean incorrect	0.613	0.622
Always correct above	0.77 (2.28% of cases)	0.917 (0% of cases)
Always incorrect below	0.529 (0% of cases)	0.529 (0% of cases)
100% accuracy above	0.77	Never reached

Figure 88 - Gains chart for model eight
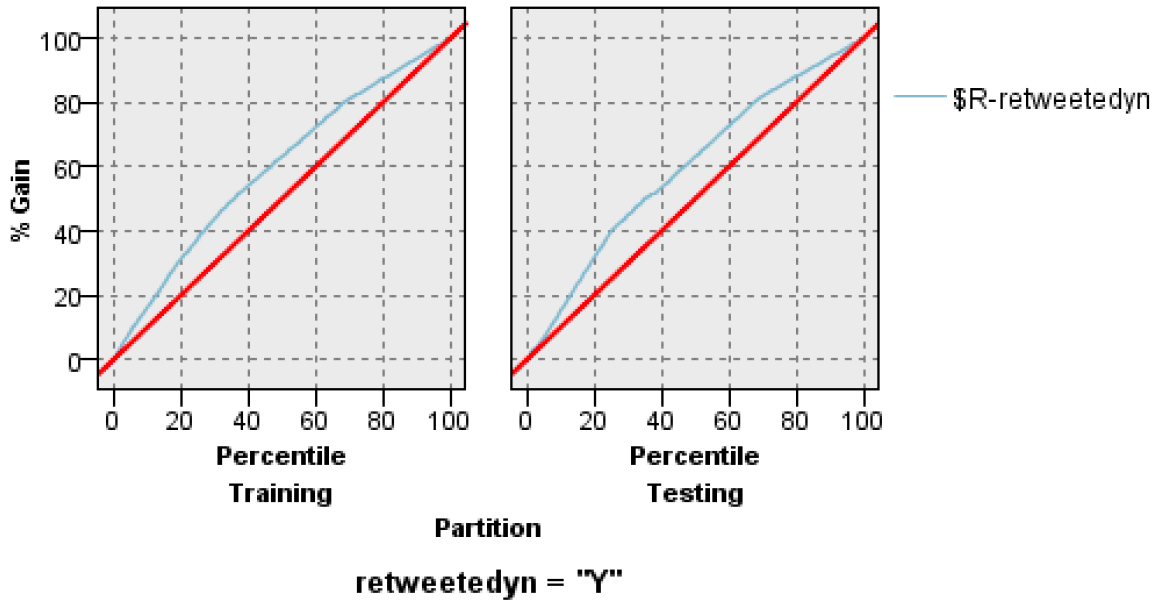
6.6.5.   Model nine: CHAID model blending all variables together

Another model (model nine - Figure 89) was constructed which blended together all the variables deemed predictive by the models built so far – the Twitter structural data, the senders' data and the manual content and sentiment-related data – in order to better understand how these factors interact. Model nine is too large to be easily viewed on one page so has been split into two – see Figure 90 and Figure 91[34].

---

[34] A guide to interpreting CHAID output is provided in Figure 49 on page 190

Figure 89 - Model nine: blended model
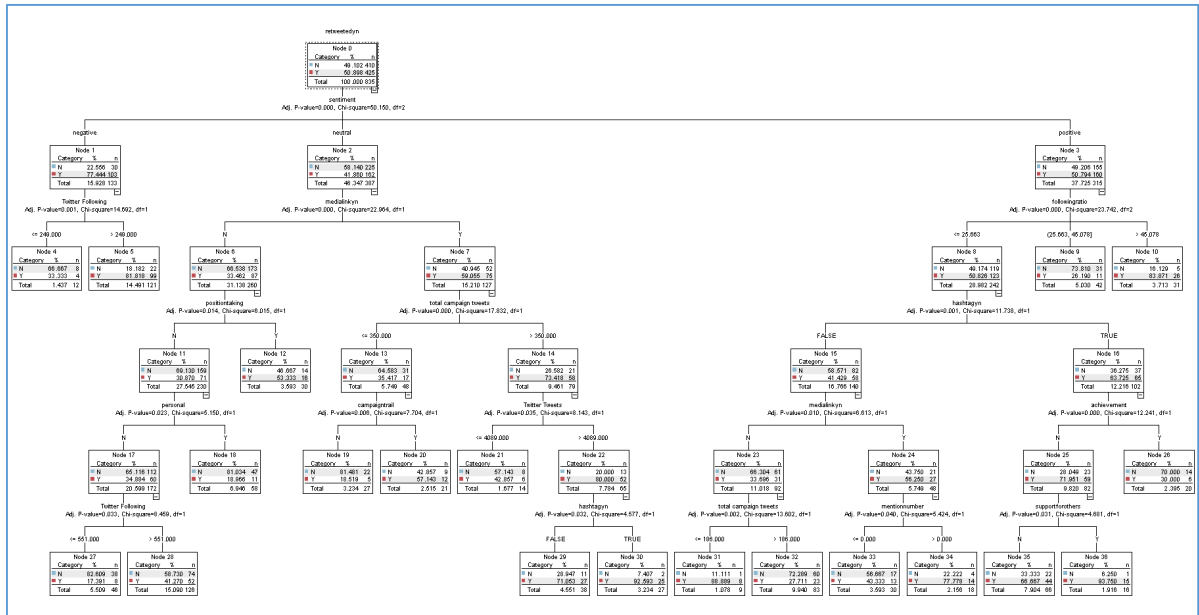
Figure 90 - Model nine: blended model part one
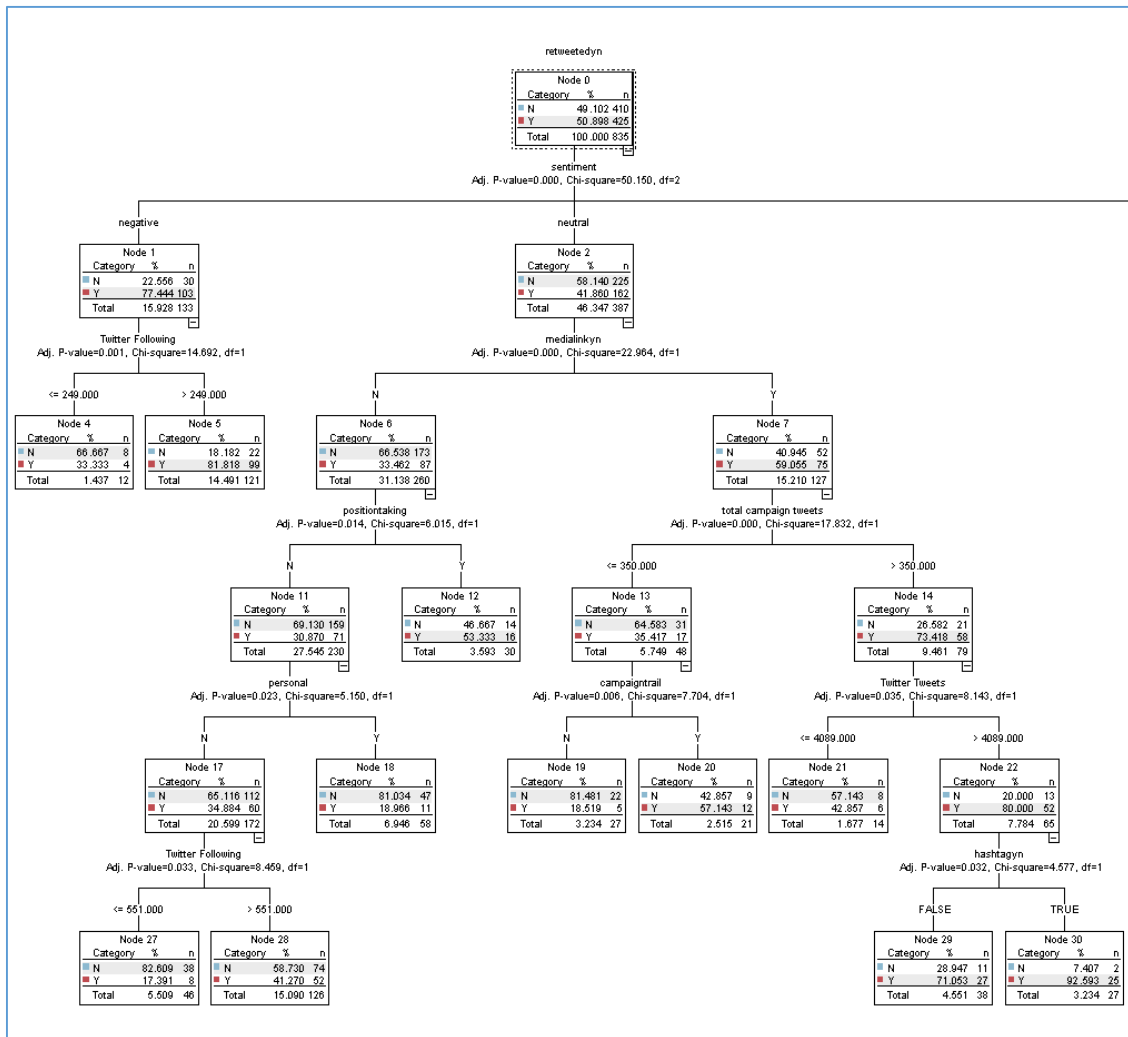
Variables included in this model

Sentiment – manually coded sentiment score

Twitter following – number of people the MP follows

Medialinkyn – contains a link to an image or video yes/no?

Positiontaking – expressing a position on an issue yes/no?

Total campaign tweets – total tweets sent during the campaign

Personal – containing personal information yes/no?

Campaigntrail – including a mention of the campaign trail yes/no?

Twitter tweets – total number of tweets during lifetime on Twitter

Hashtagyn – does the tweet contain at least one hashtag yes/no?

Following ratio – ratio of followers to followees

Achievement – contains mention of an achievement yes/no?

Total campaign tweets – total tweets sent during the campaign

Mentionnumber – number of mentions in the tweet

Supportforothers – mention of support for someone else yes/no?

Figure 91 - Model nine: blended model part two
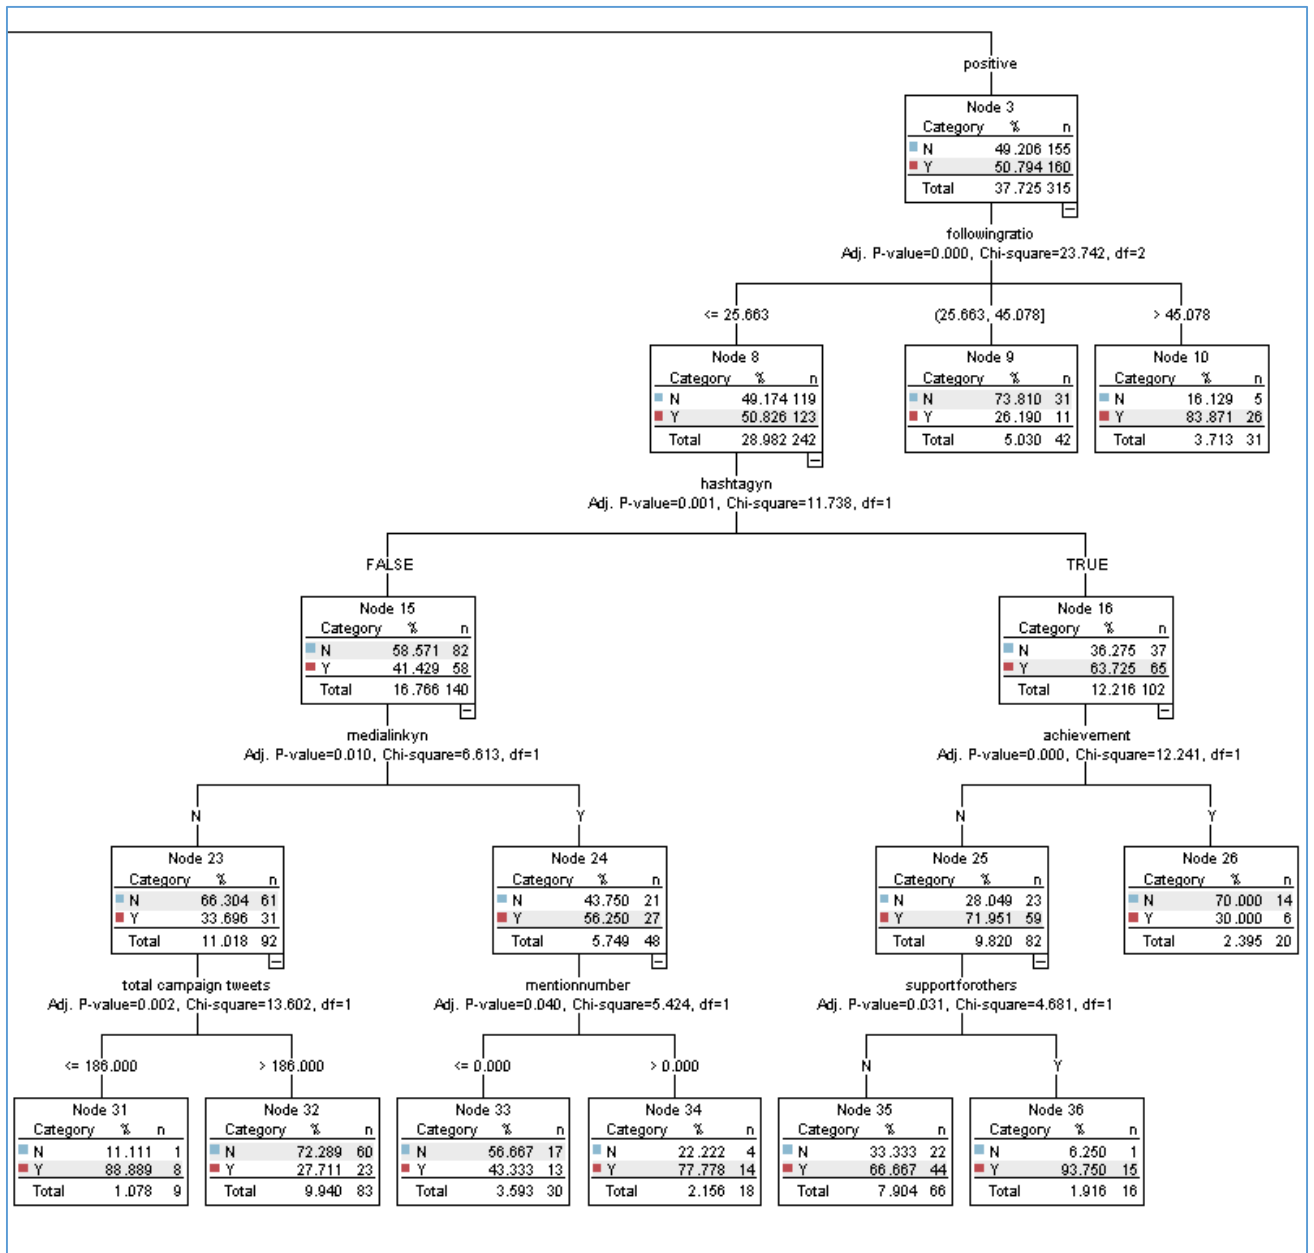
Model nine was accurate in its predictions 72.46% of the time on the training data and 65.28% of the time on the testing data (Table 45). In this case, the area under the curve was 0.784 for the training data and 0.681 for the testing data (shown visually in Figure 92). The full set of decision rules that lie behind this model can be seen in appendix six.

Table 45 - Comparison of performance on training and testing data for model nine

	Training		Testing	
Correct	605	72.46%	235	65.28%
Wrong	230	27.54%	125	34.72%
Total	835		360	

Table 46 - Evaluation of confidence scores for model nine

	Training data	Testing data
Range	0.531-0.897	0.531-0.897
Mean correct	0.73	0.728
Mean incorrect	0.673	0.698
Always correct above	0.897 (0% of cases)	0.897 (0% of cases)
Always incorrect below	0.531 (0% of cases)	0.531 (0% of cases)
93.02% accuracy above	0.818	Never reached

Figure 92 - Gains chart for model nine
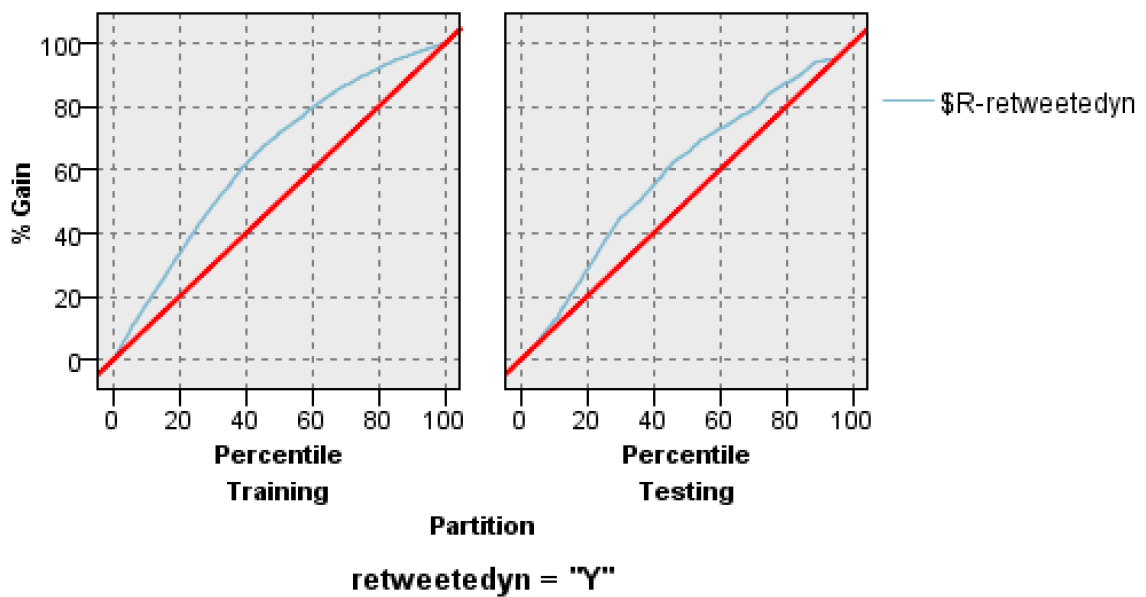
Once again sentiment emerged as the most significant predictor (Figure 93), followed by whether the tweet contains a link, how many people the sender follows, whether the tweet takes a position, the sender's ratio of followers to followees, including a hashtag, whether the tweet mentions the campaign trail, something personal or an achievement, and finally the total number of campaign tweets sent. From this, it seems that factors relating to the tweet's content have more influence than do factors relating to the author of the tweet, and that sentiment is the single most important factor by a reasonable margin.

Figure 93 - Most predictive variables in blended model nine
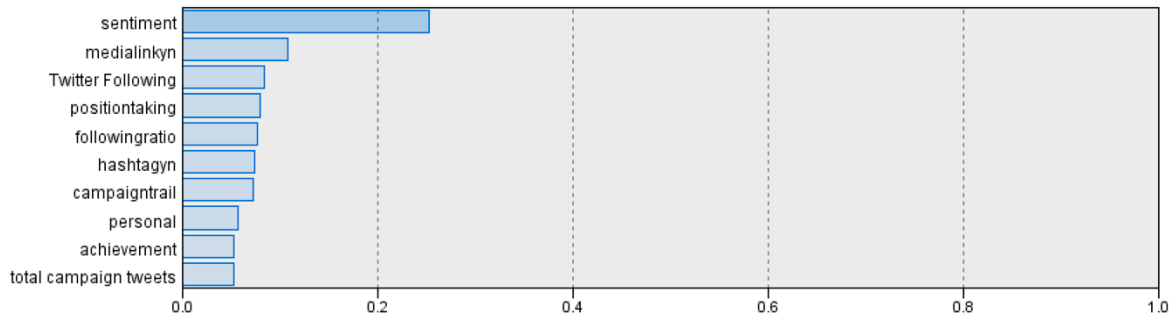
Negative tweets once again show up as the most likely tweets to be retweeted. Only one further split is added to the negative tweets, between those people who are following fewer than 249 people and those who are following more. In the first group, their negative retweets are only retweeted 33% of the time compared to 88% of the time in the second group. In this instance, the difference is likely to be because the number of people the MP is following is associated with some other variable that is not included in the model – most likely the number of followers. There is no obvious reason why following a small number of people on its own should be associated with a lower chance of getting retweeted, but following a small number of people is often associated with having a small number of followers or being relatively inactive on Twitter and sending a small number of tweets.

In the category of positive tweets, the first split is according to the sender's ratio of followers to followees, and that the results indicate that the smaller this ratio (i.e. the closer the numbers of followers and followees are), the lower the chance of positive tweets being retweeted. Those MPs with a follower/followee ratio of greater than 45 got 84% of their positive tweets retweeted (node 10) whereas those with a ratio of between 26 and 45 got only 26% of them retweeted. Those with a ratio of lower than 26 got almost exactly 50% retweeted. This result is probably indicative of the fact that the best-known MPs with the largest numbers of retweets are likely to have a high ratio of followers to followees, whereas those MPs with a smaller ratio are likely to be less well known and hence make less impact on Twitter. The results seem to suggest that if you do not have a very high ratio of followers to followees then the next best thing to do is try and keep the ratio as low as possible. There seems to be a slump in the middle where MPs have neither small enough ratios to get retweets due to being interactive, nor high enough ratios to get retweets from being well-known.

The best way to get a neutral tweet retweeted is to include a media link, be active on Twitter in terms of sending lots of tweets during the campaign and in one's total time on Twitter, and include at least one hashtag. These tweets get retweeted 93% of the time (node 30).

6.6.6. Model ten: CHAID model using all possible variables

The blended model (Figure 89) only uses the variables that the earlier CHAID models determined to be predictive and, in common with the earlier models, did not include author-related variables over which the tweet's author would have no control (number of followers, Twitter account status, gender, age and so on). Out of interest, one final model was built that included all possible variables, regardless of whether previous models determined them to be predictive or whether they related to things that the tweets' authors could influence or not (model ten – shown in full in Figure 94 and then split across Figure 95 and Figure 96)[35].

---

[35] A guide to interpreting CHAID output is provided in Figure 49 on page 190

Figure 94 - Model ten: all possible variables
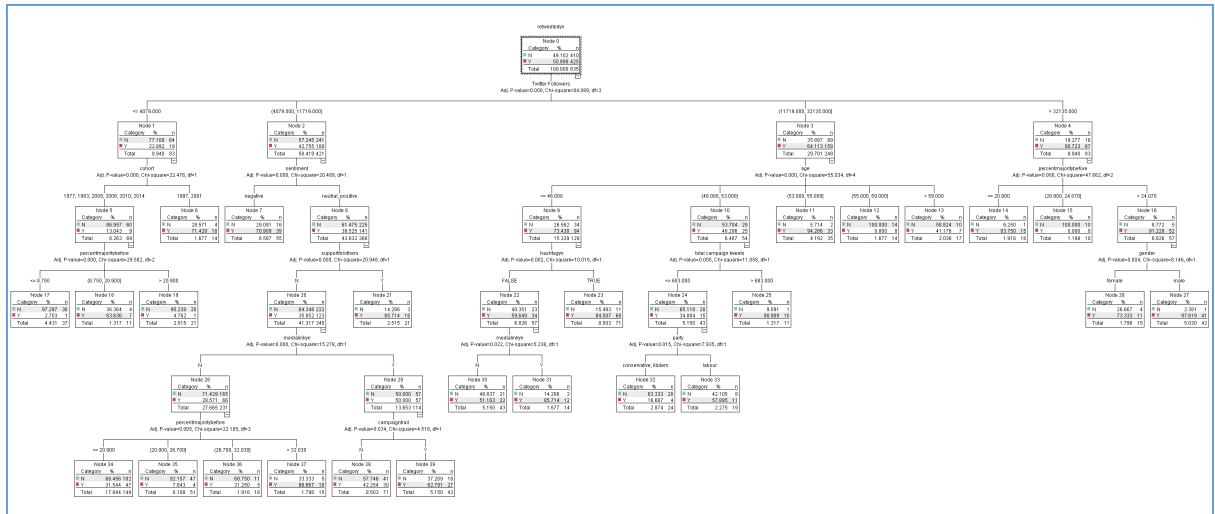
Figure 95 - Model ten: all possible variables part one

Figure 96 - Model ten: all possible variables part two
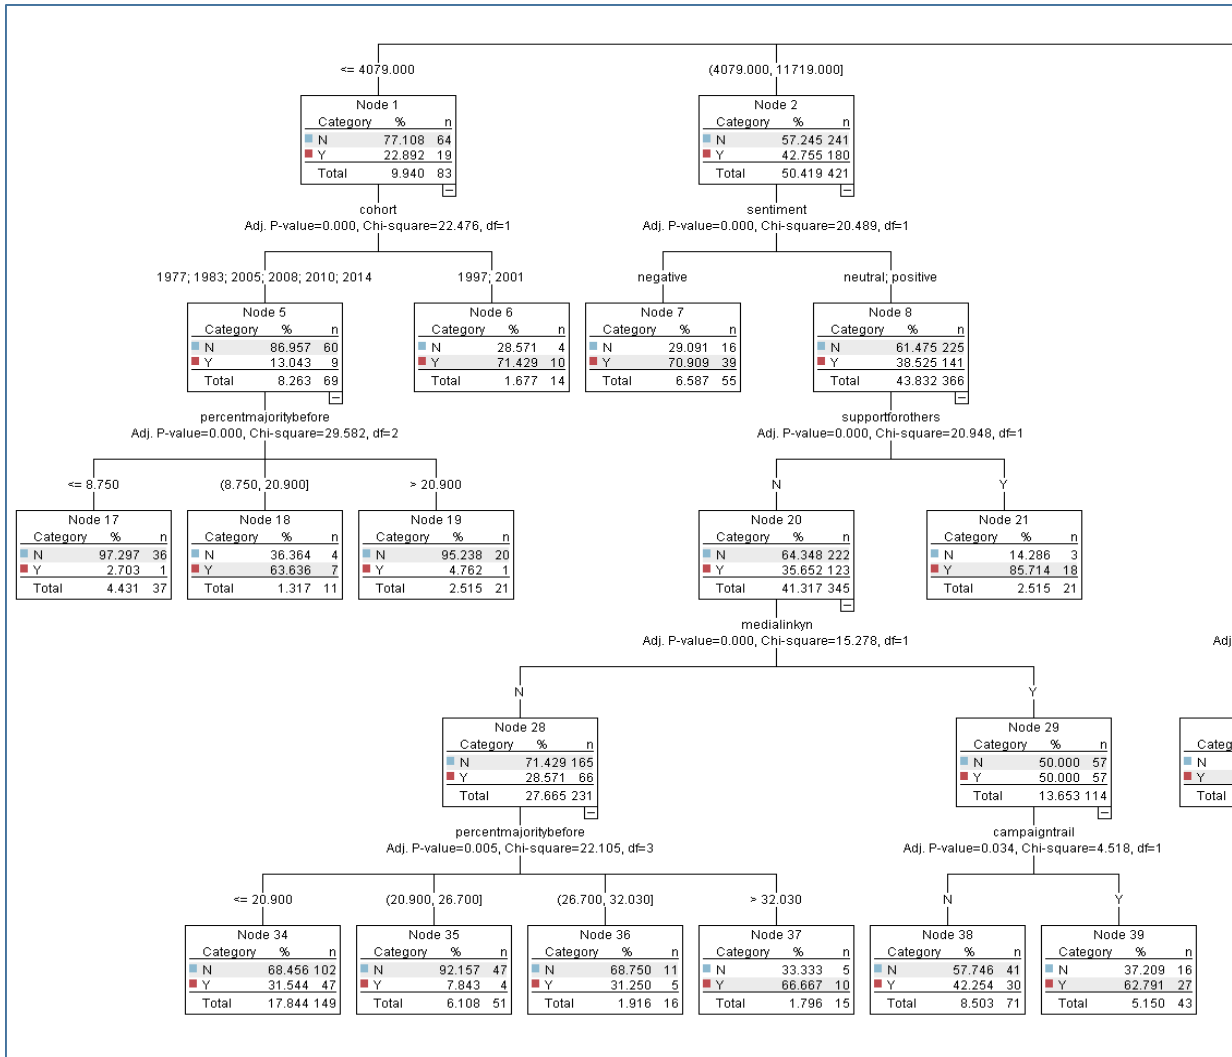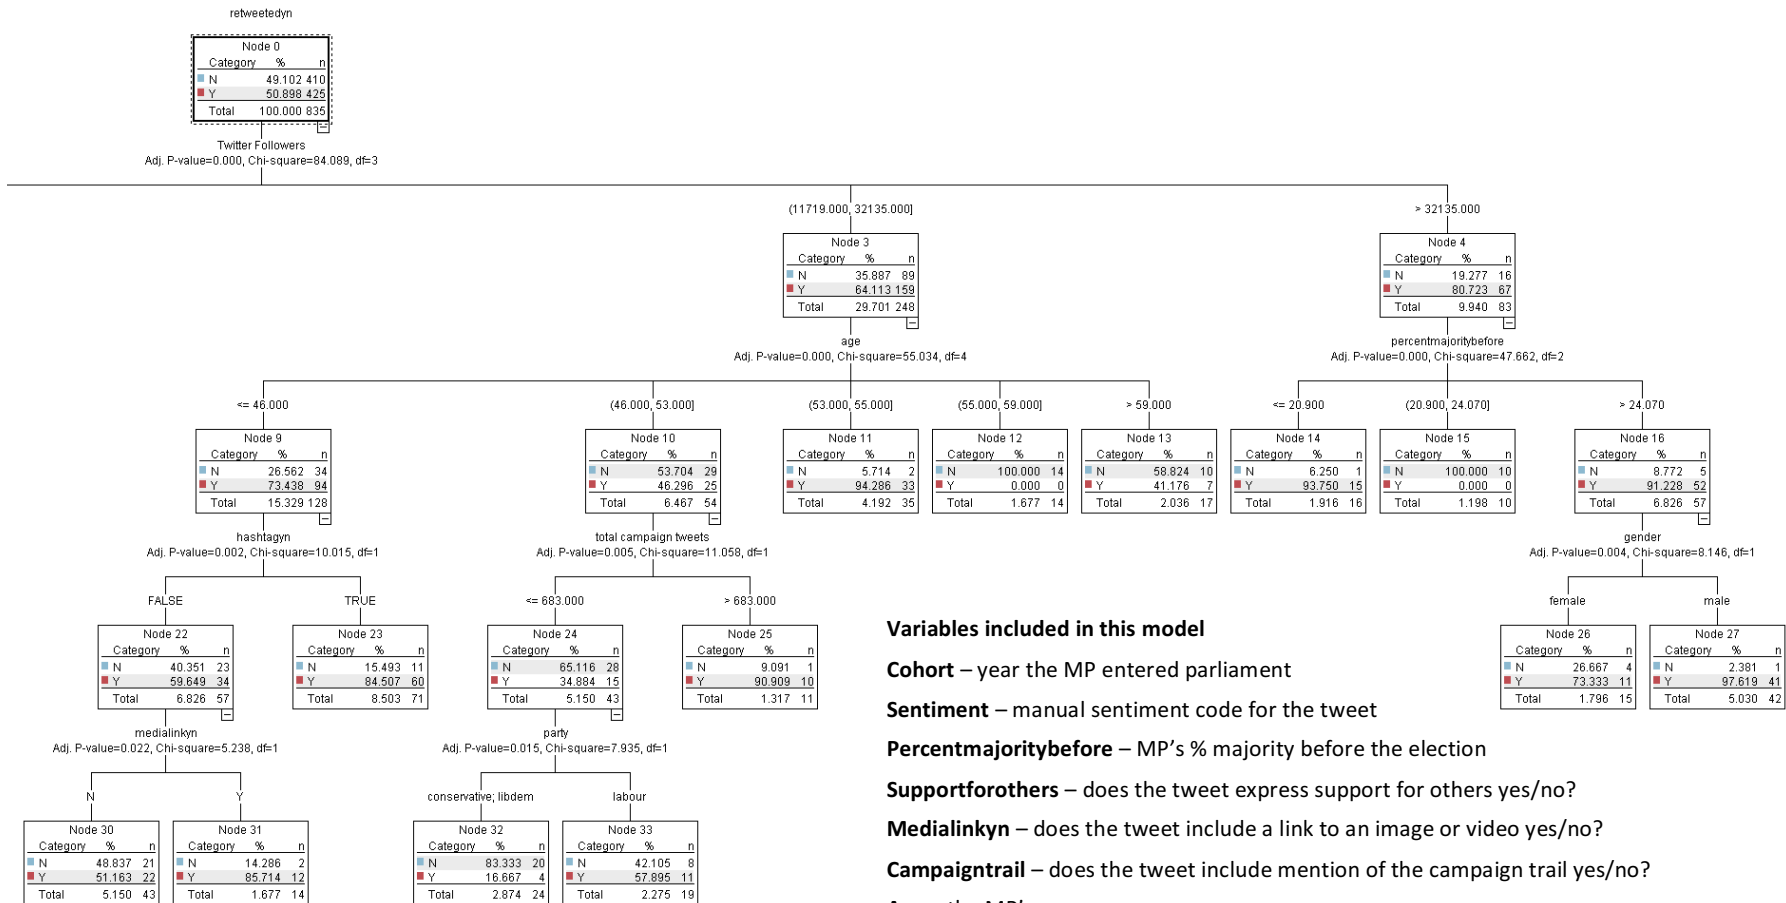
Variables included in this model

Cohort – year the MP entered parliament

Sentiment – manual sentiment code for the tweet

Percentmajoritybefore – MP's % majority before the election

Supportforothers – does the tweet express support for others yes/no?

Medialinkyn – does the tweet include a link to an image or video yes/no?

Campaigntrail – does the tweet include mention of the campaign trail yes/no?

Age – the MP's age

Hashtagyn – does the tweet include at least one hashtag yes/no?

Total campaign tweets – total number of tweets the MP sent during the campaign

Gender – MP's gender

This model once again showed the number of followers to be the most important predictor of retweets, followed by sentiment (Figure 97).

Figure 97 - Most important predictors across all variables
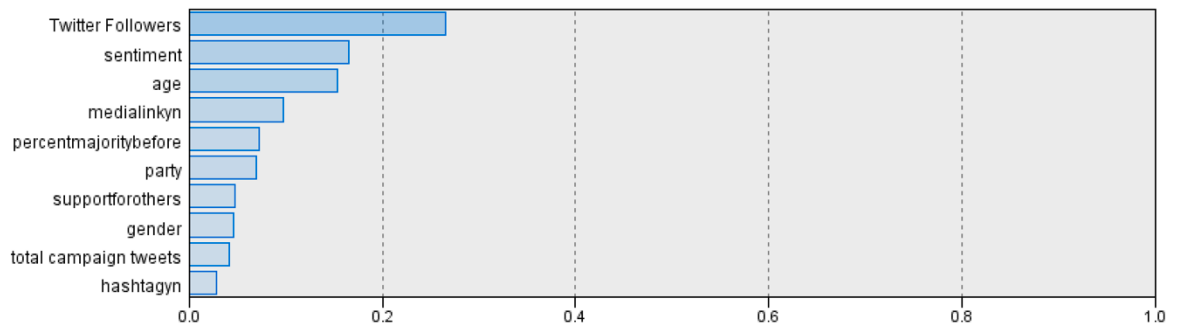
This model was correct in its predictions 76.29% on the training data and 71.67% on the testing data (Table 47). This shows that many of the predictors of retweeting are indeed things that may be beyond the direct control of the tweet's author, but some aspects of the tweet's content remain important – in particular, the tweet's sentiment which is still the second most predictive variable. A full set of decision rules for model ten can be seen in appendix six.

Table 47 - Comparison of performance on training and testing data for model ten

	Training		Testing	
Correct	637	76.29%	258	71.67%
Wrong	198	23.71%	102	28.33%
Total	835		360	

Table 48 - Evaluation of confidence scores for model ten

	Training data	Testing data
Range	0.511-0.955	0.511-0.955
Mean correct	0.774	0.768
Mean incorrect	0.662	0.696
Always correct above	0.955 (0% of cases)	0.949 (4.44% of cases)
Always incorrect below	0.511 (0% of cases)	0.511 (0% of cases)
90.84% accuracy above	0.702	0.919

In this case, the area under the curve was 0.845 for the training data and 0.783 for the testing data, visually represented in Figure 98.

Figure 98 - Gains chart for model ten
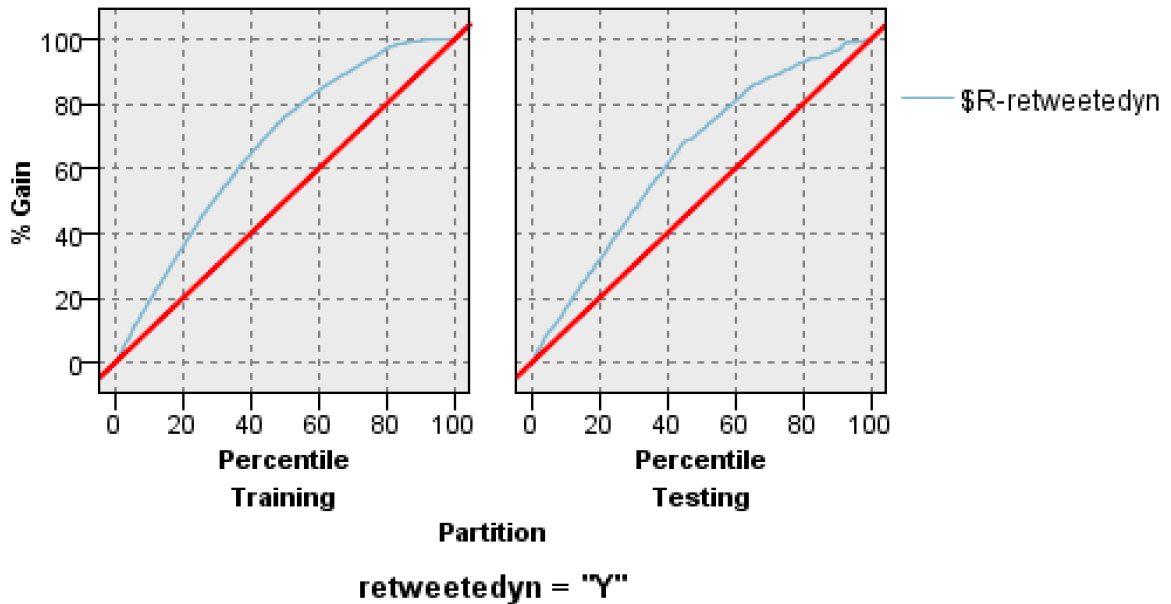
6.6.7. Comparison between models

Table 49 summarises the performance of the ten testing models that were created.

Table 49 - Summary of model performance

	Variables	Accuracy	Area under curve
Model one	Tweet structure	59.29%	0.634
Model two	Author characteristics	69.36%	0.775
Model three	Author and tweet structure	69.92%	0.783
Model four	Hashtags and machine sentiment score	57.53%	0.594
Model five	Machine generated content concepts	55.5%	0.583
Model six	Machine generated content categories	55.36%	0.557
Model seven	Manual content categories	65%	0.679
Model eight	Manual content and sentiment categories	63.31%	0.695
Model nine	Blended controllable variables	65.28%	0.681
Model ten	All possible variables	71.67%	0.783

The most predictive model in this research is correct 71.67% of the time. This is considerably better than random chance and provides valuable insight into the factors that influence retweeting in the case of politicians in the 2015 General Election campaign. The aims of this research are to enable better understanding of the factors that drive retweeting and to provide practical benefit to people who want to better understand what these factors are, rather than to build the most predictive model possible. It is not possible to model all the factors that influence retweeting as one does not have access to variables relating to the recipients of the tweets, and their characteristics will account for a significant proportion of the variance in retweeting rates between tweets. Also, it is not possible to allow effectively for the influence of context in models like these and it is clear when looking at the most highly retweeted tweets that context plays a hugely important role in determining their volume of retweets.

6.6.8. Chapter conclusion

There is no particular benchmark of predictiveness beyond which a CHAID model can be considered a success. The success or otherwise of the model is entirely determined by its usefulness to the model builder. In some cases, a model that is only a few percentage points better than random chance can be considered highly successful – if you're running large volume direct mail campaigns sending out millions of mailshots, for example, an uplift in response of only a fraction of a percent can represent many thousands of pounds of additional revenue. In another set of circumstances, such a model might be considered completely useless. Thus we cannot generalise about what constitutes a good model. The success of the model can only be evaluated on its own terms.

For models like those presented in this chapter, predictive power of between 60% and 70% can be considered successful. This is because we know from the literature that receiver characteristics play a significant role in determining whether tweets get retweeted, and those are not considered in the models presented here, so this research only considers two out of the three main influences. Additionally, no allowance can be made in these models for context, and looking at the most highly retweeted tweets shows that context is clearly a factor. Thus, we can assume that any variance in retweeting rates not accounted for by the models presented here is most likely accounted for by the characteristics of the tweet's receivers or by the tweet's context.

In the chapter that follows, the key findings of this research will be discussed in more depth and placed within the context of extant literature in the field as well as related back to the original objectives of the research in order to show how these have been met.

Chapter 7 Discussion

7.1. Chapter introduction

The purpose of this chapter is to draw out the key findings from this research and consider how they fit with what we already know about how politicians behave on Twitter and what drives retweeting. In particular, this chapter situates the findings of this research in the context of two key debates: the debate between techno-optimists and techno-pessimists, and the debate regarding the effectiveness or otherwise of negative campaigning. This research provides strong evidence that negative tweets are more engaging than positive ones, as well as providing qualified support for the techno-pessimist perspective which considers that social media is a new medium through which those who already have power can communicate, rather than offering a communication opportunity to new groups. The chapter is structured around the research objectives set in chapter one. Each will be discussed in turn.

7.2. What factors determine whether tweets are retweeted?

The first research objectives set in chapter one were to identify the factors that extant literature suggests are most likely to determine whether tweets get retweeted (objective one) along with any other factors which might also play a role in influencing whether politicians' tweets get retweeted (objective two). A review of literature relating to politicians on social media, electronic word-of-mouth and online virality (presented in chapters two and three) identified three broad categories of factor shown to influence whether tweets are retweeted.

1. The characteristics of the sender of the tweet
2. The characteristics of the tweet itself
3. The characteristics of the recipient of the tweet

These factors were summarised in the conceptual model shown in Figure 99. This chapter discusses the role that each element of the conceptual model plays in

determining whether MPs' tweets get retweeted as determined firstly by the descriptive analysis and secondly by the CHAID modelling process.

Figure 99 - Conceptual model of factors influencing retweeting
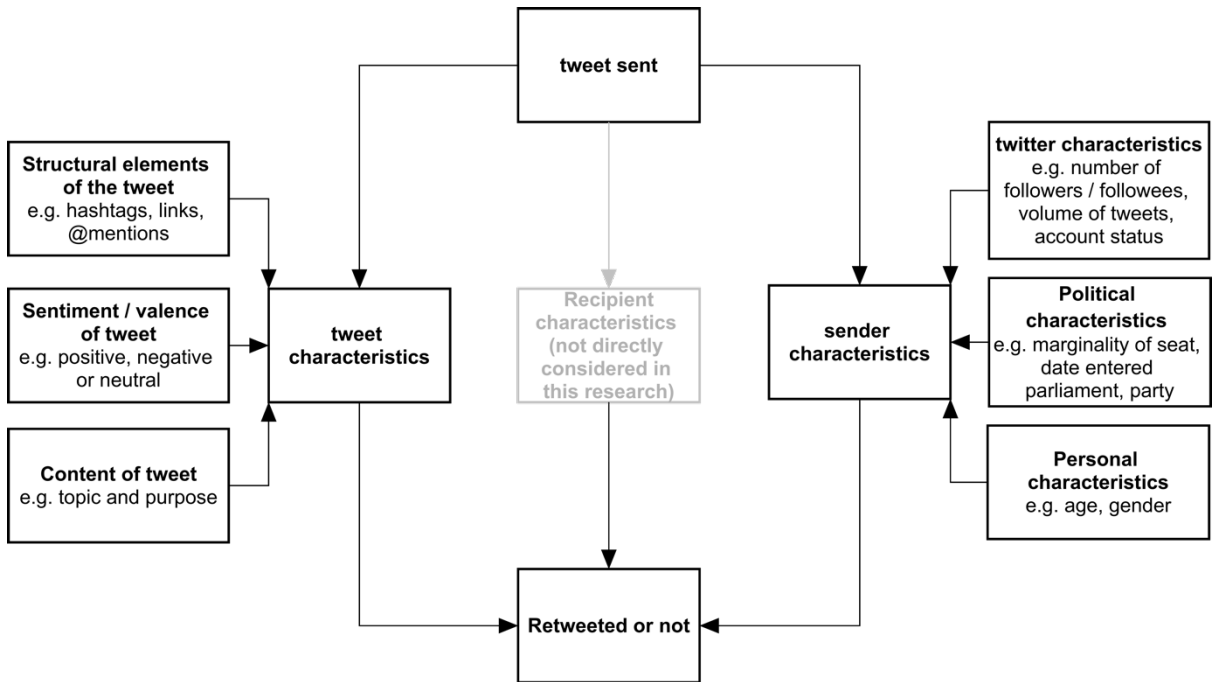
The focus of the model building, as discussed in the previous chapter, has been on including only variables over which politicians have some measure of control. Therefore, these models are less predictive than they would be if they included a broader range of factors that influence retweets, but the broader aim is to produce practical guidelines for tweeting rather than to build the most predictive model possible.

Additionally, information about tweets' recipients is not readily available through the Twitter API and so research that considered the role they play would need to be conducted using different methods from those used here. It is for this reason that almost all extant quantitative research considering retweets limits its focus to author characteristics and tweet characteristics.

That said, this research does move considerably beyond most extant research in the field by including the consideration of variables derived from manual sentiment and content analysis. As discussed in chapter four, extant research predicting retweets focuses almost exclusively on those variables which can be extracted from the Twitter API or calculated using machine-based analytics. There is some research which includes manual content analysis of politicians' tweets but the focus of this is purely on describing the characteristics of their tweets rather than on determining how effective each kind of tweet is at stimulating retweets.

7.3.   A typology of politicians' tweets

Objective three was to propose a typology of tweets sent by UK politicians during the 2015 General Election campaign and identify which work best in terms of stimulating retweets. This was a two-stage process. The first stage involved descriptive analysis of the entire population of tweets sent by MPs in the run up to the election in order to better understand who was tweeting and to identify any interesting differences in tweeting behaviour between different groups of MPs. The second stage was a review of the literature to better understand extant political tweet typologies, before building a new typology (shown in appendix five) which builds on what went before wherever possible but is also updated to take into account how politicians' Twitter behaviour has changed. This typology formed the basis of a manual content analysis of a sample of the tweets. Key findings from these two elements of the research will be discussed in turn, with particular consideration of how they relate back to extant literature.

7.4.   Changes in patterns of MPs' tweeting behaviour

MPs' tweeting behaviour has changed since the 2010 General Election campaign. In 2010 less than 20% of candidates' tweets were retweets compared to almost 50% that were original posts (Graham *et al*., 2013). By 2015, these proportions had almost reversed, with 51% of the MPs' tweets being retweets compared to 27% original tweets. It is likely that this reflects changing use of Twitter in general rather than a specific change that affects only the MPs. This is also reflected in a huge jump in

tweeting volume from the MPs compared to previous research. In the 2010 election Labour averaged 62 tweets per candidate. By the 2015 election this rose to 302, and the mean number of campaign tweets per MP was 422. This shows that there is recognition amongst MPs that Twitter is a useful communication tool. There is also presumed recognition amongst voters that following politicians on Twitter is a worthwhile activity, as shown by the fact that the median number of followers per MP now (7,919) is almost double the highest number of followers that any single MP had in 2009 (4,441 for Tom Watson).

An ongoing debate in the literature is whether politicians use Twitter as primarily a broadcast tool or as tool for two-way communication. Looking at the patterns of MPs' tweets in this data shows that 21% of their tweets are replies to other people, a drop from 31.8% in 2010 (Graham *et al.*, 2013), so on the surface it appears that perhaps the MPs are less interactive than previously. However, number of replies is not the only measure of interactivity. One could also take the number of times MPs retweet other people's tweets as being a measure of interactivity – retweeting something indicates a measure of engagement with the author of the original tweet. Thus, one could argue that the fact that over 50% of the MPs' tweets are retweets shows they are interacting with other people on Twitter and not simply broadcasting their own messages. This is a very significant rise from 2010 when less than 20% of the politicians' tweets were retweets (Graham *et al.*, 2013). Clearly, some MPs made strong attempts to interact with other people on Twitter. For example, 72% of Naomi Long's tweets (the sole representative of the Alliance Party) were replies to other people, compared to 21% across the whole dataset.

The ratio of followers to followees has also been suggested as a measure of the extent to which politicians are interactive on Twitter, with a figure of 10 or below being suggested as an indication that the politician is at least trying to listen to other people rather than to simply broadcast (Jackson and Lilleker, 2011). Of those MPs tweeting prior to the 2015 General Election, almost half (48%) had a follower / followee ratio of 10 or less, suggesting some measure of listening behaviour is reasonably widely adopted. This number has fallen somewhat from 60% in 2009 (Jackson and Lilleker,

2011) but the changes to Twitter in the intervening years render a comparison over this timespan meaningless. Many MPs now have tens or even hundreds of thousands of followers, by which point following all of them back becomes logistically impossible.

7.5.   Analysis of retweet patterns

Estimates suggest that somewhere between 6% (*Replies and retweets on Twitter*, 2010) and 36% (Enge, 2014) of all tweets get retweeted whereas nearly 85% of the MPs' tweets got retweeted. As with all variables relating to patterns of Twitter use, the distribution of number of retweets is massively skewed, with most tweets getting one or two retweets whilst a few generate huge numbers, up to a maximum of 13,919, but with a median of only four. There is also a big difference between the percentage of each MP's tweets that get retweeted – a comparative measure of Twitter performance which does not take into account retweet numbers. Out of 366 MPs, 46 got all of their tweets retweeted at least once whilst 23 managed to get less than 50% retweeted and seven managed none. There are also substantial differences between the success rates of the different parties.

This research does not attempt to predict the number of retweets that a tweet will get, but the analysis conducted suggests that the factors that influence the number of retweets are not the same as the factors that influence whether a tweet gets retweeted or not. Literature relating to virality might be of more use when it comes to explaining how many times tweets get retweeted but there are very few tweets indeed in this dataset that could meaningfully be said to have gone viral. The handful of retweets that the typical tweet gets certainly does not count as 'going viral'. Ed Miliband and David Cameron between them are responsible for 86% of the tweets that got more than 1,000 retweets. Thus, it seems that who you are is probably the most significant determinant of large retweet volumes. For most MPs then, getting thousands of retweets is not an achievable objective and a more realistic goal is to focus on getting more of their tweets retweeted.

Generally, the figures relating to retweet volumes support the techno-pessimist argument that social media in politics represents a new medium through which political business as usual can be conducted. Despite there being some MPs who have used social media as part of an effective strategy to build their public profile (for example Stella Creasy), it is still the case that those who have the loudest voices on Twitter, as measured by the number of followers they have and the volume of retweets that they attract, are all mainstream names, the party leaders plus one or two others.

The literature review and subsequent conceptual model suggest that retweeting is influenced primarily by three aspects of the sender's characteristics that are relevant to retweeting: their personal characteristics, their Twitter characteristics and their political characteristics; and by three aspects of the tweet's characteristics: its structural elements, its sentiment and its content. The findings relating to each one of these elements of the conceptual model will be discussed in turn in the section that follows.

7.5.1. Effect of MPs' personal characteristics on retweets

7.5.1.1. Demographics

Going into the 2015 election 23% of MPs were women and they represent 27% of the tweeting MPs so there is a small indication that women are more likely to be on Twitter in the first place. This gender gap has narrowed since 2009 at which point women accounted for 19% of MPs in parliament but 29% of those on Twitter (Jackson and Lilleker, 2011). However, demographic factors such as age and gender may affect whether one is on Twitter in the first place but do not appear to have a meaningful effect on how many tweets one sends, whether one's tweets get retweeted or how many times one's tweets get retweeted. Where small relationships have been detected (for example between gender and retweeting) this is almost certainly because the highest profile politicians that achieve the largest numbers of retweets are men rather than because people are less likely to retweet women's tweets because women sent them.

7.5.2. Effect of MPs' Twitter characteristics on retweets

7.5.2.1. Follower numbers

For the most part, the data provides further support for some things we already know about the way that Twitter behaviour influences retweets. For example, the percentage of an MP's tweets that are retweeted is significantly correlated with the number of followers that they have and their ratio of followers to followees. Similarly, the number of retweets that MPs achieved per campaign tweet they sent is heavily influenced by the number of followers they have, the ratio of followers to followees and also the absolute number of people they follow. However, there is no significant relationship between the total number of tweets each MP sent either during the campaign or during their entire time on Twitter and the percentage of their tweets that are retweeted or the number of retweets they generated per tweet. This shows that the key driver of retweet volume is the number of followers that one has (in support of extant literature e.g. Suh *et al.*, 2010; Yuan Huang and Zhang, 2015; Zhang, Xu and Yang, 2012), followed by the number of people one follows and the ratio of one to the other. Broadly, as one's follower numbers increase so one gets more retweets (Figure 100). Getting retweets is not just a matter of tweeting as much as you can until something sticks – building a network of followers is critical.

Figure 100 - Relationship between Twitter followers and campaign retweets generated
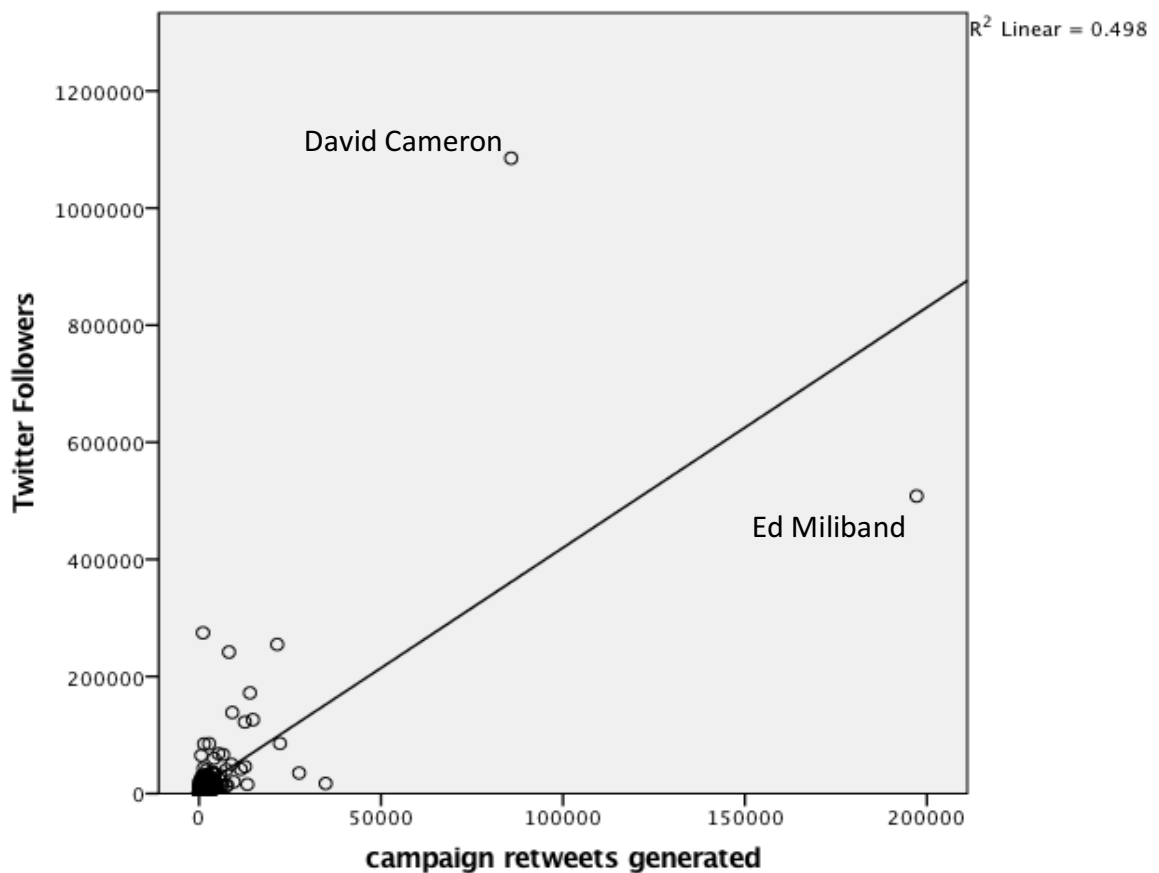
Calculating the number of retweets per follower for each MP provides further evidence for the closeness of the relationship between followers and retweets. This ranges from zero to two, with a median of 0.06 and there is little variation in the data, showing that for the most part the number of retweets per follower does not vary much, meaning that getting more followers should naturally lead to getting more retweets (Figure 101). The outlier in this case is Pete Wishart[36], the sole MP who manages to achieve two retweets for each one of his followers.

---

[36] Pete Wishart is an SNP politician who, before entering politics, had a successful career as a rock star with the Scottish band Runrig. He was a well-known figure in Scotland before becoming an MP and is not a typical case.

Figure 101 - Plot of retweets per follower for each MP
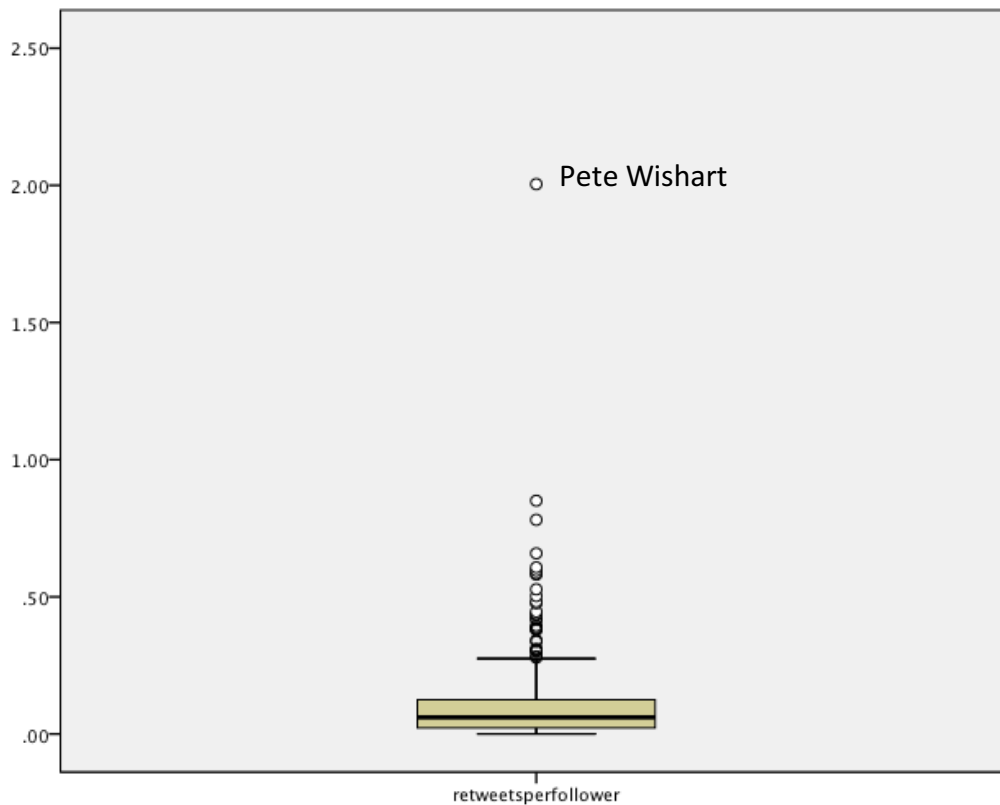
7.5.2.2. Account status

MPs with verified accounts get more retweets than those with non-verified accounts, but this is almost certainly because those with verified accounts are the best-known, highest profile politicians. This provides further support for the techno-pessimist argument that Twitter activity mirrors real life rather than providing an opportunity for lesser-known politicians to build their profiles.

7.5.3. Effect of MPs' political status on retweets

7.5.3.1. Party

The results suggest that the different parties clearly have different strategies when it comes to Twitter, and a divide between the larger parties (Labour, Conservatives and Liberal Democrats) and the smaller ones is clear. It is worth noting that all the smaller parties represented in parliament had at least one tweeting MP. This is a change since

2009 when there were only two tweeting MPs from outside the three main parties – one from the SDLP and one from the SNP (Jackson and Lilleker, 2011). As in 2009 (Jackson and Lilleker, 2011) and 2010 (Graham *et al.*, 2013), Labour continues to be over-represented on Twitter whilst the Conservatives are under-represented relative to their numbers in parliament, although the gap between the two parties has narrowed somewhat.

Although the larger parties achieve a high volume of both tweets and retweets (due to there being more of them), it is the smaller parties that are substantially more active when it comes to the number of tweets sent per person during the campaign. In the case of the Alliance, Respect, Plaid Cymru and the Green Party this is down to the tweeting of a single MP. Perhaps this could give some hope to the techno-optimists, as it demonstrates that Twitter is a medium where the parties are not limited by their size and personal campaigning is possible. Clearly, the three largest parties have campaign budgets way in excess of anything that the smaller parties can muster but on Twitter that doesn't matter – the smaller parties can punch well above their weight thanks to the efforts of a few motivated individuals. Likewise, some of the small parties did extremely well generating retweets, with the Greens, Respect, SNP and Ukip all getting more than 90% of their tweets retweeted compared to less than 50% for the Alliance and DUP. The larger parties performed well here too with retweet rates of 80% for the Liberal Democrats, 82% for the Conservatives and 87% for Labour, further evidence that, of the three main parties, Labour is perhaps making the most effective use of social media.

7.5.3.2.   Parliamentary cohort

Whilst the cohort of politicians who entered parliament in 2010 are over-represented in this dataset, they are no more active or successful on Twitter than those elected earlier, suggesting that whilst they may be more accepting of social media as a group, perhaps due to being younger or being elected for the first time firmly in the Twitter age, they are not able to translate that into more success on the medium (as measured by retweets) than politicians elected earlier. Being younger or more 'social media savvy' does not appear to confer any particular benefit in terms of being better

able to get one's tweets retweeted, further supporting the 'business as usual' position of the techno-pessimists.

7.5.3.3. *The relationship between tweets, retweets and election results*

Evidence regarding a link between Twitter behaviour and election result is limited. There is some evidence to suggest that perhaps MPs' Twitter behaviour is driven to a certain extent by their electoral position. For example, MPs who stood down at the election sent significantly less tweets than those who either won or lost. Those who went on to lose their seats sent more campaign tweets than those who held their seats, suggesting that perhaps they knew they were in a closer race and would need to work harder to campaign on their own behalves. Some additional support for this theory is provided by the fact that the MPs in safe seats sent significantly fewer tweets than those in marginal or near marginal seats. Whilst the MPs who held their seats achieved a marginally higher (but significant) number of retweets per tweet, this is more likely to be because all the highest achieving retweeted MPs held their seats (David Cameron, Ed Miliband and so on) rather than because there is a direct link between the number of retweets an MP gets and whether they then go on to hold their seat or not. There is no significant relationship between the number of retweets generated and the safeness of the MP's seat.

7.5.4. Effect of tweet characteristics on retweets

7.5.4.1. *Structural tweet elements*

Another aspect of MPs' tweeting behaviour that influences whether their tweets are retweeted is the structure of the tweets. In particular, whether structural elements such as hashtags, links and mentions are included. This research provides further support for extant research showing that including hashtags in a tweet significantly affects whether or not it is likely to be retweeted. It also supports broader research on Twitter beyond politics which shows that links, images and videos do better at getting retweets. However, it contradicts the findings of some research looking specifically at political tweets (Dang-Xuan *et al*., 2013) which found that including a link had no

effect on the chances of a tweet getting retweeted, and Liu, Liu and Li (2012) who found that including links was negatively associated with retweeting.

There is clear evidence that parties are developing more sophisticated strategies for using hashtags than was the case in 2009 when Golbeck *et al*. (2010) found that only 0.08% of Congresspeople's tweets contained hashtags. Usage now is more in line with Parmelee and Bichard's (2012) finding that 60% of politicians' tweets contained hashtags during the 2010 congressional campaigns. That said, only 5% of Labour tweets contained the hashtag #votelabour and 7% contained #labourdoorstep despite the fact that tweets containing these hashtags got retweeted 97% and 94% of the time respectively and were being promoted as the main campaign-related hashtags by the central party, so clearly hashtag use is not universal.

There are also some interesting differences in hashtag use between the parties, pointing to differences in social media strategies. Labour, the Conservatives and the Liberal Democrats were broadly equally likely to use hashtags (40%, 42% and 44% of their tweets contained hashtags respectively). Only one of the smaller parties made any more use of hashtags than this - 69% of the Green Party's tweets include at least one hashtag. This provides qualified support for Gainous and Wagner's (2014) finding that challenger parties were twice as likely to use hashtags in their tweets as were incumbent parties. Clearly, there is work to be done in helping MPs from other small parties better understand the value of hashtags in getting their tweets out to people.

Having said all of that, it is important to recognise that those tweets that were most frequently retweeted do not include hashtags, images or links. Looking at the ten most retweeted tweets, only one contains a hashtag, four contain links to pictures and one contains a mention of someone else. Overall five out of the ten do not include any links, hashtags or mentions. Such elements play an important role in determining whether most tweets are retweeted but the fact that they are much less likely to feature in the most highly retweeted tweets provides further evidence that the factors that influence retweet volume are not the same as those that influence whether or not a tweet is retweeted at least once. Clearly in the case of the most retweeted

tweets, the thing that is driving the retweet volume is the context in which they were sent, something that is almost impossible to quantitatively measure.

Additionally, when examining the influence of author-related variables compared with the tweet structural variables, those models using author characteristics performed better than those including only tweet structural information, showing that more of the variance in retweet rates is explained by the characteristics of the person sending the tweet than by the structural elements of the tweet.

7.5.5. MPs' tweets categorised by content and sentiment

A sample of the tweets sent by MPs in the 2015 General Election campaign was subject to both computer-based and manual sentiment and content analysis. The findings from analysis of these variables will now be discussed.

7.5.5.1. How sentiment influences retweeting

This research has revealed some interesting findings with regard to the impact of sentiment on retweeting. There are two aspects to these findings. One is methodological and relates to the effectiveness of different ways of categorising sentiment in tweets, the other is practical and relates to the effect that sentiment has on retweeting. Each will be discussed in turn.

As part of this research, three methods of sentiment analysis were compared. Two used machine-based sentiment coding (based on algorithms built into Brandwatch and SPSS Text Analytics) and one was manual coding. The first useful finding is the amount of variation between the two machine algorithms and the human coder. These comparisons show that SPSS was much more willing to flag tweets as positive or negative whereas the Brandwatch algorithm was much more conservative and tended to assume most tweets were neutral. Of course it makes a difference that Brandwatch has a neutral category, which SPSS does not, so perhaps this means that SPSS is more likely to 'force' tweets into either positive or negative categories as it does not have the option of neutral (although it does treat positive and negative as two different

variables, each one is a yes/no flag so it does have the option of coding a tweet as neither positive nor negative).

Levels of agreement between Brandwatch and the manual sentiment analysis were low. Both agreed on how a tweet should be classified in only 54% of cases which is little better than random chance. In the case of negative tweets, 59% of Brandwatch's negative tweets were also deemed negative by the human coder, 51% of Brandwatch's neutral tweets were agreed to be neutral by the human coder and 74% of Brandwatch's positive tweets were also coded as positive by the human coder. However, as with the SPSS algorithm, the manual coder was much more likely to allocate a positive or negative code than was Brandwatch. In the sample of 1,212 tweets that were manually coded, Brandwatch identified only 32 as negative compared to 203 identified as negative by the human coder. Brandwatch identified 992 as neutral compared to 553 by the human coder, and 171 as positive compared to 439 by the human coder.

A similar comparison between SPSS and the human coder reveals that only 42% of the tweets that SPSS thought were negative were also coded as negative by the human coder, compared to agreement over 52% of its positive tweets. SPSS coded 253 tweets as negative compared to 203 for the human coder. It categorised 627 as positive compared to 439 for the human coder. Generally, it seems that Brandwatch was more conservative in its assessments of sentiment and hence erred more on the side of suggesting tweets were neutral whereas the human coder, with a better grasp of nuance and context, was able to more accurately assess whether tweets were positive or negative. SPSS, on the other hand, without the benefit of a neutral category, tended to overestimate the numbers of tweets in both the positive and negative categories.

Each of the three sets of sentiment codes were tested to see if they influenced the chances of a tweet being retweeted at all. The results indicate that the manually coded sentiment categories are much more robust than either those of Brandwatch or SPSS. Just short of 80% of the manually coded negative tweets were retweeted

(compared to 58% of Brandwatch's negative tweets and 55% of SPSS's) which shows that the manually coded category was picking up more strongly on a real theme in the data – negative content is more likely to get retweeted.

Earlier research (Klotz, 1998; Bimber and Davis, 2003) showed that politicians were very unlikely to campaign negatively online, largely because it is much harder to distance yourself from a negative attack on your opponent if it is made on your own website. Looking specifically at Twitter, Jackson and Lilleker (2011) found that MPs were very unlikely to use it for attacking their opponents. However, other research suggests that, as online content is generally preaching to the converted, politicians are more likely to be negative online as they do not risk alienating undecided voters (Druckman, Kifer and Parkin, 2010). The research presented here shows that some politicians (not all) now are campaigning negatively on their own behalves and that this can be a successful strategy in terms of generating Twitter retweets.

The second useful finding is that not only are negative tweets significantly more likely to get retweeted in the first place, but they are also more likely to get retweeted a large number of times, with negative tweets attracting a median of more than twice as many retweets as either positive or neutral tweets. These findings show not only that MPs are more likely to go negative online than has previously been thought, but also that citizens appear to respond more readily to negative content coming from MPs via social media. This is a new finding – there does not appear to be any previous research that looks at this particular aspect of political negativity online. This finding could also be taken as further support for the techno-pessimists' position as it suggests that in 2015 the online campaign reflected what was generally seen as a very negative and fear-based offline campaign (Shrimsley, 2015) rather than offering anything different.

Evidence is mixed regarding the influence of sentiment on people's propensity to pass along information online. Berger and Milkman (2012) found that positive content was more likely to go viral than negative. Although they were not looking specifically at tweets, the research presented here contradicts their finding. Looking at the 28

tweets in this sample that got more than 100 retweets (the only tweets that could be said to have gone viral in any meaningful sense) shows that 64% of them were negative, with the others evenly split between positive and neutral. Broadening this out to all tweets with 16 or more retweets (defined by SPSS as extreme values in this dataset) a similar pattern holds, with 50% of these tweets being negative compared to 30% positive and 20% neutral. This compares to 17% of the tweets being negative in the dataset as a whole. Thus, it can be concluded that negative tweets from MPs are considerably more likely to get large numbers of retweets than are positive or neutral tweets.

Further descriptive analysis determined that not all MPs were equally like to 'go negative' in their tweets. There is a statistically significant difference between Labour and the Conservatives as regards sentiment – almost 19% of Labour's tweets are negative compared to 16% of the Conservatives (chi-square = 10.980, p (two-tailed) = .004). The Liberal Democrats were much less likely to go negative – only 9% of their tweets were negative. Men were almost twice as likely to post negative tweets as women – 20% of the men's tweets were negative compared to 10% of the women's (chi-square = 19.065, p (two-tailed) < .001). There is also a significant relationship between marginality of seat and propensity to go negative. Those in marginal seats were significantly less likely to send attacking tweets (8% compared to 16% for near marginal MPs and 19% for MPs in safe seats – chi-square = 11.493, p (two-tailed) = .003) or to send negative tweets (8% compared to 17% for near marginal and 19% for safe seats – chi-square = 16.313 p (two-tailed) = .003).

Some individual MPs have a higher propensity to go negative than others. Selecting only those tweets that are coded as negative or as attacks or as fear appeals (237 in total), one can see that 101 MPs were responsible for these, meaning that 265 MPs did not send any negative tweets. Within that group of 101, most MPs were responsible for one or two of these at most – the mean number of negative or attacking tweets per MP was 2.34. However, Barry Sheerman sent 21 of them, Brandon Lewis sent 15 and Karl Turner sent 11 (no other MPs went into double digits) indicating that personal propensity to go negative varies significantly between MPs. Of

the variables under the sender's control, the sentiment of the tweet is by far the most powerful in predicting whether a tweet will be retweeted.

7.5.5.2.   How content influences retweeting

There was no existing content coding schema suited to the specific circumstances studied here and existing coding schemas based on other contexts did not fully capture the richness of MPs' tweeting so a new coding schema was developed. How MPs use Twitter now has changed substantially from how they used it in elections of 2010 and that is reflected in the wider range of codes needed. For example, tweeting about television programmes was not something that MPs did in 2010. It has emerged as a much more recent phenomenon, and one that MPs have embraced just as other Twitter users.

The most common type of MPs' tweet by a considerable margin is the local tweet – one which makes some explicit mention of the MP's constituency. Just under 40% of the coded tweets feature a constituency mention of some kind. This provides support for Jackson and Lilleker's (2011) suggestion that MPs use Twitter as a way of personally campaigning and boosting their impression of local service and is in line with their finding that tweets talking about MPs' local work were the most common. Also in common with Jackson and Lilleker (2011), MPs in safe seats were less likely to send local tweets than are those in marginal and near marginal seats however this difference is not statistically significant at the 0.05 level (chi-square = 2.912, df1, p (two-tailed) = .088). There are, however, statistically significant differences between the parties when it comes to using local tweets. This tactic is most often used by Conservative MPs, 44% of whose tweets include a local mention (compared to 35% for Labour and 30% for the Liberal Democrats, chi-square = 12.640, df2, p (two-tailed) = .002). However, these local tweets are no more or less likely to be retweeted than any other kind of tweet.

The second most common form of tweet is a mention of some aspect of the campaign trail. These tweets keep MPs' followers informed about where they are and what they

are doing. They give a clear impression of the MP as being active and engaged. These could be seen as another form of local tweet as they tend to involve the MP saying something about what kind of campaign activities they are doing in their local area on a particular day. All the major parties are equally likely to use this kind of tweet and there are no other obvious significant differences between particular groups of MPs and their use of campaign trail tweets. One could say that the campaign trail / local tweet combination is the default tweet of choice. This is in line with previous research examining American politicians and showing that just over 40% of their tweets were some kind of update from the campaign trail (Parmelee and Bichard, 2012). Campaign trail tweets are more likely to be retweeted than local tweets, being retweeted at a rate of 61% compared to 50% for the sample as a whole. This is probably because a local tweet could just literally be a mention of the weather in the constituency and does not necessarily have to include any mention of the election, whereas campaign tweets clearly have an election focus and hence are perhaps more likely to be retweeted by other members of the campaign team.

The third most common form of tweet is the attack tweet. This is a change from 2009 when Jackson and Lilleker (2011) found that MPs were very unlikely to use Twitter for attacking others. Attack tweets were the least common form of tweet in their dataset. In this dataset there were 208 attack tweets in the manually coded sample, 17.5% of the total, and 39 fear appeals (just over 3% of the total). These two categories overlap almost exactly, all but three of the fear appeals were also flagged as being attack tweets, providing further evidence of the robustness of the coding scheme. The fear appeals were, by some margin, the most likely tweets to be retweeted, being retweeted almost 90% of the time. More general attack tweets got retweeted 74% of the time, again significantly higher than the average (remembering that 50% of the tweets in the sample were retweeted). The only other category which came close in terms of being retweeted was tweets expressing support for others, of which 77% were retweeted. This is easier to explain as generally the person for whom the support is being expressed would then go on to retweet the tweet.

As regards the attack tweets, there are some differences between the three main parties in terms of propensity to use this tactic[37]. Labour MPs were the most likely to use attack tweets with just over 18% of their tweets falling into this category, compared to 17.5% of the Conservatives' tweets and 9% of the Liberal Democrats' tweets. However, these differences are not statistically significant at the 0.05 level, perhaps due to the relatively small sample size in this case (chi-square = 5.069, df 2, p (two-tailed) = .079).

As regards fear appeals, the Conservatives were more likely to use them – 5% of their tweets were in this category compared to just over 2% of Labour's tweets, and this difference is statistically significant (chi-square = 6.088, df 2, p (two-tailed) = .014). There are no Liberal Democrat fear appeals in the sample of tweets coded. This finding reflects the Conservatives' wider campaign strategy in the 2015 General Election in which they made extensive use of fear appeals to persuade people that the economy would not be safe under a Labour government and that a vote for Labour would let in the SNP (Shrimsley, 2015).

There are some gender differences as regards propensity to use attack tweets too – almost 20% of the men's tweets are attacking compared to 11% of the women's (chi-square = 14.336, df 1, p (two-tailed) <.001). A similar pattern can be observed with fear appeals – they account for 4% of the men's tweets compared to 1.5% of the women's (chi-square = 4.753, df 1, p (two-tailed) = .029). Those in safe seats were also significantly more likely to use attack tweets than those in marginal or near marginal seats. Almost 20% of the tweets from MPs in safe seats were attacking tweets, compared to 16% in near marginal seats and 8% in marginal seats (chi-square = 11.493, df 2, p (two-tailed) = .003). Perhaps MPs in safe seats felt freer to indulge in negativity in their tweets whilst those in marginal seats focused more on campaigning on behalf of themselves and presenting a positive face to the world. There were no

---

[37] Note, it is not possible to include the smaller parties in this analysis as too few of their tweets are included in the sample of tweets that was manually coded.

significant differences regarding propensity to use fear appeals according to marginality of seat due to the low sample size.

Further evidence for the robustness of the manual content categories is shown by the fact that 85% of the tweets manually coded as attack tweets were also coded as negative, along with 80% of the fear appeals, even though the sentiment coding and the content category coding were done as two separate, unrelated processes. In contrast the Brandwatch algorithm flagged only 7% of attack tweets as negative, and did not code any of the fear appeal tweets as negative. In both cases, more than 90% were instead coded as neutral.

After local tweets, campaign trail tweets and attacking tweets, the next most common category was personal tweets. Previous research has found that politicians shared very little of themselves online. In Parmelee and Bichard's (2012) research only 8% of American politicians' tweets contained a personal reference and Lawless (2012) found only 5% of Congresspeople's tweets included any personal information. However, over 12% of the tweets considered here contained a personal element – typically references to family, meals out, sport or music. Only 31% of these tweets were retweeted. Labour MPs sent more personal tweets (15.5% of their total tweets) than the Conservatives (9.5%) or Liberal Democrats (7%) (Chi-square = 12.435, df 2, *p* (two-tailed) = .002) suggesting perhaps a more personal campaigning strategy being adopted by Labour MPs. There are no other significant differences between other variables and personal tweets.

A similar number of MPs used Twitter as a tool for publically thanking people, most commonly campaign volunteers and other supporters. Thanking tweets were no more or less likely to be retweeted than average and there were no obvious significant relationships between propensity to send thanking tweets and any other variables.

The next most common category was position-taking tweets and these were slightly more likely to be retweeted than average (61% retweeted compared to 50% in the sample as a whole). There were no significant relationships detected between the sending of position-taking tweets and any other variables. Previous research (Graham

et al., 2013) found that UK MPs made relatively little mention of either their own position or their party's position on issues in their tweets (5.6% and 1.7% of tweets respectively). In this case, nearly 11% of the coded tweets took a position on a political issue, suggesting that perhaps MPs are becoming more confident at using Twitter to express their views.

There are three other types of tweet that have a significant impact on the chances of a tweet being retweeted. They are tweets about events, tweets expressing support for others and tweets in response to something the MP has seen in the media. Event tweets are a particular kind of campaign tweet that make mention of a specific campaign-related event that the MP has attended such as a hustings or debate. Only 38.5% of these tweets are retweeted. However, tweets expressing support for others are highly likely to be retweeted – 77% of them are retweeted – presumably because the person who is mentioned in the tweet will typically retweet it at least once. The last category associated with differences in retweet rates is the media response category. The majority of these tweets related to something seen on television, most referred to the party leaders' debate or to the BBC's Question Time programme. Using a hashtag relating to the media is a good way of exposing a wider audience of people to your tweet and being part of a broader conversation. Clearly MPs recognise this because 71% of the tweets in this category included a hashtag, compared to just 37% of the total tweets in the sample (chi-square = 50.497, df 1, *p* (two-tailed) < .001). Breaking this down further, there were 34 different hashtags used in these media response tweets (bearing in mind that a single tweet can contain more than one hashtag). The most commonly used by far was #leadersdebate used 31 times. There were also another 9 hashtags that were some version of this (e.g. #leaderdebates, #bbcdebate). All but 10 of the hastags used in these tweets contained some mention of either the television debates or of Question Time. One could argue that the use of these hashtags by MPs indicates willingness to engage in discussion with people on issues of relevance to the wider campaign and so provides evidence that not all MPs are operating exclusively in broadcast mode when tweeting. There was no particular pattern discernable when looking at which MPs sent these tweets – in total 47 MPs

were responsible for the 92 tweets in this category, most sending just one or two each. There is one exception to this however – Barry Sheerman – who sent 14 of these tweets (almost 15% of the total), substantially more than his nearest rival, Jeremy Hunt, who sent five.

7.6.   Predicting retweets using CHAID modelling

The fourth objective of the research was to build CHAID models to predict whether tweets would be retweeted or not. Chapter six presents the findings from a series of ten CHAID models, developed using different combinations of variables. In common with other retweet prediction research, this research uses only variables relating to the author of the tweet and the tweet itself to form the basis of its predictions. However, there are two critical differences between this research and other retweet prediction research. The first is that this research uses manually coded content and sentiment variables. Most extant research in this field either does not use content or sentiment variables at all or only uses machine-generated variables. The second difference is that this research only includes in the models variables that are within the tweet author's control. If the aim were simply the build the most predictive model possible then a wider range of variables would be included, but the aim is to provide practical guidelines to help MPs understand how to make better use of Twitter and, as such, there is limited value in including in the models variables over which they have no control.

7.6.1.   The influence of the structure of the tweet

The first CHAID model used only variables relating to the structural characteristics of the tweets and was correct in 59.29% of its predictions. Whether or not a tweet contains a link to a picture or image is the most predictive variable, followed by whether or not it contains a hashtag. Mentions (both whether or not a mention is included, and how many mentions) and the inclusion of any kind of URL are much less predictive but still included in the model. Particular nodes of this model contain very high proportions of retweets, for example tweets without hashtags but with media links and mentioning more than one person are retweeted 80% of the time, and those

with hashtags, including at least one mention and a media link are retweeted 74% of the time. Conversely, tweets without hashtags or media links but which include another kind of link are only retweeted 35% of the time, and those without hashtags or any kind of link but which include at least one mention are only retweeted 36% of the time. This shows that hashtags and media links are the most powerful elements of the tweet's structure when it comes to influencing retweeting.

7.6.2. The influence of the characteristics of the sender

More influential than the structure of the tweet are the characteristics of its sender. Whether a tweet is retweeted or not can be correctly predicted 70% of the time using the mean number of tweets the MP sent per day of the campaign, the number of people they are following, the total number of tweets they have sent since they have been on Twitter and their ratio of followers to followees. Of these, tweets per day is the most predictive. However, because these are all continuous variables, the resulting CHAID model is difficult to interpret, acknowledged as a limitation of this method. CHAID groups continuous variables into clusters, always looking for the groupings with the clearest splits relating to the predictor variable (whether the tweet is retweeted or not). For continuous variables this can lead to some seemingly fairly arbitrary groupings, without an obvious pattern in their relation to the predictor. For example, this model split the data into nine groups according to the differing numbers of tweets sent. Whilst the number of tweets goes up as one moves through the groups, the percentage of retweets does not, so it is not possible to simply look at the model and conclude that more tweets leads to a higher percentage of retweets.

Blending together author variables and tweet structural variables into a single model does not significantly increase the predictive power of the model, suggesting that the characteristics of the tweet's author are more important in determining retweeting than are the structural characteristics of the tweet. Further support for this position comes from the fact that those tweets which attracted the highest numbers of retweets tended not to include any of the structural elements that the CHAID models deemed to be important (hashtags, mentions, links). This also provides further support for the theory that the factors which govern whether a tweet gets retweeted

or not are not the same factors as determine how many times a tweet gets retweeted, in line with extant research (Morchid *et al.*, 2014).

7.6.3. The influence of content and sentiment

The CHAID modelling of manual content and sentiment variables showed that the sentiment of the tweet is by some margin the variable that best predicts retweeting, and that this holds true across models using different combinations of variables. Of the things that the tweeter has control over, sentiment is the most powerfully associated with retweeting. The only variable in the dataset that is more predictive of retweeting is the number of followers one has, which has not been included in most models on the grounds that it is beyond the tweeter's control. Previous attempts to predict retweets have found sentiment to be the least predictive variable (Lemahieu *et al.*, 2015) but when sentiment is used as a variable in extant retweeting research it is always based on machine analysis. The research presented here shows the limitations of machine sentiment analysis as compared to manual analysis when it comes to accurately assessing the valence of MPs' tweets.

Negative tweets in the model were retweeted almost 70% of the time and particular kinds of negative tweets were even more highly retweeted – for example, attack tweets that were also fear appeals were retweeted almost 90% of the time. These percentages indicate that sending a certain kind of negative tweet almost guarantees a retweet. However, it is also possible to get a positive tweet retweeted. All those positive tweets that were also calls to vote were retweeted. Note, a simple neutral call to vote alone is not enough – only 56% of the total call to vote tweets were retweeted. To guarantee retweeting of a call to vote the MP needs to add a positive emotional spin rather than simply passing on the information. Positive tweets expressing support for others were retweeted 77% of the time and 100% of positive campaign trail tweets were retweeted. For neutral tweets, the best way to get them retweeted is to express support for others (retweeted at almost exactly the same rate as the positive support for others tweets, indicating that the important thing here is the expression of support for others rather than the valence of the tweet). Overall, this research provides support for the findings of Berger and Milkman (2012) and

Dang-Xuan *et al.* (2013) that content with high emotionality is more likely to get passed on than more neutral content but contradicts Berger and Milkman's (2012) contention that positive content is more likely to get shared.

7.6.4. Blending factors together

When all the potentially impactful factors that MPs could control were blended together into a single model (model nine), tweet sentiment remained the most predictive variable. Of the ten variables model nine deemed significant, five were derived from the manual content variables, two related to the structure of the tweet (including a link to an image or video, and including a hashtag) and three related to the Twitter behaviour of the sender (number of people following, following / follower ratio and total campaign tweets). The final model, model ten, threw in all possible variables including those that had earlier been excluded because they represent things that the senders of the tweets could not influence. Model ten then showed followers as the most significant predictor. This is not surprising as the number of people that follow you determines to a large extent how many people see your tweet. It is logical to assume that, generally speaking, the more people who see your tweet, the more chance it has of being retweeted. However, sentiment is the second most predictive variable, only a little way behind number of followers in terms of its predictive power. This shows that, whilst MPs cannot influence how many people follow them, they can still have a significant influence over the content of the tweet which is highly predictive of both whether the tweet gets retweeted and how many times it gets retweeted. Adding a hashtag or a media link into the tweet, or using your tweet to express support for others all increase the chances of being retweeted.

7.7.   Demonstrate a new method for predicting retweets

Objective five of this research was to demonstrate a new methodological approach to predicting retweets – CHAID analysis – that could be of use to social media researchers working in other fields or to commercial marketers. The research has demonstrated that CHAID can be used as an effective tool for predicting whether tweets will be retweeted and to identify the factors that are most significant in driving

those predictions. CHAID offers significant benefits over other methods regularly used, primarily regression analysis, machine learning techniques and neural networks. Whilst other decision tree approaches have been used before to predict retweets (Bakshy *et al.*, 2011; Uysal and Croft, 2011; Kupavskii *et al.*, 2012) no examples have been found of CHAID being used this way.

The results of CHAID analysis are generally easy to interpret. This is most clearly the case when the models are built using selections of categorical variables. As has already been discussed, including continuous variables in the models can improve their predictive power (in this case, at least) but makes the results harder to interpret. However, even taking this limitation into account, it is still fair to say that CHAID results are more transparent and easily interpretable by the non-analyst than alternatives such as regression modelling and neutral networks. This is important because if findings from research like this are to have any practical benefit they need to be explainable and for this reason, as Gayo-Avello (2012) argues, 'black box' approaches to prediction should be avoided.

Much extant research is aimed at predicting in advance how likely a particular tweet is to be retweeted so that, for example, Twitter could tailor someone's Twitter feed to show them only those tweets that an algorithm suggests they are likely to find interesting (e.g. Uysal and Croft, 2011; Webberley *et al*., 2016). For models like that there is no need to be able to pick out which factors drive the decision to retweet – all that matters is developing an accurate propensity score. To date, the focus of predictive retweeting literature is really on developing as accurate a model as possible, rather than using the modelling process as a tool to understand which tweets are popular and why. This research has demonstrated how using CHAID can make the results of such modelling much more accessible so that the focus can be on how and why particular tweets are retweeted and others are not, whilst still building models that are predictively accurate, rather than simply on building the most accurate model possible with no thought for the factors that drive the predictions.

CHAID offers a level of flexibility that other predictive techniques do not. This research focuses on predicting a categorical variable (do tweets get retweeted or not) but CHAID works with both categorical and continuous variables so further research could be conducted using the same method to predict the number of retweets tweets get. Both categorical and continuous predictors can be included in a single model. CHAID is also a more robust technique than many others as it does not require data to be normally distributed and, as shown here, continuous variables relating to Twitter behaviour tend to be non-normally distributed.

7.8. Practical advice for politicians wishing to effectively use Twitter

The final aim (objective six) of this research was to provide practical advice for political parties or individual politicians who wish to use Twitter as part of their campaign communication strategy, regarding how to best harness the power of Twitter to engage with citizens. These will be presented in the recommendations section of the final chapter of this thesis.

7.9. Chapter conclusion

This chapter has demonstrated how each of the objectives set for this research has been met. The findings of the research have been related back to the literature and it has been shown how this research contributes methodologically by demonstrating how CHAID predictive models can be used as an effective method of predicting retweets. An additional methodological contribution comes from the comparison of the effectiveness of manual content and sentiment analysis with computer-based coding. This research shows that manual sentiment coding of Twitter data is significantly more robust than either of the machine-based coding sentiment analysis methods that it was compared against.

This chapter also shows how the research contributes to two key areas of debate in the literature – the debate between techno-optimists and techno-pessimists, and the debate about the effectiveness or otherwise of negative campaigning. The evidence presented here tends to provide more support for the techno-pessimist perspective

as, for the most part, it looks as though those who achieve the most success on Twitter in terms of retweets are those politicians who are best-known offline as well. That said, there are perhaps some glimmers of hope for techno-optimists with evidence that smaller parties are able to achieve more in terms of the percentage of their tweets that are retweeted than are larger parties, and there are one or two individual politicians who are able to punch above their weight in terms of measures of retweet success. The bulk of the evidence, however, supports the 'business as usual' techno-pessimist perspective. The most retweeted MPs are the best-known. Of the 57 tweets that achieved more than 1,000 retweets, 86% were sent by either David Cameron or Ed Miliband. No unknown MPs with low public profiles achieved a significant volume of retweets. It is who you are rather than what you say that drives large retweet volumes.

As regards negative campaigning, the results show that negative and attacking tweets are much more likely to be retweeted than positive or neutral ones, as well as to generate a larger volume of retweets. However, at the end of the day context trumps everything – those tweets with the largest numbers of retweets tended not to meet any of the criteria outlined here insofar as they did not include hashtags, links or mentions of other people and yet still achieved massive volumes of retweets.

The following chapter will address the final objective of the research by using these findings as the basis for a set of recommendations for MPs who wish to get more of their tweets retweeted, before outlining the limitations of this research and presenting a more detailed discussion of the contributions to knowledge that it makes.

Chapter 8 Recommendations and conclusions

8.1.   Chapter introduction

This research has demonstrated how CHAID can be used both to predict whether or not politicians' campaign tweets will be retweeted and to better understand which factors drive retweeting in this context. This final chapter briefly summarises the key results before addressing the final research objective (objective six) by translating those results into a series of practical recommendations for MPs wishing to engage people with more of their tweets. The limitations of the research will be discussed before the contributions to knowledge are summarised.

8.2.   Summary of results

The conceptual model developed from the literature (Figure 102) proposed that three broad categories of factor would influence whether politicians' tweets were retweeted, namely the characteristics of the sender of the tweet, of the tweet itself and of the recipient of the tweet. This research has focused on the first two of these.

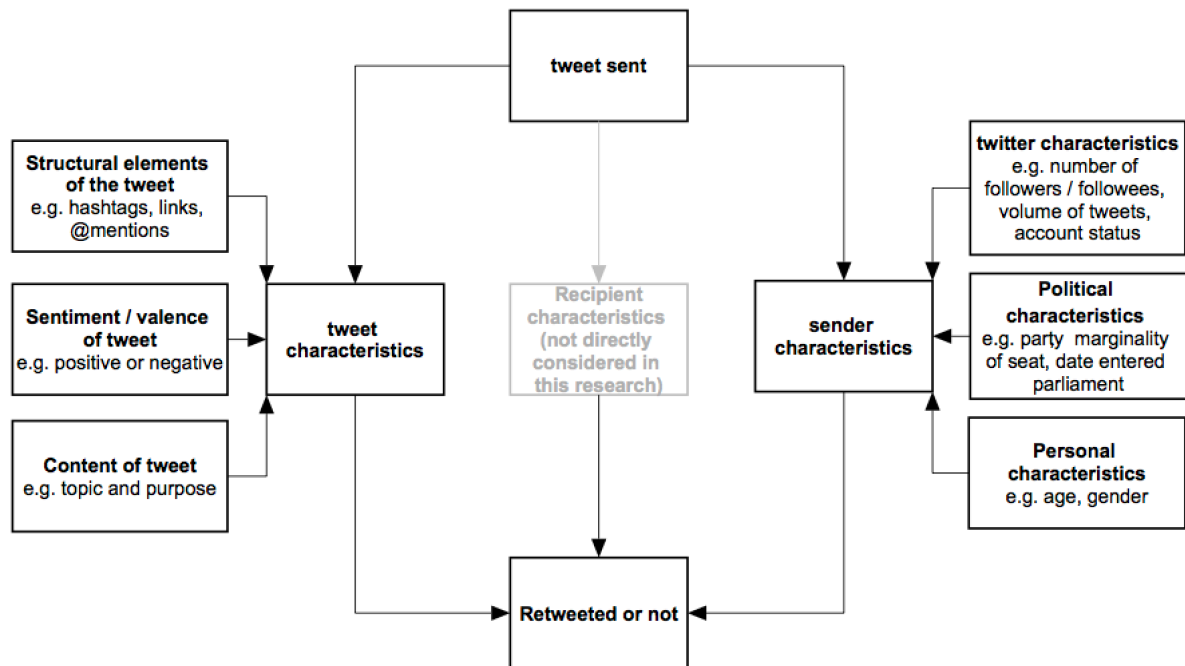*Figure 102 - Conceptual model of factors influencing retweeting*

Table 50 and Table 51 present a summary of how the factors identified in the conceptual model influence retweeting, based on the findings of both the initial descriptive analysis and of the CHAID predictive modelling phase.

Table 50 - How the characteristics of the tweet's sender influence retweeting

Factor	Influence on retweeting	Thesis section
Twitter characteristics	There is a significant relationship between the number of followers an MP has and both the number of retweets per campaign tweet that they achieve and the percentage of their tweets that get retweeted.	Section 6.3.1.2 and 7.5.2.1
	Number of followers is the single most predictive variable of whether a tweet gets retweeted or not.	Section 6.6.6
	There is a significant relationship between the number of people an MP follows and both the number of retweets per campaign tweet they achieve and the percentage of their tweets that get retweeted.	Section 6.3.1.2
	There is a significant relationship between an MP's ratio of followers to followees and both the number of retweets per campaign tweet they achieve and the percentage of their tweets that get retweeted.	Section 6.3.1.2
	The number of tweets sent per day during the campaign is the most predictive of the variables relating to the author that they can control however the relationship between tweets per day and retweets is not linear.	Section 6.5.2
	Tweets sent from verified accounts are more likely to be retweeted than those from non-verified accounts however verification status has no influence on the number of times a tweet gets retweeted	Section 6.3.1.3 and 7.5.2.2
Political characteristics	The sender's party affiliation has a significant influence on the chances of a tweet getting retweeted and on how many times the tweet is retweeted.	Section 6.3.3.1 and 7.5.3.1
	Respect, Ukip and the Green Party got all their tweets retweeted at least once compared to only 47% for the Alliance and 40% for the DUP.	Section 6.3.3.1 and 7.5.3.1
	There are substantial differences between the parties with regard to retweet volumes, with the SNP achieving almost 30 retweets per tweet compared to less than one for the Alliance and Plaid Cymru.	Section 6.3.3.1
	Election outcome is related to the number of retweets – MPs who held their seats achieved a higher median number of retweets per tweet than those who lost their seats or stood down.	Section 6.3.3.2
	There is no significant relationship between safeness of seat and number of retweets generated per tweet.	Section 6.3.3.2
	There is no significant relationship between the year in which an MP entered parliament (their parliamentary cohort) and how likely their tweets are to be retweeted.	Section 6.3.3.3
Personal characteristics	MPs in the age group 35-44 are significantly more active on Twitter than other age groups however age group does not affect how many retweets they get.	Section 6.3.2.2
	Tweets sent by men are slightly more likely to get retweeted than those sent by women, however gender does not influence the number of retweets a tweet gets.	Section 6.3.2.1

Table 51 - How the characteristics of the tweet influence retweeting

Factor	Influence on retweeting	Thesis section
Structural elements	Tweets with hashtags are more likely to be retweeted than those without, and also achieve a higher number of retweets per tweet. Whether or not there is a hashtag in a tweet is the most predictive of the tweet structural variables.	Sections 6.4.1 and 6.5.1
	Using a particular hashtag (generally official party campaign hashtags or those related to popular TV programmes) can virtually guarantee a retweet.	Section 6.4.1 and 6.5.4
	Including at least one mention in a tweet significantly improves its chances of being retweeted.	Section 6.4.2
	Tweets that include links are more likely to get retweeted than those without, and this effect is more pronounced if the link is to an image or a video.	Section 6.4.3
Sentiment	The two sources of machine-generated sentiment variables vary considerably from each other and from the manual content codes, both in terms of how they code the tweets and in terms of how good they are at predicting whether a tweet will be retweeted or not.	Section 7.5.5.1
	Both sets of machine-generated sentiment variables suggest that negative tweets are more likely to get retweeted than positive tweets or neutral tweets.	Section 7.5.5.1
	The manually coded sentiment variable shows that negative tweets are both much more likely to get retweeted than either positive or neutral tweets and that they are more likely to achieve a higher volume of retweets.	Section 7.5.5.1
	Almost 80% of the manually coded negative tweets were retweeted compared to just over 50% of the positive ones (in a sample of tweets split 50/50 between retweeted and not retweeted).	Section 6.6.1
	Negative tweets represent 17% of the tweets in the manually coded sample but comprise 50% of the most retweeted tweets (those with 15 or more retweets).	Section 7.5.5.1
	Sentiment is the most predictive of the manually derived content-related variables.	Section 6.6.4
Content	The content categories attack, fear appeal, media response, position-taking, support for others and campaign trail are all positively associated with retweeting.	Section 6.6.2
	The content categories of personal tweets and event-related tweets are negatively associated with retweeting.	Section 6.6.2
	The tweet's content significantly influences not only whether it gets retweeted but also how many times it gets retweeted.	Section 6.6.2
	After sentiment, the category of attack tweets is the second most predictive content-related category. Tweets that combine an attack with a fear appeal are almost guaranteed to be retweeted.	Section 6.6.3
	Attack tweets represent 42% of the most retweeted tweets (those with 15 or more retweets) compared to 17.2% in the dataset as a whole.	Section 7.5.5.2

8.3. Recommendations for MPs

Objective six of the research was to use the findings to put forward a set of practical recommendations for MPs. In the light of the findings just outlined, practical advice for MPs who wish to boost their chances of getting their tweets retweeted boils down to a few elements.

- The single best way to improve both your chances of getting retweeted and the number of retweets you get is to grow your follower base. Whilst the number of followers you have is not completely within your control, there are steps you can take to improve the situation, for example by following more people yourself (follows are generally reciprocated, particularly if they come from someone's local MP), retweeting other people's tweets (people will often follow people who retweet them), and promoting your Twitter handle as widely as possible across other promotional media (for example, in your LinkedIn profile, on your Facebook page, on any printed literature that you distribute, on banners, posters and badges).

- The number of people that you follow is positively associated with the chances of your tweets getting retweeted and the number of times that they get retweeted, as is your ratio of followers to followees, so it is important not to focus all your efforts on getting more followers but also to make sure that you follow back.

- Tweets with hashtags get retweeted more than those without. Hashtags effectively bring your tweets to the attention of more people than simply those who follow you and so judicious use of hashtags can also form part of an effective strategy to build up follower numbers.

- Use your party's campaign hashtags as well as general election hashtags wherever relevant as these are associated with very high levels of retweeting.

- Build up your Twitter profile and retweet numbers by engaging with other people in discussions relating to media events. Popular political hashtags such

as #BBCQT attract high levels of retweets and put your tweets in front of a wide audience.

- Tweets that include links get more retweets than those without. In particular, links to images or videos are associated with high levels of retweeting.

- Retweeting other people's materials will bring yourself to their attention and increase your follower numbers. Retweets are often reciprocated so it is worth retweeting the tweets of your colleagues as this increases the chances that they will retweet your tweets, thus bringing them to the attention of a wider audience. Some parties do this extremely effectively and manage to achieve 100% retweet rates for their tweets.

- Negative tweets are most likely to be retweeted, particularly attack tweets and fear appeal tweets. However, it is possible to get high numbers of tweets retweeted using other tactics – one does not have to go negative in order to generate retweets (and indeed most MPs do not). Tweets thanking other people and mentioning them are likely to get retweeted (although these types of tweets rarely generate a substantial volume of retweets). Position-taking tweets tend to do better than tweets which do not state a position. Positive tweets calling for people to vote or mentioning attendance at campaign events have high retweet rates. If your goal is purely to generate retweets, then it is best to steer clear of personal tweets as these are negatively associated with retweets. As already discussed, TV response tweets generate high levels of retweets, as do tweets expressing support for someone else.

8.4.   Implications for marketing practice

Although this research focuses on political tweets, the findings could have wider implications for marketing practitioners in other fields. The focus of much social media marketing advice is on manipulating the structural elements of a tweet – be sure to include a hashtag, don't forget to add a link and so on – in order to encourage retweeting. However, this research indicates that the content of the tweet also plays a very significant role in determining which tweets people find sufficiently engaging to

retweet. Brand marketers could benefit from conducting a similar content and sentiment analysis of their own tweets in order to determine what types of content are most effective at stimulating engagement. In particular, the finding that negative tweets are substantially more likely to be retweeted than positive or neutral tweets has implications for brand marketers who are keen to avoid 'Twitter storms' or the spreading of negative word-of-mouth when their customers tweet about them. Social media listening tools such as Brandwatch can be a useful way for companies to keep track of what people are saying about them. If negative brand mentions are more likely to be passed along than positive ones, then that suggests that brands would benefit from understanding what circumstances might motivate customers to send negative tweets in order that they can take steps to prevent those circumstances occurring. For example, if social media listening and content analysis helps a delivery company to understand that late deliveries and rude drivers tend to stimulate customers to tweet negatively, steps can be taken to address these two problems at an operational level.

8.5.   Limitations of the research

8.5.1. Limited generalisability

Results from research based on Twitter cannot be generalised any more widely than Twitter. Twitter users are not representative of the general population and so conclusions about how Twitter users behave cannot be used to predict how the wider population might behave. Research suggests that Twitter users are younger than the general population and, certainly in an American context, that they are more likely to be Democrats than to be Republicans (Gayo-Avello, 2011). However, this weakness of Twitter data in general is not a substantial weakness for research reported in this thesis as no attempt is being made here to generalise beyond Twitter behaviour to the wider population. This research is about how people behave *on Twitter* so a research population entirely comprised of tweets from particular Twitter users is a valid approach, and Twitter users form a community with its own set of practices and which is valuable to examine in its own right (boyd, Golder and Lotan, 2010). That said, there

are of course other limitations to the generalisability of this research. This research examines the factors which influence whether or not politicians' tweets get retweeted. One cannot generalise beyond this to assume that the same factors would be relevant when considering how a non-politician can maximise the chances of their tweets being retweeted, or what might work for a different group of politicians, or what might work during a non-campaign time rather than during campaign time. Ultimately, however, this is inductive research which builds theory rather than tests it so generalisability is not one of its primary aims.

8.5.2. Predicting whether tweets get retweeted rather than how many times

This research focuses on predicting whether a tweet is retweeted rather than on how many times. However, the descriptive analysis phase did shed some light on factors that influence retweet volume. Although most tweets get retweeted at least once, there are MPs and parties in the dataset who perform badly as far as getting any retweets are concerned and would benefit from a greater understanding of what gets retweeted and why. Examining which tweets get retweeted is the starting point for broader research which could then examine the extent to which the factors determined to predict a single retweet can also be used to predict volume of retweets, or whether the factors that influence whether a tweet gets retweeted differ from those that influence how many times it gets retweeted.

8.5.3. Other statistical methods could be more predictive

It is possible that a more predictive model could have been built using different methods of prediction. Prediction levels of 80% or higher have been achieved in extant literature using different methods. However, those researchers often have access to datasets that include millions, even billions of tweets, where the sole aim of the research is to build the model without the need to turn the results into practical recommendations, thus those models use variables relating to things that tweet authors cannot control. In contrast, this research is based on a relatively small number of tweets and aims to shed light on the factors that influence retweets in a specific set of circumstances using only variables that tweet authors can then adapt to improve their performance. This research also differs from extant research insofar as it aims to

demonstrate a method that could potentially be used by non-expert statisticians in other contexts, producing results that are relatively easy to interpret.

8.5.4. Assumption that retweets are a good thing

This research, in common with most others in the field, assumes that retweets are desirable. However, as discussed earlier in this thesis, a large number of retweets is not necessarily a good thing. This research does not make any distinction between desirable and undesirable retweets. That said, all the tweets in this dataset with a large volume of retweets were desirable, meaning it is safe to assume they were retweeted in order to spread their message more widely rather than to mock or make fun. This research also focuses just on automatic retweets (those which occur when the recipient presses the 'retweet' button in their Twitter app) and people who intend to mock the original tweet are much more likely to use the modified retweet method to add commentary to the original tweet in order to make their intention clear. Further research could be conducted looking specifically at what makes politicians' tweet go 'negatively viral'.

8.5.5. Speed of change on the internet

Twitter operates in 'internet time' (Karpf, 2012) which means that it is constantly changing. These findings are a snapshot of the situation at the time when the data was collected. As Twitter changes constantly we cannot assume that the findings relating to the 2015 General Election can be compared to findings relating to the 2010 General Election, nor that they will necessarily shed any light on how people might use Twitter when it comes to the next General Election. However, this is a limitation of all similar research in this field and a general challenge facing researchers working on any aspect of social media or internet communications.

8.5.6. Focus on campaign periods only

It may be that the factors which influence Twitter behaviour during an election campaign differ from those which influence it during non-campaign times. The nature of politicians' tweets may be different when not part of a specific election campaign and, similarly, the levels of engagement they can stimulate from citizens may also

differ outside campaign times. We cannot draw any conclusions about what retweeting politicians' tweets means in terms of other aspects of a citizen's behaviour such as, for example, voting for the politician concerned or taking some other kind of political action on the basis of having engaged with their tweets. However, given that it is not possible to do all the things that one might like to do, focusing on political Twitter use during election times is probably a good place to start as these are the times when it has the most potential to influence the democratic process and also when politicians are most likely to use it.

8.5.7. Focus on behaviour rather than attitudes

This research focuses on better understanding of the factors that determine the extent to which people engage with politicians' tweets. In line with most literature in this field, retweeting is deemed to indicate engagement. The focus is tightly on factors that influence the chances of a tweet being retweeted. This does not tell us anything about what further impact that tweet may have had on the person who retweeted it. We do not know if retweeting and the kind of engagement that implies are linked to a greater propensity to vote for a candidate or engage with politics in other ways. Further research is needed to determine the extent to which Twitter influences politics more broadly as it is beyond the scope of this research to address that question. Additionally, this work is focused on retweets as a measure of Twitter engagement but there are other measures too which are not considered here – in particular liking tweets, mentioning the tweet's sender, modified retweets and adding the tweet's sender to a list. Further research would be required to better understand the role that these measures play.

8.5.8. Focus on MPs standing for re-election

For practical reasons this research examines only the Twitter behaviour of MPs standing for re-election and thus wider conclusions cannot be drawn about the behaviour of challenger candidates in election campaigns. The behaviour of challenger candidates may be different and may reveal other aspects of Twitter behaviour that could correlate with election results.

8.5.9. We do not know who is doing the retweeting

Without knowing who is doing the retweeting, it is hard to evaluate the true value of a retweet as a campaign tool. Clearly there are substantial differences between the parties in terms of retweet volumes and we can speculate that this may be to do with differing social media strategies between the parties – some parties are likely to be more effective at retweeting each others' tweets than others are but it is unclear what value there is in having your tweet retweeted many times by your fellow party members. A retweet from a fellow party member may not be as valuable as a retweet from a constituent or a journalist. In this research, as in most extant research on the topic, the retweet is taken to indicate a measure of engagement with the original tweet but a retweet from a fellow campaign worker probably indicates a lower form of engagement than a retweet from a private citizen with whom one does not have a pre-existing relationship. Further research could look not only at whether the MPs' tweets get retweeted but at the reach of those retweets and whether there are different patterns in terms of who does the bulk of the retweeting.

8.5.10. Coding by purpose rather than topic

The content analysis coding approach taken in this research was based around understanding the purpose and intent behind each tweet. This approach does not include any consideration of the topic of the tweet (tax, NHS, housing, immigration and so on) and these factors may also influence how likely particular tweets are to get retweeted. That said, a relatively small number of tweets in the sample analysed here (just under 11%) include discussion of policy or position-taking so if a meaningful analysis of topic were to be conducted then a larger sample of tweets would be needed.

8.5.11. Other variables could have been used

Both the strength and, in some ways, the weakness of Twitter data is the vast number of possible variables that are available or can be appended to the data or calculated from the existing data. This research uses a large number of variables from a range of different sources. However, there are always more variables that could have been included and every paper one reads includes mention of another variable that the

researchers thought might be relevant but which has not been considered here. For reasons of practicality and time the researcher had to limit the number of variables added into the dataset in order to keep it to a manageable size. The final dataset for analysis here had over 80 variables in it, not including the hundreds of concept and category variables created by SPSS. Even with that number, the researcher spent many days down various rabbit holes looking for possible relationships between different combinations of variables. Had the dataset included many more variables, the project could have gone on forever. The choice of variables to include was largely governed by those that other researchers most commonly used or which seemed to have the most theoretical resonance based on the literature review. To keep the project to a manageable size it was necessary to be fairly ruthless about not adding additional variables past a certain point. However, it is undoubtedly the case that there are other variables out there that could have been appended or calculated which might have improved the predictive power of these models.

8.5.12. No consideration of the role played by the tweet's recipient

This study does not include any consideration of the receiver's characteristics and so does not claim to present a whole picture of the factors that influence whether MPs' tweets are retweeted. This limitation is shared with most research in this field as data related to tweets and their authors is readily available via the Twitter API whereas information about the receivers of those tweets cannot easily be accessed in this way. That said, one goal of this research is to generate practical guidelines that politicians can use in order to be more effective on Twitter. These guidelines can only be useful if they focus on the elements of Twitter that individual tweeters have some control over, namely the content of their tweets and the way they themselves behave on Twitter. Politicians have no control at all over who follows them so an understanding of how followers' characteristics influence retweeting, while interesting in the abstract, would not further inform practical tweeting guidelines. The content of the tweet is the thing that the author has the most direct control over and hence, whilst acknowledging that other factors influence retweeting too, it is content-related variables that are the prime focus of the research.

8.6.   Contribution to knowledge

This research contributes to knowledge both by extending what we know about why particular tweets are retweeted and by demonstrating a new method which can be used to predict retweets and to identify the factors that drive retweeting in particular contexts.

8.6.1.   Demonstration of new method of predicting retweets

This research makes a methodological contribution, demonstrating that CHAID modelling can be used to predict whether tweets are retweeted and to identify those factors that most influence that outcome. An extensive review of the retweet prediction literature has not uncovered any examples of CHAID models used in this way. Here the focus is on predicting whether politicians' tweets will get retweeted, but the same CHAID approach could also be used in other contexts such as by commercial marketers wanting to better understand patterns of brand retweeting. Not only does this research show that CHAID modelling can be used in this new context, but it also demonstrates that CHAID offers some substantial benefits to marketers keen to learn practical lessons from the modelling process rather than simply aiming to build the most predictive models possible. Prime amongst these is that the output from a CHAID model provides clear decision rules which can be relatively easily translated into operational guidelines for marketers wishing to optimise their retweeting.

8.6.2.   Use of manual content and sentiment analysis

There is very limited consideration of tweet content in extant retweet prediction research, beyond use of content-related variables that are easily available from the Twitter API such as whether the tweet contains a hashtag or a link. Where content analysis is used, it is automated. To date, there does not appear to be any use of manual content analysis as part of retweet prediction modelling, certainly not in the context of political tweets. This research provides further evidence that the content of the tweet is critical in determining whether it is retweeted, and builds new methodological knowledge by demonstrating that manually coded content and

sentiment categories can be highly predictive of whether a tweet gets retweeted or not.

8.6.3. Differences between manual and machine-based sentiment analysis

This research provides further evidence that machine-based sentiment analysis has serious limitations, particularly when it comes to assessing the sentiment of tweets. Tweets are extremely short, contain a lot of noise and regularly make use of irony and sarcasm – all things with which machine sentiment analysis struggles. Political tweets in particular require a high level of contextual understanding to determine accurately whether the author intends them to be positive or negative. This research clearly demonstrates the weakness of two different machine sentiment analysis approaches when compared with sentiment analysis done by a human coder. All three sets of variables came to the same conclusion – that negative tweets are more likely to get retweeted than positive or neutral tweets – but the manual content analysis was much more effective at accurately identifying negative tweets and hence was able to demonstrate the strength of this effect much more strongly.

8.6.4. Typology of politicians' tweets

This research contributes to literature considering political use of social media by putting forward a new typology of politicians' tweets. This typology builds on extant research categorising politicians' tweets, most of which was conducted around elections held between 2009 and 2012. The newly proposed typology builds on what already exists to take into account substantial changes in the way that Twitter use has changed, for example by including categories for tweets responding to the mainstream media, a relatively recent Twitter phenomenon and one that extant typologies do not consider at all. Additionally, this research proposes a more refined categorisation of types of negative tweet by distinguishing between general attack tweets and more specific fear appeal tweets.

8.6.5. Retweeting of politicians' tweets

Extant research on politicians' Twitter use primarily focuses on describing the characteristics of politicians' tweets, without really addressing the question of what

then happens to those tweets once sent. This research moves our understanding of politicians' Twitter behaviour forward by demonstrating a clear link between the type of tweet and the chances of it being retweeted. Specifically, it shows that negative tweets are much more likely to get retweeted than positive tweets and that fear appeal tweets, in particular, almost always get retweeted. This is a particularly interesting finding in the light of current debates about the role of fear in political discourse as sparked by the Scottish Independence Referendum campaign and, more recently, the Brexit debate.

8.6.6. Addressing the attitude / behaviour gap in political marketing research

What extant research there is examining how the public respond to politicians' tweets uses interviews and surveys asking people to explain which tweets they responded to and why rather than on measuring which tweets they actually responded to. This research contributes significantly to our understanding of how people respond to politicians' tweets by measuring how they *actually* respond to them rather than how they *say* they respond to them.

8.6.7. Predicting retweets in a specific context

Extant research on predicting retweets tends to use massive samples of randomly selected tweets as the basis for its modelling. This then does not tell us anything about the factors that are likely to influence the retweeting of particular types of tweets such as, in this case, those sent by politicians. This research builds new knowledge about retweeting by using a population of a particular type of tweets as the basis for its analysis, rather than basing its analysis on a random sample of all tweets.

8.6.8. Contribution to negative campaigning literature

This research contributes to knowledge in the area of negative political campaigning by demonstrating firstly that politicians are regularly using negative campaigning tactics such as opponent attacks and fear appeals in their personal tweets, and secondly that these tweets are much more likely to get retweeted than any other kinds of tweets they might send. This contributes to the ongoing debate about

whether negative campaigning tactics alienate voters or not. The fact that these kinds of tweets are by far the most likely to get retweeted suggests that they can be effectively used to engage with citizens on Twitter, although further research is needed to better understand how this then affects other kinds of political behaviour such as voting.

8.6.9. Contribution to literature on virality

Evidence from the literature on virality regarding whether people are more or less likely to pass on positive or negative content is mixed. The research presented here provides clear evidence that as far as politicians' tweets go, negative content works substantially better than positive content, not only at getting a tweet retweeted in the first place also in stimulating a significant volume of retweets.

8.7. Future research

This research opens up the possibility of several fruitful lines of future enquiry.

- **Comparison of modelling methods** – the current dataset could be used to compare the predictive effectiveness of different modelling methods. The focus here is on CHAID for the reasons already explained but the same data could also be used as the basis for model building using other popular methods for the purposes of comparison.

- **Application of CHAID approach to other contexts** – the CHAID method described in this thesis could be deployed in different contexts, political or otherwise.

- **Focus on retweet volume** – further research could be undertaken, potentially with the same dataset, to extend this modelling process to consider retweet volume in more depth than has been possible here, particularly as this research suggests that perhaps different factors drive retweet volume than drive whether or not a tweet gets retweeted at all.

- **Consideration of non-campaign periods** – additional research could consider whether MPs' Twitter use differs depending on whether they are in a

campaign period or not. Additional tweets from the same MPs but collected outside of the election period could be used for comparison.

- **Changes over time** – a fruitful line of enquiry could be to look at how MPs' patterns of tweeting change over time, conducting a longitudinal study using the same methods throughout rather than, as at present, comparing different studies conducted at different times using different methods.

- **Consideration of responders' characteristics** – future research could include more direct consideration of the role played by the characteristics of the recipients of the tweets perhaps by conducting some network analysis or through the means of survey research or qualitative methods.

- **Wider use of new tweet categories** – the typology of tweet types proposed here could be used as the basis of wider research considering the influence of the content of political tweets, for example research examining links between Twitter activity and election outcome or other measures of voter behaviour. Additional research could be conducted using the tweet categories that were not found to be predictive of retweeting here to determine whether perhaps they influence tweet response in other ways (for example by stimulating replies or favourites).

- **Other measures of Twitter engagement** – retweeting is only one measure of Twitter engagement. This dataset could be used for further research examining what factors influence whether people reply to politicians' tweets or favourite them.

- **Use of different sentiment analysis tools** – there is a wide range of different sentiment analysis tools available. Neither SPSS Text Analytics nor Brandwatch are optimised for political communication. This research could be extended by comparing the sentiment scores of other automated sentiment scoring tools that are more specifically focused on scoring political content.

8.8.   Chapter conclusion

This research examines how MPs used Twitter during the 2015 UK General Election with a focus on how effectively they engaged with their followers. Specifically, the research identifies the factors which influenced whether or not politicians' tweets got retweeted. Twitter offers politicians the possibility of campaigning on their own behalves, relatively free from constraints imposed by their party's central office. This freedom comes with both benefits and risks, and a better understanding of how to use Twitter to effectively engage with citizens would be to most politicians' benefit. Techno-optimists would argue that Twitter offers the possibility of a strong voice to politicians who might struggle to get heard on traditional media, however this research provides very limited support for that view. The balance of evidence presented here suggests that in fact Twitter largely represents business as usual – the politicians with the loudest voices on Twitter tend almost exclusively to be those with the highest profiles in traditional media as well.

One aim of the research was to better understand what kind of political content is most likely to be shared by MPs' followers and here the results are unequivocal. Negative tweets, in particular those which attack others or which employ fear appeals, are by some margin both the most likely to be retweeted and to attract a significant volume of retweets. There is, however, some hope for those that despair at the thought of political social media communication descending into a sea of relentless negativity. MPs currently using these tactics are in the minority – most MPs did not send any negative, attacking or fear-based tweets. Additionally, the CHAID modelling process revealed that it is equally possible to get positive tweets retweeted if they are carefully crafted. The thought that negative content works best on Twitter may be a depressing one, but it also suggests that the view of demobilisation theorists that negative content leads to people becoming disengaged is not correct – on the contrary, in the dataset examined here negative content leads to the highest levels of engagement and in these times of ever lower election turnouts, surely anything that encourages political engagement from the public must ultimately be a good thing.

References

Achananuparp, P., Lim, E.-P., Jiang, J. and Hoang, T.-A. (2012) 'Who is retweeting the tweeters? Modeling, originating and promoting behaviors in the Twitter network', *ACM Transactions on Management Information Systems*, 3(3), p. 13:1–13:30. doi: 10.1145/2361256.2361258.

Ackroyd, S. and Fleetwood, S. (2000) 'Realism in contemporary organisation and management studies', in Ackroyd, S. and Fleetwood, S. (eds) *Realist perspectives on management and organisations*. London: Routledge, pp. 3–23.

Adams, A. and McCorkindale, T. (2013) 'Dialogue and transparency: a content analysis of how the 2012 presidential candidates used Twitter', *Public Relations Review*. Elsevier Inc., 39(4), pp. 357–359. doi: 10.1016/j.pubrev.2013.07.016.

Adi, A., Erickson, K. and Lilleker, D. G. (2013) 'Elite tweets: analyzing the Twitter communication patterns of Labour Party peers in the House of Lords', *Policy and Internet*, 6(1), pp. 1–27. doi: 10.1002/1944-2866.POI350.

Agarwal, A., Xie, B., Vovsha, I., Rambow, O. and Passonneau, R. (2011) 'Sentiment analysis of Twitter data', in *Proceedings of the Workshop on Language in Social Media (LSM 2011)*. Portland, Oregon, pp. 30–38. Available at: http://dl.acm.org/citation.cfm?id=2021109.2021114%5Cnpapers3://publication/uuid/83CA53FE-43D1-4BD5-BCF2-D55B82CF0F99.

Agranoff, C. and Tabin, H. (2011) *Socially elected: how to win elections using social media*. Pendant Publishing.

Ahn, H. and Park, J.-H. (2015) 'The structural effects of the sharing function on Twitter networks: focusing on the retweet function', *Journal of Information Science*, 41(2), pp. 354–365. doi: 10.1177/1045389X14554132.

Ampofo, L., Anstead, N. and O'Loughlin, B. (2012) 'Trust, confidence, credibility – citizen responses on Twitter to opinion polls during the 2010 UK General Election', in Loader, B. D. and Mercea, D. (eds) *Social media and democracy – innovations in participatory politics*. Abingdon: Routledge.

Ansolabehere, S., Iyengar, S., Simon, A. and Valentino, N. (1994) 'Does attack advertising demobilize the electorate?', *American Political Science Review*, 88(4), pp. 829–838. Available at: http://www.jstor.org/stable/2082710.

Antonetti, P., Baines, P. and Walker, L. (2015) 'From elicitation to consumption: assessing the longitudinal effectiveness of negative emotional appeals in social marketing', *Journal of Marketing Management*, 31(9–10), pp. 940–969. doi: 10.1080/0267257X.2015.1031266.

Apté, C., Damerau, F. and Weiss, S. M. (1994) 'Automated learning of decision rules for text categorization', *ACM Transactions on Information Systems*, 12(3), pp. 233–

251. doi: 10.1145/183422.183423.

Bachmann, I. and de Zuniga, H. G. (2013) 'News platform preference as a predictor of political and civic participation', *Convergence: The International Journal of Research into New Media Technologies*, 19(4), pp. 496–512. doi: 10.1177/1354856513493699.

Bakshy, E., Hofman, J., Mason, W. and Watts, D. (2011) 'Everyone's an influencer: quantifying influence on Twitter', in *Proceedings of the Fourth ACM International Conference on Web Search and Data Mining*. Hong Kong, pp. 65–74. doi: 10.1145/1935826.1935845.

Berger, J. and Milkman, K. L. (2010) *Social transmission, emotion and the virality of online content*. Available at: http://robingandhi.com/wp-content/uploads/2011/11/Social-Transmission-Emotion-and-the-Virality-of-Online-Content-Wharton.pdf.

Berger, J. and Milkman, K. L. (2012) 'What makes online content viral?', *Journal of Marketing Research*, 49(2), pp. 192–205. doi: http://dx.doi.org/10.1509/jmr.10.0353.

Bhaskar, R. (2008) *A realist theory of science*. 2nd edn. Abingdon: Routledge.

Bienkov, A. (2012) *Astroturfing: what is it and why does it matter?*, *The Guardian*. Available at: http://www.theguardian.com/commentisfree/2012/feb/08/what-is-astroturfing (Accessed: 6 April 2016).

Bimber, B. and Copeland, L. (2013) 'Digital media and traditional political participation over time in the US', *Journal of Information Technology & Politics*, 10(2), pp. 125–137. doi: 10.1080/19331681.2013.769925.

Bimber, B. and Davis, R. (2003) *Campaigning online: the internet and US elections*. New York: Oxford University Press.

Bode, L. and Dalrymple, K. E. (2014) 'Politics in 140 characters or less: campaign communication, network interaction and political participation on Twitter', *Journal of Political Marketing*. doi: 10.1080/15377857.2014.959686.

Boffey, D. (2014) *Ukip tells members 'don't go on Twitter'*, *The Observer*. Available at: http://www.theguardian.com/politics/2014/dec/20/ukip-tells-members-dont-go-twitter-nigel-farage (Accessed: 24 January 2015).

Bond, R. M., Fariss, C. J., Jones, J. J., Kramer, A. D. I., Marlow, C., Settle, J. E. and Fowler, J. H. (2012) 'A 61-million-person experiment in social influence and political mobilization', *Nature*. Nature Publishing Group, 489(7415), pp. 295–298. doi: 10.1038/nature11421.

Bosker, B. and Grandoni, D. (2012) *9 quirkiest facts about Twitter: gazing into the soul of the Twittersphere*, *The Huffington Post*. Available at: http://www.huffingtonpost.com/2012/10/10/quirkiest-facts-twitter-users_n_1956260.html (Accessed: 17 May 2013).

Boyd, D., Golder, S. and Lotan, G. (2010) 'Tweet, tweet, retweet: conversational

aspects of retweeting on Twitter', in *Proceedings of the Annual Hawaii International Conference on System Sciences*. Kauai, Hawaii. doi: 10.1109/HICSS.2010.412.

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), pp. 388–405. doi: 10.2307/3647684.

Bruns, A. and Burgess, J. (2012) 'Researching news discussion on Twitter', *Journalism Studies*, 13(February 2015), pp. 801–814. doi: 10.1080/1461670X.2012.664428.

Buch-Hansen, H. (2013) 'Social network analysis and critical realism', *Journal for the Theory of Social Behaviour*, 44(3), pp. 306–325. doi: 10.1111/jtsb.12044.

Burgess, J. and Bruns, A. (2012) '(Not) the Twitter election: the dynamics of the #ausvotes conversation in relation to the Australian media ecology', *Journalism Practice*, 6, pp. 1–40. doi: 10.1080/17512786.2012.663610.

Calantone, R. J. and Warshaw, P. R. (1985) 'Negating the effects of fear appeals in election campaigns', *Journal of Applied Psychology*, 70(4), pp. 627–633.

Canhoto, A. I. and Padmanabhan, Y. (2015) ' 'We (don't) know how you feel' – a comparative study of automated vs. manual analysis of social media conversations', *Journal of Marketing Management*, 31(9–10), pp. 1141–1157. doi: 10.1080/0267257X.2015.1047466.

Cha, M., Benevenuto, F., Haddadi, H. and Gummadi, K. (2012) 'The world of connections and information flow in Twitter', *IEEE Transactions on Systems, Man and Cynernetics - Part A: Systems and Humans*, 42(4), pp. 991–998. doi: 10.1109/TSMCA.2012.2183359.

Cha, M., Haddadi, H., Benevenuto, F. and Gummadi, K. P. (2010) 'Measuring user influence in Twitter: the million follower fallacy', in *Proceedings of the Fourth International AAAI Conference on Weblogs and Social Media*, pp. 10–17.

Chen, G. M. (2011) 'Tweet this: a uses and gratifications perspective on how active Twitter use gratifies a need to connect with others', *Computers in Human Behavior*. Elsevier Ltd, 27(2), pp. 755–762. doi: 10.1016/j.chb.2010.10.023.

Chevalier, J. A. and Mayzlin, D. (2006) 'The effect of word-of-mouth on sales: online book reviews', *Journal of Marketing Research*, 43(3), pp. 345–354. doi: 10.1509/jmkr.43.3.345.

Chou, H.-Y., Lien, N.-H., H-Y, C., N-H, L., Chou, H.-Y. and Lien, N.-H. (2011) 'What does a negative political ad really say? The effects of different content dimensions', *Journal of Marketing Communications*, 17(4), pp. 281–295. doi: 10.1080/13527260903546213.

Coakley, J. (2008) 'Militant nationalist electoral support: a measurement dilemma', *International Journal of Public Opinion Research*, 20(2), pp. 224–236. doi: 10.1093/ijpor/edn014.

Collier, A. (1994) *Critical realism: an introduction to Roy Bhaskar's philosophy*. London: Verso.

Compston, S. (2014) *Identifying and understanding retweets*, *Gnip*. Available at: http://support.gnip.com/articles/identifying-and-understanding-retweets.html (Accessed: 6 April 2016).

Conover, M. D., Ratkiewicz, J., Francisco, M., Gonc, B., Flammini, A. and Menczer, F. (2011) 'Political polarization on Twitter', in *International AAAI Conference on Web and Social Media*. North America. Available at: http://www.aaai.org/ocs/index.php/ICWSM/ICWSM11/paper/view/2847.

Conway, M. (2006) 'The subjective precision of computers: a methodological comparison with human coding in content analysis', *Journalism & Mass Communication Quarterly*, 83(1), pp. 186–200. doi: 10.1177/107769900608300112.

Cook, L. (2013) *How is social media changing politics?*, *Friends of Europe: debating Europe*. Available at: http://www.friendsofeurope.org/media/uploads/2014/10/FoE-Report-SocialMedia-2013-Web.pdf.

Cui, G., Lui, H.-K. and Guo, X. (2012) 'The effect of online consumer reviews on new product sales', *International Journal of Electronic Commerce*, 17(1), pp. 39–58. doi: 10.2753/JEC1086-4415170102.

Curtis, S. (2014) *Technology firms explore 'second screen' opportunity*, *The Telegraph*. Available at: http://www.telegraph.co.uk/technology/news/10815317/Technology-firms-explore-second-screen-opportunity.html (Accessed: 9 August 2015).

Dang-Xuan, L., Stieglitz, S., Wladarsch, J. and Neuberger, C. (2013) 'An investigation of influentials and the role of sentiment in political communication on Twitter during election periods', *Information, Communication & Society*, 16(5), pp. 795–825. doi: 10.1080/1369118X.2013.783608.

Davies, R. (2014) *Briefing: social media in election campaigning*. Available at: http://www.europarl.europa.eu/RegData/bibliotheque/briefing/2014/140709/LDM_BRI(2014)140709_REV1_EN.pdf.

Dean, D. (2005) 'Fear, negative campaigning and loathing: the case of the UK election campaign', *Journal of Marketing Management*, 21(9), pp. 1067–1078. doi: 10.1362/026725705775194111.

Dean, D., Croft, R. and Pich, C. (2015) 'Toward a conceptual framework of emotional relationship marketing: an examination of two UK political parties', *Journal of Political Marketing*, 14, pp. 19–34. doi: 10.1080/15377857.2014.990849.

Dell (2013) *Statistics: methods and applications*, *documents.software.dell.com*. Tulsa: Dell. Available at: http://documents.software.dell.com/statistics/current/textbook (Accessed: 12 January 2016).

Dermody, J. and Scullion, R. (2000) 'Perceptions of negative political advertising: meaningful or menacing? An empirical study of the 1997 British General Election

campaign', *International Journal of Advertising*, 19(2), pp. 201–223. doi: 10.1080/02650487.2000.11104795.

Devlin, R., Reeves, A., Assibey, A. and Madlani, R. (2015) *The class of 2015: examining the background of the newly elected MPs*, *Industry and Parliament Trust*. Available at: http://www.ipt.org.uk/the-class-of-2015-examining-the-background-of-the-newly-elected-mps.aspx (Accessed: 24 July 2016).

Dillard, J. P. and Anderson, J. W. (2004) 'The role of fear in persuasion', *Psychology & Marketing*, 21(11), pp. 909–926. doi: 10.1002/mar.20041.

Druckman, J. N., Kifer, M. J. and Parkin, M. (2010) 'Timeless strategy meets new medium: going negative on congressional campaign websites, 2002–2006', *Political Communication*, 27(1), pp. 88–103. doi: 10.1080/10584600903502607.

Dunne, Á., Lawlor, M.-A. and Rowley, J. (2010) 'Young people's use of online social networking sites – a uses and gratifications perspective', *Journal of Research in Interactive Marketing*, 4(1), pp. 46–58. doi: 10.1108/17505931011033551.

Easterby-Smith, M., Thorpe, P. R., Jackson, P. P. and Lowe, D. A. (2008) *Management research: theory and practice*. 3rd edn. London: Sage Publications Ltd.

Eckler, P. and Bolls, P. (2011) 'Spreading the virus: emotional tone of viral advertising and its effect on forwarding intentions and attitudes', *Journal of Interactive Advertising*, 11(2), p. 1. doi: 10.1080/15252019.2011.10722180.

Effing, R., Hillegersberg, J. Van and Huibers, T. (2011) 'Social media and political participation: are Facebook, Twitter and YouTube democratizing our political systems?', in Tambouris, E., Macintosh, A., and De Bruijn, H. (eds) *Electronic Participation: Third IFIP WG 8.5 International Conference, ePart 2011. Proceedings.* Delft, The Netherlands, August 29 - September 1: Springer Berlin Heidelberg, pp. 25–35. Available at: http://link.springer.com/chapter/10.1007%252F978-3-642-23333-3_3.

Eltantawy, N. and Wiest, J. B. (2011) 'The Arab Spring: social media in the Egyptian revolution: reconsidering resource mobilization theory', *International Journal of Communication*, 5, p. 18. Available at: http://ijoc.org/index.php/ijoc/article/view/1242.

Enge, E. (2014) *Twitter engagement unmasked: a study of more than 4M tweets*, *Stonetemple.com*. Available at: https://www.stonetemple.com/twitter-engagement-umasked/ (Accessed: 20 May 2016).

Enli, G. S. and Skogerbø, E. (2013) 'Personalized campaigns in party-centred politics', *Information, Communication & Society*, 16(5), pp. 757–774. doi: 10.1080/1369118X.2013.782330.

Esomar (2011) *Esomar guideline on social media research*. Available at: https://www.esomar.org/uploads/public/knowledge-and-standards/codes-and-guidelines/ESOMAR-Guideline-on-Social-Media-Research.pdf.

Faber, R. J., Tims, A. R. and Schmitt, K. G. (1993) 'Negative political advertising and voting intent: the role of involvement and alternative information sources', *Journal of Advertising*, 22(4), pp. 67–76. Available at: http://www.jstor.org/stable/4188900.

Fast, A., Dursun, D., Hill, T., Elder, J., Miner, G. and Nisbet, B. (2012) *Practical text mining and statistical analysis for non-structured text data applications*. Oxford: Academic Press.

Field, A. (2009) *Discovering statistics using SPSS*. 3rd edn. London: Sage.

Finkel, S. E. and Geer, J. G. (1998) 'A spot check: casting doubt on the demobilizing effect of attack advertising', *American Journal of Political Science*, 42(2), pp. 573–595. doi: 10.2307/2991771.

Frequently asked questions: MPs (no date) *parliament.uk*. Available at: http://www.parliament.uk/about/faqs/house-of-commons-faqs/members-faq-page2/ (Accessed: 24 July 2016).

Fridkin, K. L. and Kenney, P. (2011) 'Variability in citizens' reactions to different types of negative campaigns', *American Journal of Political Science*, 55(2), pp. 307–325. doi: 10.1111/j.1540-5907.2010.00494.x.

Gainous, J. and Wagner, K. (2014) *Tweeting to power: the social media revolution in American politics*. Kindle edi. New York: Oxford University Press.

Gayo-Avello, D. (2011) 'Don't turn social media into another "literary digest" poll', *Communications of the ACM*, 54(10), pp. 121–128. doi: 10.1145/2001269.2001297.

Gayo-Avello, D. (2012) '"I wanted to predict elections with Twitter and all I got was this lousy paper" - a balanced survey on election prediction using Twitter data', *Computers and Society*, pp. 1–13. Available at: http://arxiv.org/abs/1204.6441.

Gleibs, I. H. (2014) 'Turning virtual public spaces into laboratories: thoughts on conducting online field studies using social network sites', *Analyses of Social Issues and Public Policy*, 14(1), pp. 352–370. doi: 10.1111/asap.12036.

Go, A., Bhayani, R. and Huang, L. (2009) 'Twitter sentiment classification using distant supervision', *Processing*, 150, pp. 1–6. doi: 10.1016/j.sedgeo.2006.07.004.

Goel, S., Anderson, A., Hofman, J. and Watts, D. (2015) 'The structural virality of online diffusion', *Management Science*, 62(1), pp. 180–196. doi: 10.1145/2229012.2229058.

Golbeck, J., Grimes, J. M. and Rogers, A. (2010) 'Twitter use by the US Congress', *Journal of the American Society for Information Science and Technology*, 61(8), pp. 1612–1621. doi: 10.1002/asi.

Graham, T., Broersma, M. and Hazelhoff, K. (2013) 'Twitter as an instrument for connected representation', in Gerodimos, R., Jackson, D., Lilleker, D., and Scullion, R. (eds) *The media, political participation and empowerment*. London: Routledge.

Graham, T., Broersma, M., Hazelhoff, K. and Haar, G. van't (2013) 'Between broadcasting political messages and interacting with voters: the use of Twitter during

the 2010 UK General Election campaign', *Information, Communication & Society*, 16(5), pp. 692–716. doi: 10.1080/1369118X.2013.785581.

Grant, W. J., Moon, B. and Busby Grant, J. (2010) 'Digital dialogue? Australian politicians' use of the social network tool Twitter', *Australian Journal of Political Science*, 45(4), pp. 579–604. doi: 10.1080/10361146.2010.517176.

Hall Jamieson, K. (1992) *Dirty politics: deception, distraction and democracy*. New York: Oxford University Press.

Hansen, L. K., Arvidsson, A., Nielsen, F. A., Colleoni, E. and Etter, M. (2011) 'Good friends, bad news: affect and virality in Twitter', in *Proceedings of the International Workshop on Social Computing, Network and Services*. Crete, Greece, pp. 1–14. Available at: https://arxiv.org/abs/1101.0510.

Helm, T. (2013) *Obama's tech wizard hones his media weapons to help Miliband get to No 10*, *The Guardian*. Available at: http://www.theguardian.com/politics/2013/nov/30/obama-miliband-no-10-mitt-romney-labour-campaign (Accessed: 30 November 2013).

Hemphill, L., Otterbacher, J. and Shapiro, M. (2013) 'What's Congress doing on Twitter?', *Proceedings of the 2013 Conference on Computer-supported Cooperative Work - CSCW '13*. New York, USA: ACM Press, p. 877. doi: 10.1145/2441776.2441876.

Hennig-Thurau, T., Gwinner, K. P., Walsh, G. and Gremler, D. D. (2004) 'Electronic word-of-mouth via consumer-opinion platforms: what motivates consumers to articulate themselves on the internet?', *Journal of Interactive Marketing*, 18(I), pp. 38–52. doi: 10.1002/dir.10073.

Hennig-Thurau, T., Wiertz, C. and Feldhaus, F. (2014) 'Does Twitter matter? The impact of microblogging word of mouth on consumers' adoption of new movies', *Journal of the Academy of Marketing Science*, pp. 375–394. doi: 10.1007/s11747-014-0388-3.

Hermida, A., Lewis, S. C. and Zamith, R. (2012) 'Sourcing the Arab spring: a case study of Andy Carvin's sources during the Tunisian and Egyptian revolutions', in *International Symposium on Online Journalism*. Austin, Texas. Available at: https://online.journalism.utexas.edu/2012/papers/Hermida.pdf.

Himmelweit, H. T., Biberian, M. J. and Stockdale, J. (1978) 'Memory for past vote: implications of a study of bias in recall', *British Journal of Political Science*, 8(3), pp. 365–375. doi: 10.1017/S0007123400001411.

Hirschl, T. A., Booth, J. G. and Glenna, L. L. (2009) 'The link between voter choice and religious identity in contemporary society: bringing classical theory back in', *Social Science Quarterly*, 90(4), pp. 927–944. doi: 10.1111/j.1540-6237.2009.00670.x.

Ho, R. (1998) 'The intention to give up smoking: disease versus social dimensions', *Journal of Social Psychology*, 138(3), pp. 368–380. doi: 10.1080/00224549809600389.

Holsti, O. R. (1969) *Content analysis for the social sciences and humanities*. Addison-

Wesley.

Homer, P. M. and Batra, R. (1994) 'Attitudinal effects of character-based versus competence-based negative political communications', *Journal of Consumer Psychology*, 3(2), pp. 163–185. doi: 10.1016/S1057-7408(08)80003-4.

Hong, L., Dan, O. and Davison, B. D. (2011) 'Predicting popular messages in Twitter', in *Proceedings of the 20th International Conference on Worldwide Web (WWW)*. Hyderabad, India, pp. 57–58. doi: 10.1145/1963192.1963222.

Hunt, S. D. (1990) 'Truth in marketing theory and research', *Journal of Marketing*, 54(3), pp. 1–15. doi: 10.2307/1251812.

IBM (2012) *Introduction to IBM SPSS Modeler and data mining student guide*. IBM.

IBM (2013) *Introduction to IBM SPSS Modeler Text Analytics*. 15th edn. IBM.

IBM (2014) *Predicting continuous targets using SPSS Modeler*. IBM SPSS.

Jackson, N. and Lilleker, D. (2011) 'Microblogging, constituency service and impression management: UK MPs and the use of Twitter', *The Journal of Legislative Studies*, 17(1), pp. 86–105. doi: 10.1080/13572334.2011.545181.

Jansen, B. J. and Zhang, M. (2009) 'Twitter power: tweets as electronic word of mouth', *Journal of the American Society for Information Science and Technology*, 60(11), pp. 2169–2188. doi: 10.1002/asi.

Jin, S.-A. A. and Phua, J. (2014) 'Following celebrities' tweets about brands: the impact of Twitter-based electronic word-of-mouth on consumers' source credibility perception, buying intention and social identification with celebrities', *Journal of Advertising*, 43(2), pp. 181–195. doi: 10.1080/00913367.2013.827606.

Kaid, L. L. and Johnston, A. (2001) *Videostyle in presidential campaigns: style and content of televised political advertising*. Westport: Praeger Publishers.

Karpf, D. (2012) 'Social science research methods in internet time', *Information, Communication & Society*, 15(5), pp. 639–661. doi: 10.1080/1369118X.2012.665468.

Kelly, H. (2014) *The power of one wrong tweet*, *CNN*. Available at: http://edition.cnn.com/2013/04/23/tech/social-media/tweet-ripple-effect/ (Accessed: 11 February 2015).

Kereszturi, G. (2014) 'The impact of social media on online political engagement using sentiment analysis in the 2014 Romanian presidential elections', in *Academy of Marketing conference*. Limerick, Ireland. Available at: http://programme.exordo.com/am2015/delegates/presentation/293/.

Kim, E., Sung, Y. and Kang, H. (2014) 'Brand followers' retweeting behavior on Twitter: how brand relationships influence brand electronic word-of-mouth', *Computers in Human Behavior*, 37, pp. 18–25. doi: 10.1016/j.chb.2014.04.020.

King, J. D. and McConnell, J. B. (2003) 'The effect of negative campaign advertising on

vote choice: the mediating influence of gender', *Social Science Quarterly*, 84(4), pp. 843–859. doi: 10.1046/j.0038-4941.2003.08404012.x.

King, R. A., Racherla, P. and Bush, V. D. (2014) 'What we know and don't know about online word-of-mouth: a review and synthesis of the literature', *Journal of Interactive Marketing*. Elsevier B.V., 28(3), pp. 167–183. doi: 10.1016/j.intmar.2014.02.001.

Klotz, R. (1998) 'Virtual criticism: negative advertising on the internet in the 1996 Senate races', *Political Communication*, 15(3), pp. 347–365. doi: 10.1080/105846098198939.

Kotler, P., Levy, S. J. and Luck, D. J. (1969) 'Broadening the concept of marketing too far', *Journal of Marketing*, 33(1), pp. 53–63. Available at: http://www.ncbi.nlm.nih.gov/pubmed/12309673.

Kouloumpis, E., Wilson, T. and Moore, J. (2011) 'Twitter sentiment analysis: the good, the bad and the OMG!', in *Proceedings of the 5th International Conference on Weblogs and Social Media*. Barcelona, Spain, pp. 538–541. Available at: http://www.aaai.org/ocs/index.php/ICWSM/ICWSM11/paper/viewPDFInterstitial/2857/3251.

Kred.com (no date) *Kred.com*. Available at: http://kred.com/rules (Accessed: 12 February 2015).

Krippendorff, K. (2013) *Content analysis: an introduction to its methodology*. 3rd edn. Thousand Oaks, California: Sage.

Kupavskii, A., Ostroumova, L., Umnov, A., Usachev, S., Serdyukov, P., Gusev, G. and Kustarev, A. (2012) 'Prediction of retweet cascade size over time', *Proceedings of the 21st ACM International Conference on Information and Knowledge Management*, pp. 2335–2338. doi: 10.1145/2396761.2398634.

Labour Shadow Cabinet and ministers' resignations - the letters in full (2016) *The Telegraph*. Available at: http://www.telegraph.co.uk/news/0/labour-shadow-cabinet-resignations---the-letters-in-full/ (Accessed: 23 July 2016).

Larsson, A. O. and Hallvard, M. (2011) 'Who tweets? Tracking microblogging use in the 2010 Swedish election campaign', in *Proceedings of ECIS*. Shanghai, China. Available at: http://aisel.aisnet.org/ecis2011/.

Larsson, A. O. and Kalsnes, B. (2014) '"Of course we are on Facebook" - use and non-use of social media among Swedish and Norwegian politicians', *European Journal of Communication*, 29(6), pp. 653–667. doi: 10.1177/0267323114531383.

Lau, R. R., Sigelman, L., Heldman, C. and Babbitt, P. (1999) 'The effects of negative political advertisements: a meta-analytic assessment', *The American Political Science Review*, 93(4), pp. 851–875. doi: 10.2307/2586117.

Lau, R. R., Sigelman, L. and Rovner, I. B. (2007) 'The effects of negative political campaigns: a meta-analytic reassessment', *Journal of Politics*, 69(4), pp. 1176–1209. doi: 10.1111/j.1468-2508.2007.00618.x.

Lawless, J. L. (2012) 'Twitter and Facebook: new ways for members of Congress to send the same old messages?', in Fox, R. L. and Ramos, J. M. (eds) *Citizens, elections and governing in the new media era*. Cambridge: Cambridge University Press, pp. 206–232. doi: http://dx.doi.org/10.1017/CBO9781139059893.013.

Lazarsfeld, P., Berelson, B. and Gaudet, H. (1968) *The people's choice: how the voter makes up his mind in a presidential campaign*. 3rd edn. Columbia University Press.

Lazarsfeld, P. and Katz, E. (1955) *Personal influence: the part played by people in the flow of mass communications*. New York: The Free Press.

Lee, B. (2014) 'Window dressing 2.0: constituency-level web campaigns in the 2010 UK General Election', *Politics*, 34(1), pp. 45–57. doi: 10.1111/1467-9256.12029.

Lee, E.-J. and Shin, S. Y. (2014) 'When the medium is the message: how transportability moderates the effects of politicians' Twitter communication', *Communication Research*, 14(8), pp. 1088–1110. doi: 10.1177/0093650212466407.

Lee, K. (2015) *What 1 million tweets taught us about how people tweet successfully*, *buffersocial*. Available at: https://blog.bufferapp.com/twitter-data-1-million-tweets (Accessed: 6 April 2015).

Leetaru, K. H. (2012) *Data mining methods for the content analyst*. Abingdon: Routledge.

Lemahieu, R., Van Canneyt, S., De Boom, C. and Dhoedt, B. (2015) 'Optimizing the popularity of Twitter messages through user categories', in *2015 IEEE International Conference on Data Mining Workshop (ICDMW)*. Atlantic City, New Jersey, pp. 1396–1401. doi: 10.1109/ICDMW.2015.39.

Lewis, K. (2014) *The Lidl surprises that keep on giving*, *Intermarketing.co.uk*. Available at: http://www.intermarketing.co.uk/blog/item/lidl-surprises-keep-giving/ (Accessed: 24 January 2015).

Lewis, S. C., Zamith, R. and 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), pp. 34–52. doi: 10.1080/08838151.2012.761702.

Lilleker, D. G. and Koc-Michalska, K. (2013) 'Online political communication strategies: MEPs, e-representation and self-representation', *Journal of Information Technology & Politics*, 10(2), pp. 1–40. doi: 10.1080/19331681.2012.758071.

Liu, Y., Kliman-Silver, C. and Mislove, A. (2014) 'The tweets they are a-changin': evolution of Twitter users and behavior', *Proceedings of the Eighth International AAAI Conference on Weblogs and Social Media*, pp. 305–314. Available at: http://www.aaai.org/ocs/index.php/ICWSM/ICWSM14/paper/download/8043/8131.

Liu, Z., Liu, L. and Li, H. (2012) 'Determinants of information retweeting in microblogging', *Internet Research*, 22(4), pp. 443–466. doi: 10.1108/10662241211250980.

Lotan, G., Graeff, E., Ananny, M., Gaffney, D., Pearce, I. and Boyd, D. (2011) 'The revolutions were tweeted: information flows during the 2011 Tunisian and Egyptian revolutions', *International Journal of Communications*, 5, pp. 1375–1406. Available at: http://ijoc.org/index.php/ijoc/article/view/1246/613.

Ma, L., Lee, C. S. and Goh, D. H.-L. (2013) 'Investigating influential factors influencing users to share news in social media: a diffusion of innovations perspective', in Downie, J. S. and McDonald, R. H. (eds) *Proceedings of the 13th ACM/IEEE-CS Conference on Digital Libraries*. Indianapolis, pp. 403–404. Available at: http://dl.acm.org/citation.cfm?id=2467696.

Mahmud, J., Chen, J. and Nichols, J. (2014) *Why are you more engaged? Predicting social engagement from word use*, *Computing Research Repository*. Available at: http://arxiv.org/abs/1402.6690.

Malhotra, A., Malhotra, C. K. and See, A. (2012) 'How to get your messages retweeted', *MIT Sloan Management Review*, 53(Winter), pp. 61–66. Available at: http://sloanreview.mit.edu/article/how-to-get-your-messages-retweeted/.

Manovich, L. (2012) 'Trending: the promises and challenges of big social data', in Gold, M. K. (ed.) *Debates in the digital humanities*. Minneapolis: University of Minnesota Press. Available at: http://manovich.net/index.php/projects/trending-the-promises-and-the-challenges-of-big-social-data.

Martin, P., Brians, C. L., Wattenberg, M. P. and Brians, C. L. (1999) 'Negative campaign advertising: demobilizer or mobilizer?', *American Political Science Review*, 93(4), pp. 891–899. doi: 10.2307/2586119.

Mason, R. (2014) *Online political ads test boundaries while electoral law lags behind*, *The Guardian*. Available at: http://www.theguardian.com/politics/2014/apr/07/online-political-ads-electoral-law (Accessed: 23 February 2016).

Mason, R. (2016) *Corbyn: Labour must use social media to fight rightwing press attacks*, *The Guardian*. Available at: http://www.theguardian.com/politics/2016/may/03/corbyn-labour-must-use-social-media-to-fight-rightwing-press-attacks (Accessed: 16 July 2016).

McKelvey, K., DiGrazia, J. and Rojas, F. (2014) 'Twitter publics: how online political communities signaled electoral outcomes in the 2010 US house election', *Information, Communication & Society*, 17(4), pp. 436–450. doi: 10.1080/1369118X.2014.892149.

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), pp. 80–98. doi: 10.1177/107769900007700107.

Merritt, S. (1984) 'Negative political advertising: some empirical findings', *Journal of Advertising*, 13(3), pp. 27–38. doi: 10.1080/00913367.1984.10672899.

Mitchell, A., Holcomb, J. and Weisel, R. (2016) *Election 2016: campaigns as a direct*

source of news, *PewResearchCenter*.

Mizerski, R. W. (1982) 'An attribution explanation of the disproportionate influence of unfavorable information', *Journal of Consumer Research*, 9(December), pp. 301–310. Available at: http://www.jstor.org/stable/2488625.

Momoc, A. (2010) 'Online Negative Campaign in the 2004 Romainian Presidential Elections', *Styles of Communication*, (2), pp. 89–99.

Momoc, A. (2012) 'The presidential candidates on Twitter during the 2009 Romanian elections', *Romanian Journal of Communication and Public Relations*, 14(1), pp. 21–37.

Morchid, M., Dufour, R., Bousquet, P.-M., Linarès, G. and Torres-Moreno, J.-M. (2014) 'Feature selection using principal component analysis for massive retweet detection', *Pattern Recognition Letters*, 49(0), pp. 33–39. doi: 10.1016/j.patrec.2014.05.020.

Moreno, M. a, Goniu, N., Moreno, P. S. and Diekema, D. (2013) 'Ethics of social media research: common concerns and practical considerations', *Cyberpsychology, behavior and social networking*, 16(9), pp. 708–13. doi: 10.1089/cyber.2012.0334.

Morris, D. (2011) *Vote.com*. Kindle. Los Angeles: Renaissance Books.

Muenchen, R. (2014) *The popularity of data analysis software, r4stats.com*. Available at: http://r4stats.com/articles/popularity/ (Accessed: 4 March 2015).

Newman, N. (2010) *#UKelection 2010, mainstream media and the role of the internet: how social and digital media affected the business of politics and journalism*. Available at: http://reutersinstitute.politics.ox.ac.uk/sites/default/files/Social Media and the Election.pdf.

Nielsen, R. K. and Vaccari, C. (2013) 'Do people "like" politicians on Facebook? Not really. Large-scale direct candidate-to-voter online communication as an outlier phenomenon', *International Journal of Communications*, 7, pp. 2333–2356. Available at: http://ijoc.org/index.php/ijoc/article/view/1717.

Nunez, M. (2016) *Former Facebook workers: we routinely suppressed conservative news*, *Gizmodo*. Available at: http://gizmodo.com/former-facebook-workers-we-routinely-suppressed-conser-1775461006 (Accessed: 10 May 2016).

Oboler, A., Welsh, K. and Cruz, L. (2012) 'The danger of big data: social media as computational social science', *First Monday*, 17(7). doi: 10.5210/fm.v17i7.3993.

Pak, A. and Paroubek, P. (2010) 'Twitter as a corpus for sentiment analysis and opinion mining', in Calzolari, N., Choukri, K., Maegaard, B., Mariani, J., Odijk, J., Piperidis, S., Tapias, M., and Rosner, M. (eds) *Proceedings of the Seventh International Conference on Language Resources and Evaluation*. Valletta, Malta: European Language Resources Association, pp. 1320–1326. doi: 10.1371/journal.pone.0026624.

Pariser, E. (2011) *Beware the online 'filter bubbles'*, *TED2011*. Available at: https://www.ted.com/talks/eli_pariser_beware_online_filter_bubbles (Accessed: 26 July 2016).

Park, C. S. (2013) 'Does Twitter motivate involvement in politics? Tweeting, opinion leadership and political engagement', *Computers in Human Behavior*. Elsevier Ltd, 29(4), pp. 1641–1648. doi: 10.1016/j.chb.2013.01.044.

Parmelee, J. H. and Bichard, S. L. (2012) *Politics and the Twitter revolution: how tweets influence the relationship between political leaders and the public*. Plymouth: Lexington Books.

Petrovic, S., Osborne, M. and Lavrenko, V. (2011) 'RT to win! Predicting message propagation in Twitter', in *Proceedings of the Fifth International Conference on Weblogs and Social Media - ICWSM '11*, pp. 586–589. Available at: http://homepages.inf.ed.ac.uk/miles/papers/icwsm11.pdf.

Phelps, J. E., Lewis, R., Mobilio, L., Perry, D. and Raman, N. (2004) 'Viral marketing or electronic word-of-mouth advertising: examining consumer responses and motivations to pass along email', *Journal of Advertising Research*, 44(4), pp. 333–348. doi: 10.1017/S0021849904040371.

Pierce, D. (2016) *What you need to know about Twitter's algorithmic timeline*, *Wired*. Available at: http://www.wired.com/2016/02/what-you-need-to-know-about-twitters-algorithmnic-timeline/ (Accessed: 28 April 2016).

Plotkowiak, T. and Stanoevska-Slabeva, K. (2013) 'German politicians and their Twitter networks in the Bundestag election 2009', *First Monday*, 18(5–6). Available at: http://firstmonday.org/ojs/index.php/fm/article/view/3816.

Porter, A. (2010) *General Election 2010: a battle between hope and fear*, *The Telegraph*. Available at: http://www.telegraph.co.uk/news/election-2010/7561219/General-election-2010-a-battle-between-hope-and-fear.html (Accessed: 6 May 2014).

Prince, R. (2014) *Toddlers mesmerised by surreal world of unboxing videos*, *The Telegraph*. Available at: http://www.telegraph.co.uk/news/worldnews/northamerica/usa/11112511/Toddlers-mesmerised-by-surreal-world-of-unboxing-videos.html (Accessed: 24 January 2014).

Ramesh, R. (2012) *Grant Shapps's Wikipedia page was edited to remove byelection gaffe*, *The Guardian*. Available at: http://www.theguardian.com/politics/2012/sep/11/grant-shapps-wikipedia-edited-gaffe (Accessed: 24 January 2015).

Raynauld, V. and Greenberg, J. (2014) 'Tweet, click, vote: Twitter and the 2010 Ottawa Municipal Election', *Journal of Information Technology & Politics*, 11, pp. 412–434. doi: 10.1080/19331681.2014.935840.

Reige, A. (1998) 'Realism rules OK: scientific paradigms in marketing research', in Deans, B. and Deans, K. (eds) *Australia and New Zealand Marketing Academy Conference (ANZMAC98)*. University of Otago, Dunedin, New Zealand, pp. 1947–1959.

Replies and retweets on Twitter (2010) *Sysomos*. Available at:

https://sysomos.com/inside-twitter/twitter-retweet-stats (Accessed: 20 May 2016).

Richards, S. (2015) *Political leadership is impossible in the age of social media - as Jeremy Corbyn has found*, *The Independent*. Available at: http://www.independent.co.uk/voices/political-leadership-is-impossible-in-the-age-of-social-media-as-jeremy-corbyn-has-found-a6763871.html (Accessed: 17 June 2016).

Ridout, T. N., Shah, D. V, Goldstein, K. M. and Franz, M. M. (2004) 'Evaluating measures of campaign advertising exposure on political learning', *Political Behavior*, 26(3), pp. 201–226. doi: 10.1023/B:POBE.0000043453.96025.a8.

Riffe, D., Lacy, S. and Fico, F. (2014) *Analysing media messages: using quantitative content analysis in research*. 3rd edn. Abingdon: Routledge.

Rogers, S. (2014) *What fuels a tweet's engagement?*, *Twitter Media Blog*. Available at: https://blog.twitter.com/2014/what-fuels-a-tweets-engagement (Accessed: 20 February 2015).

Ronson, J. (2015) *So You've Been Publicly Shamed*. London: Picador.

Rose, K. (2014) *UK social media statistics for 2014*, *Social Media Today*. Available at: http://www.socialmediatoday.com/content/uk-social-media-statistics-2014 (Accessed: 1 March 2015).

Ross, K. and Burger, T. (2014) 'Face to face(book): social media, political campaigning and the unbearable lightness of being there', *Political Science*, 66(1), pp. 46–62. doi: 10.1177/0032318714534106.

Rudat, A., Buder, J. and Hesse, F. W. (2014) 'Audience design in Twitter: retweeting behavior between informational value and followers' interests', *Computers in Human Behavior*. Elsevier Ltd, 35, pp. 132–139. doi: 10.1016/j.chb.2014.03.006.

Ruiter, R., Abraham, C. and Kok, G. (2001) 'Scary warnings and rational precautions: a review of the psychology of fear appeals', *Psychology & Health*, 16(6), pp. 613–630. doi: 10.1080/08870440108405863.

Sæbø, Ø. (2011) 'Understanding Twitter use among parliament representatives: a genre analysis', *Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)*, 6847 LNCS, pp. 1–12. doi: 10.1007/978-3-642-23333-3_1.

Shrimsley, R. (2015) *General election: where fear and negative campaigning pervaded*, *Financial Times*. Available at: https://next.ft.com/content/a944ddf6-f3cc-11e4-99de-00144feab7de (Accessed: 12 April 2016).

Silverman, J. (2015) *'Pics or it didn't happen' – the mantra of the Instagram era*, *The Guardian*. Available at: http://www.theguardian.com/news/2015/feb/26/pics-or-it-didnt-happen-mantra-instagram-era-facebook-twitter (Accessed: 25 February 2015).

Small, T. A. (2012) 'What the hashtag? A content analysis of Canadian politics on

Twitter', in Loader, B. D. and Mercea, D. (eds) *Social media and democracy – innovations in participatory politics*. Abingdon: Routledge.

Smith, M. A., Rainie, L., Himelboim, I. and Shneiderman, B. (2014) *Mapping Twitter topic networks: from polarized crowds to community clusters*. Available at: http://www.pewinternet.org/2014/02/20/mapping-twitter-topic-networks-from-polarized-crowds-to-community-clusters.

So, J., Prestin, A., Lee, L., Wang, Y., Yen, J. and Chou, W.-Y. S. (2015) 'What do people like to "share" about obesity? A content analysis of frequent retweets about obesity on Twitter', *Health Communication*, 236(July 2015), pp. 1–14. doi: 10.1080/10410236.2014.940675.

Sobh, R. and Perry, C. (2006) 'Research design and data analysis in realism research', *European Journal of Marketing*, 40(11/12), pp. 1194–1209. doi: 10.1108/03090560610702777.

Sohn, D. (2009) 'Disentangling the effects of social network density on electronic word-of-mouth (eWOM) intention', *Journal of Computer-Mediated Communication*, 14(2), pp. 352–367. doi: 10.1111/j.1083-6101.2009.01444.x.

Steinberg, T. (2016) *Tom Steinberg profile page*, *Facebook*. Available at: https://www.facebook.com/tom.steinberg.503/posts/10157028566365237?pnref=story (Accessed: 26 July 2016).

Sternthal, B., Dholakia, R. and Leavitt, C. (1978) 'The persuasive effect of source credibility: tests of cognitive response', *Journal of Consumer Research*, 4(4), pp. 252–260. doi: 10.1086/208704.

Stewart, H. and Asthana, A. (2016) *David Cameron says his EU campaign is project fact, not project fear*, *The Guardian*. Available at: http://www.theguardian.com/politics/2016/feb/29/cameron-dismisses-claims-project-fear-eu-remain-negative (Accessed: 18 May 2016).

Stieglitz, S. and Dang-Xuan, L. (2013a) 'Emotions and information diffusion in social media—sentiment of microblogs and sharing behavior', *Journal of Management Information Systems*, 29(4), pp. 217–248. doi: 10.2753/MIS0742-1222290408.

Stieglitz, S. and Dang-Xuan, L. (2013b) 'Social media and political communication: a social media analytics framework', *Social Network Analysis and Mining*, 3(4), pp. 1277–1291. doi: 10.1007/s13278-012-0079-3.

Strong, J. T. and Dubas, K. M. (1993) 'The optimal level of fear-arousal in advertising: an empirical study', *Journal of Current Issues and Research in Advertising*, 15(2), pp. 93–99. doi: 10.1080/10641734.1993.10505006.

Struhl, S. (2015) *Practical text analytics*. London: Kogan Page.

Suh, B., Hong, L., Pirolli, P. and Chi, E. H. (2010) 'Want to be retweeted? Large scale analytics on factors impacting retweet in Twitter network', in *Social Computing (SocialCom), 2010 IEEE Second International Conference on*. Minneapolis, MN: IEEE,

pp. 177–184. doi: 10.1109/SocialCom.2010.33.

Sunstein, C. (2001) *Republic.com*. Woodstock, Oxfordshire: Princeton University Press.

Swirsky, E. S., Hoop, J. G. and Labott, S. (2014) 'Using social media in research: new ethics for a new meme?', *The American Journal of Bioethics*, 14(10), pp. 60–61. doi: 10.1080/15265161.2014.948302.

Thelwall, M., Buckley, K., Paltoglou, G. and Cai, D. (2010) 'Sentiment in short strength detection informal text', *Journal of the American Society for Information Science and Technology*, 61(12), pp. 2544–2558. doi: 10.1002/asi.

Towner, T. L. and Dulio, D. a. (2012) 'New media and political marketing in the United States: 2012 and beyond', *Journal of Political Marketing*. Routledge, 11(1–2), pp. 95–119. doi: 10.1080/15377857.2012.642748.

Tumasjan, A., Sprenger, T. O., Sandner, P. G., Welpe, I. M. and . (2010) 'Election forecasts with Twitter: how 140 characters reflect the political landscape', *Social Science Computer Review*, 29(4), pp. 402–418. doi: 10.1177/0894439310386557.

Turnsek, M. and Jankowski, N. W. (2008) 'Social media and politics: theoretical and methodological considerations in designing a study of political engagement', *Social Science Research Network*. London, pp. 1–26. doi: 10.2139/ssrn.1629098.

Tweetminister (2011) *Tweetminister predicts: findings*. Available at: https://www.scribd.com/doc/31208748/Tweetminster-Predicts-Findings (Accessed: 25 July 2016).

Twitter (2014) *Twitter terms of service*. Available at: https://twitter.com/tos (Accessed: 19 February 2015).

Twitter (no date) *Twitter usage*. Available at: https://about.twitter.com/company (Accessed: 19 January 2015).

Uysal, I. and Croft, W. B. (2011) 'User oriented tweet ranking: a filtering approach to microblogs', in *International Conference on Information and Knowledge Management, Proceedings*. Glasgow, pp. 2261–2264. doi: 10.1145/2063576.2063941.

Vergeer, M., Hermans, L. and Sams, S. (2010) 'The voter is only a tweet away: microblogging during the 2009 European Parliament election campaign in the Netherlands', in *Open SIUC Political Network Conference*. Available at: http://opensiuc.lib.siu.edu/pnconfs_2010/13.

Vergeer, M., Hermans, L. and Sams, S. (2011) 'Online social networks and micro-blogging in political campaigning: the exploration of a new campaign tool and a new campaign style', *Party Politics*, 19(3), pp. 477–501. doi: 10.1177/1354068811407580.

Viner, K. (2016) *How technology disrupted the truth*, *The Guardian*. Available at: https://www.theguardian.com/media/2016/jul/12/how-technology-disrupted-the-truth?INTCMP=sfl (Accessed: 26 July 2016).

Webberley, W. M., Allen, S. M. and Whitaker, R. M. (2016) 'Retweeting beyond

expectation: inferring interestingness in Twitter', *Computer Communications*. Elsevier Ltd., 73, pp. 229–235. doi: http://dx.doi.org/10.1016/j.comcom.2015.07.016.

Weber, R. P. (1990) *Basic content analysis: Sage University paper series on quantitative applications in the social sciences*. 2nd edn. Newbury Park, CA: Sage.

Westerman, D., Spence, P. R. and Van Der Heide, B. (2012) 'A social network as information: the effect of system generated reports of connectedness on credibility on Twitter', *Computers in Human Behavior*. Elsevier Ltd, 28(1), pp. 199–206. doi: 10.1016/j.chb.2011.09.001.

Wicks, R. H. and Souley, B. (2003) 'Going negative: candidate usage of internet web sites during the 2000 presidential campaign', *Journalism & Mass Communication Quarterly*, 80(1), pp. 128–144. doi: 10.1177/107769900308000109.

Williams, S. A., Terras, M. and Warwick, C. (2013) 'What do people study when they study Twitter?', *Journal of Documentation*, 69, pp. 1–74. doi: 10.1108/JD-03-2012-0027.

Wilson, C. and Dunn, A. (2011) 'The Arab Spring: digital media in the Egyptian revolution: descriptive analysis from the Tahrir data set', *International Journal of Communication*, 5, p. 25. Available at: http://ijoc.org/index.php/ijoc/article/view/1180.

Yang, Z., Guo, J., Cai, K., Tang, J., Li, J., Zhang, L. and Su, Z. (2010) 'Understanding retweeting behaviors in social networks', in *CIKM '10 Proceedings of the 19th ACM International Conference on Information and Knowledge Management*. New York, pp. 1633–1636. doi: 10.1145/1871437.1871691.

Ye, Q., Law, R., Gu, B. and Chen, W. (2011) 'The influence of user-generated content on traveler behavior: an empirical investigation on the effects of e-word-of-mouth to hotel online bookings', *Computers in Human Behavior*. Elsevier Ltd, 27(2), pp. 634–639. doi: 10.1016/j.chb.2010.04.014.

Yoon, H. Y. and Park, H. W. (2014) 'Strategies affecting Twitter-based networking pattern of South Korean politicians: social network analysis and exponential random graph model', *Quality and Quantity*, 48, pp. 409–423. doi: 10.1007/s11135-012-9777-1.

Yoon, K. A. K., Pinkleton, B. E. and Ko, W. (2005) 'Effects of negative political advertising on voting intention: an exploration of the roles of involvement and source credibility in the development of voter cynicism', *Journal of Marketing Communications*, 11(2), pp. 95–112. doi: 10.1080/1352726042000315423.

Yuan Huang, W. H. and Zhang, K. (2015) 'Research of micro-blog diffusion effect based on analysis of retweet behavior', in *Proc. 2015 IEEE 14th International Conference on Cognitive Informatics and Cognitive Computing*, pp. 255–261. doi: 10.1007/s12650-016-0347-9.

Zamora Medina, R. and Zurutuza Munoz, C. (2014) 'Campaigning on Twitter: towards

the "personal style" campaign to activate political engagement during the 2011 Spanish General Elections', *Communication & Society*, 27(1), pp. 83–106. Available at: http://www.unav.es/fcom/communication-society/es/articulo.php?art_id=481.

Zhang, Y., Xu, Z. and Yang, Q. (2012) 'Predicting popularity of messages in Twitter using a feature-weighted model', *International Journal of Advanced Intelligence*. Available at: http://www.nlpr.ia.ac.cn/2012papers/gjhy/gh154.pdf.

Zimmer, M. and Proferes, N. J. (2014) 'A topology of Twitter research: disciplines, methods, and ethics', *Aslib Journal of Information Management*, 66, pp. 250–261. doi: 10.1108/AJIM-09-2013-0083.

Zinkhan, G. M. and Hirschheim, R. (1992) 'Truth in marketing theory and research: an alternative perspective', *Journal of Marketing*, 56(2), p. 80. doi: 10.2307/1252044.

Appendices

1. Developments in UK political parties' use of the internet since 1997

The first internet election in the UK was in 1997. The Labour Party launched its website www.labour.org.uk in November 1996 and the BBC launched its Politics 97 website as an experiment to see whether such news could be effectively delivered by way of the internet[38]. The internet was a very peripheral part of this campaign. As Figure 103 shows, party websites were basic and operated in broadcast mode with no voter interaction.

Figure 103 - Snapshot of Labour.org.uk on launch in 1996

By 2001 most political websites had added basic facilities for interaction with voters. For example, the Labour Party's 2001 website (Figure 104) includes several calls to action on the homepage, inviting visitors to enter their postcode to find out what Labour has been doing in their area, and encouraging them to sign up for the party's

---

[38] A basic version of this site can still be seen at http://www.bbc.co.uk/news/special/politics97/

email newsletter. Visitors could also join, register as campaign volunteers and donate to the party from this site. This represents a significant step forward from 1997 but the site still operated soley in broadcast mode with no facilities for visitors to communicate directly with the party. Studies of American political websites of the same time show that they took a very similar approach, focusing heavily on preaching to the converted with the emphasis on reinforcing existing opinions, eliciting donations, encouraging people to become activists and getting out the vote rather than on communicating with undecided voters (Bimber and Davis, 2003).

Figure 104 - Labour Party website at the time of the 2001 General Election
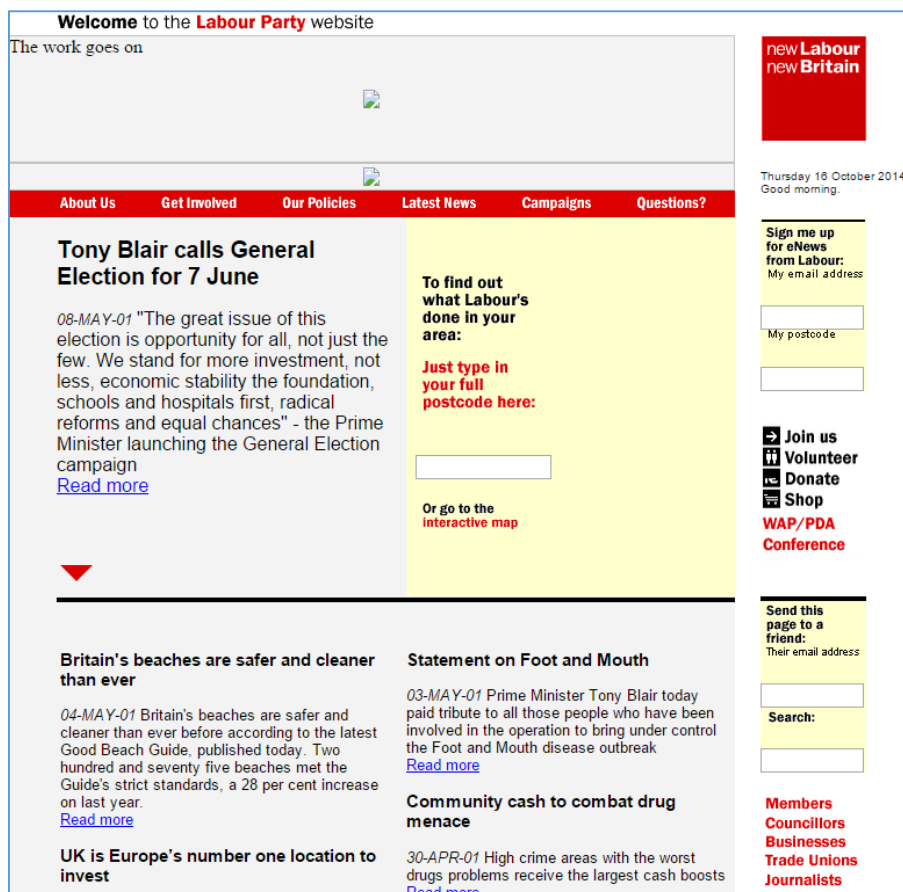
Labour's 2005 site (Figure 105) operated along very similar lines. Although the design is more in keeping with the slick websites we are more used to now, the site was still almost entirely a broadcast platform with no additional opportunities for visitors to interact. Social media was still in its infancy with sites such as Myspace only launching in 2003. Facebook did not launch in the UK until October 2005 and until September

2006 was only accessible to university students. Twitter was not launched until March 2006. Alan Johnson, the Labour MP and former minister, was the first UK politician to be active on Twitter, using it in his campaign to become deputy leader of the Labour Party in 2007 (Jackson and Lilleker, 2011).

Figure 105 - Labour Party website at the time of the 2005 General Election

The election of 2010 was Britain's first true internet election (Newman, 2010), and was the first election in which parties were able to make use of web 2.0 technologies such as social media. By 2010 the power of the internet and social media in particular to enable politicians to connect with voters had been amply demonstrated by Barrack Obama's presidential campaign of 2008. As can be seen from Labour's site (Figure 106), by 2010 political parties had embraced social media. The site features both

Facebook and Twitter icons. However, the focus is still primarily on one-way broadcast communication. The homepage social media icons encourage visitors to follow the party on social media, rather than to talk about it on social media themselves. It is only when one looks at deeper pages in the site, for example the pages explaining different policies, that visitors are invited to share the policy articles with their own social media networks.

Figure 106 - Labour Party website at the time of the 2010 General Election

For the 2015 General Election the Labour Party website was much more sophisticated in its exploitation of the possibilities of social media (Figure 107). There are prominent links from the homepage to Twitter, Facebook, Google+, YouTube and Instagram. Every article on the site concludes with a request for the viewer to share it with their own social networks. A prominent article on the homepage invites viewers to follow

the party's campaign on Twitter using the hashtags #forthemany, #LabourDoorstep and #VoteLabour.

Figure 107 - Labour Party website at the time of the 2015 General Election
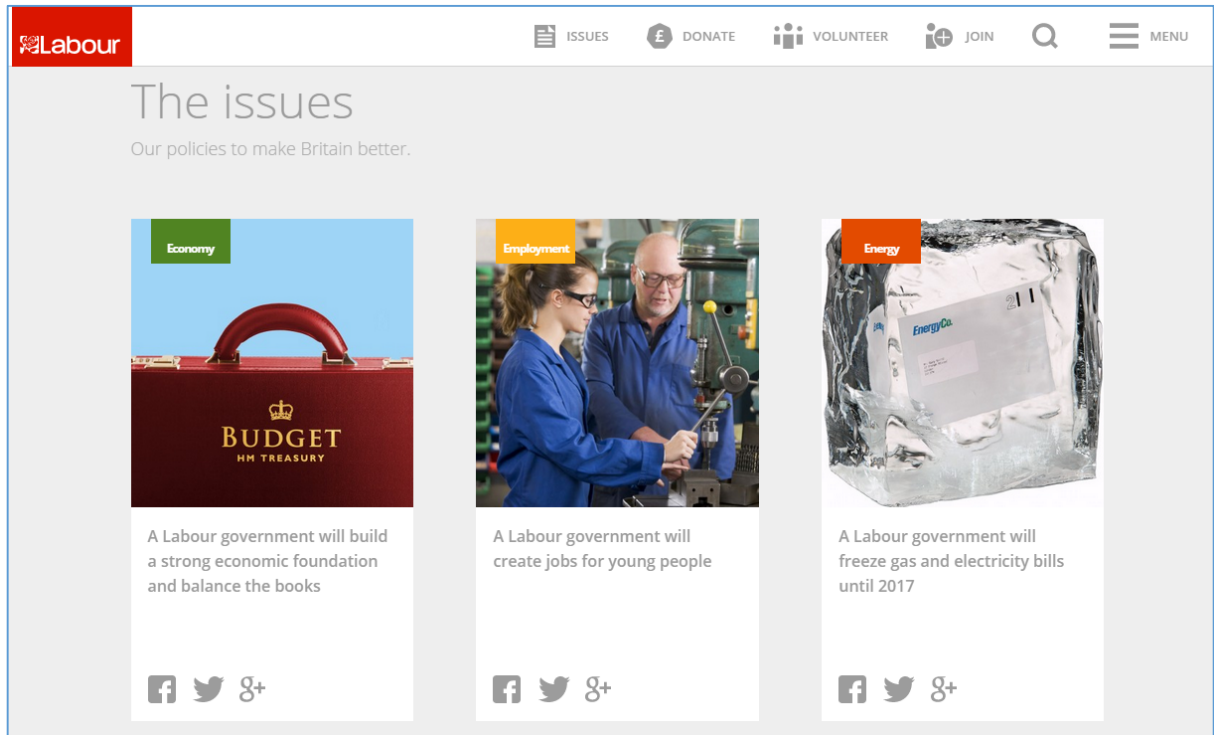
A similar progression can be seen when viewing the websites of the other main political parties in the UK.

2. Full list of variables used in the analysis

Variable	Source	Information contained in the variable
Author variables		
Author	Twitter	The Twitter handle of the author.
Party	Appended	The party to which each politician belongs.
Marginality	Appended	The extent to which each MPs seat was considered to be marginal, calculated according to Finer *et al.*'s 1961 model of marginality as cited in Jackson and Lilleker (2011) and used in their research. This approach categorises a majority of 11% of votes over the next nearest candidate as safe, between 5.1% and 10.9% as near-marginal and anything less than 5% as marginal
Result	Appended	The outcome of the election for each MP (held the seat, lost the seat or stood down at the election so didn't fight the seat).
Gender	Brandwatch / appended	The gender of the tweet's author (where known) – in cases where the gender was unknown then additional research was done and the correct gender assigned to every case.
Kred influence	Brandwatch	Kred is an influence score based on how frequently someone is retweeted, replied, mentioned and followed on Twitter (Kred.com, no date).
Kred outreach	Brandwatch	Kred outreach score is a measure of the extent to which someone engages with others and helps them spread their messages, calculated based on the number of times an author retweets, replies and mentions others (Kred.com, no date).
Twitter followers	Twitter	The number of followers that the author of the tweet had at the time that the data was collected.
Twitter following	Twitter	The number of people whom the tweet author followed at the time that the data was collected.
Following ratio	Calculated	Number of followers divided by the number of people followed, giving an indication of the extent to which politicians listen as well as speak.
Twitter tweets	Twitter	The total number of tweets that the tweet author had sent at the time that the data was collected.
Twitter verified	Twitter	Whether the account has been Twitter verified or not.
Total campaign tweets	Calculated	The total number of tweets sent by the MP during the campaign.

Variable	Source	Information contained in the variable
Campaign retweets generated	Calculated	Total number of retweets generated by all the tweets that an individual sent during the campaign.
Sum of campaign replies	Calculated	Total number of replies generated by all the tweets that an individual sent during the campaign.
Campaign retweets per tweet	Calculated	Average number of retweets generated by the MP's campaign tweets.
Tweets per day	Calculated	Total number of tweets sent during the campaign divided by the number of days of the campaign
Mean retweets	Calculated	Mean number of retweets each MP generates per campaign tweet
Tweet structural data		
Snippet	Twitter	The full text of the tweet
Thread entry type	Twitter	Each tweet falls into one of three categories: a share (retweet), reply (reply to someone else's tweet) or post (an original tweet).
Twitter reply count	Twitter	The number of times that the tweet has been replied to.
Impressions	Brandwatch	The sum of all the followers of authors who tweeted this tweet.
Reach	Brandwatch	The sum of Kred's influence score for all users who have retweeted the tweet.
urlyn	Calculated	Whether or not the tweet contains a link of any kind.
medialinkyn	Calculated	Whether or not the tweet contains a link to a video or picture.
Twitter reply to	Twitter	If a tweet is a reply, the person to whom the reply is directed.
Twitter retweet of	Twitter	In the case of retweets, a link to the original tweet which was retweeted.
Twitter retweets	Twitter	The number of times that the tweet was retweeted.
hashtagyn	Calculated	Whether the tweet contains a hashtag.
hashtagnumber	Calculated	How many hashtags the tweet contains.
mentionnumber	Calculated	How many @mentions the tweet contains.

Variable	Source	Information contained in the variable
Tweet content variables		
Brandwatch sentiment	Brandwatch	Whether the tweet is positive, negative or neutral as determined by Brandwatch.
Manual sentiment	Manual analysis	Whether the tweet is positive, negative or neutral as determined by human coders.
#bbcdebate	Calculated	Whether the tweet contains the hashtag #bbcdebate – one of the hashtags used during the televised leaders' debate.
#leadersdebate	Calculated	Whether the tweet contains the hashtag #leadersdebate – another popular hashtag used during the leaders' debate.
#bbcqt	Calculated	Whether the tweet contains the hashtag #bbcqt – this is the hashtag of the BBC's Question Time programme.
#votelabour	Calculated	Whether the tweet contains the hashtag #votelabour
#voteconservative	Calculated	Whether the tweet contains the hashtag #voteconservative
#GE2105	Calculated	Whether the tweet contains the hashtag #GE2015
#labourdoorstep	Calculated	Whether the tweet contains the hashtag #labourdoorstep
SPSS sentiment	SPSS	Whether tweet is positive or negative as determined by SPSS Text Analytics
SPSS concept variables	SPSS	Series of variables created by SPSS Text Analytics based on the concepts that it finds in the data (not listed here)
SPSS category variables	SPSS	Series of variables created by SPSS Text Analytics based on the categories that it finds in the data (not listed here)

3. An introduction to Twitter

This appendix provides an overview of how Twitter works in order to set the scene for the research presented in this thesis. The aim of the appendix is to ensure that the reader is familiar with core Twitter functionality and those aspects of Twitter's operation which are particularly relevant to this research project.

Twitter, founded in 2006, enables asynchronous communication between users who have 140 characters, ostensibly to answer the question 'What's happening?'. It can be accessed via web browser or mobile phone and so is simple to use, particularly for anyone familiar with texting. As of March 2016 Twitter reports that it has 320 million active users sending 500 million tweets per day, of which 80% are sent via mobile (Twitter, no date). A user's Twitter messages (tweets) are available to view on their personal profile page and in the Twitter feeds of the people who follow them. Unlike Facebook, Twitter is an open and outward-facing medium. The default setting is that tweets are publically viewable and anyone can follow anyone else.

You do not need to know someone in order to follow them on Twitter. It is possible to have a private Twitter account (in which case followers have to be approved by the account holder) but only a small number of users opt for this – 11.8% in 2012 (Bosker and Grandoni, 2012) – so the vast majority of Twitter accounts are publically accessible. Anyone can respond to a tweet – you do not have to be connected to an individual in order to direct a tweet at them. Because of this open structure, individuals can follow celebrities, politicians, journalists and other well-known figures and interact with them directly in a way that would not be possible through other media.

There are three ways in which a tweet can be categorised – as a post, a reply or a retweet. A post is an original tweet, authored by the person who sent it. It could contain elements such as a mention of another person or a link to material provided by someone else. Figure 108 shows an example of a post. The author of the tweet is Alistair Carmichael. He mentions @nick_clegg in this tweet but the tweet is not a reply

to something that @nick_clegg has written and so is categorised as a post rather than a reply.

Figure 108 - Example of a post
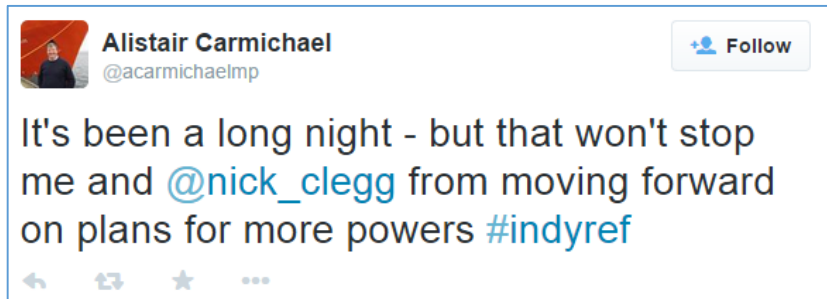
A reply is a tweet which one individual sends in reply to another individual. Twitter only counts a tweet as a reply if the recipient's name (i.e. the person to whom the author is replying) appears at the start of the tweet. A tweet which mentions another person, but where their name does not appear at the start of the tweet (as in the example from Alistair Carmichael, above), does not count as a reply. The tweet shown in Figure 109 is an example of a reply. Angela Constance has replied to a tweet from @Ianrj. His handle appears at the start of the tweet, indicating that the reply is directed at him in response to something that he has tweeted.

Figure 109 - Example of a reply

The third kind of tweet is the retweet. There are two types of retweet – auto retweets and modified retweets. Twitter has a built-in retweet feature so a user can simply click a button and immediately retweet any tweet they like. This is known as an auto-retweet. For example, Figure 110 shows that Alex Fergusson has retweeted a tweet originally sent by @rogerlwhite. He has not added any additional comments of his own to the original tweet, he has simply pushed the 'retweet' button so this is an auto retweet.

Figure 110 - Example of a retweet
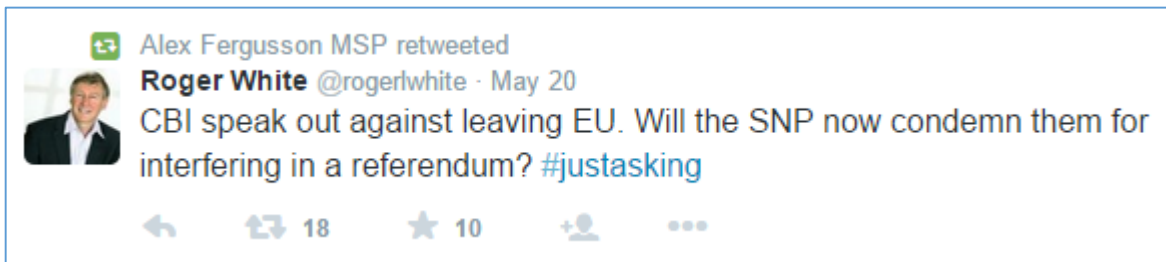
It is also possible for a user to modify a tweet before they retweet it, perhaps to add some comments of their own. This is known as a modified retweet and Twitter convention is generally that these tweets begin 'MT' to indicate that they are modified retweets (see the example in Figure 111 in which Mike Crockart has added his own comments to a tweet originally sent by @VSOUK before passing it on to his own followers as a modified retweet). Hardly any of the MPs' tweets in the dataset examined in this research used the MT convention – only 28 out of 154,565 tweets were MTs (compared to 91,644 retweets).

Figure 111 - Example of a modified retweet
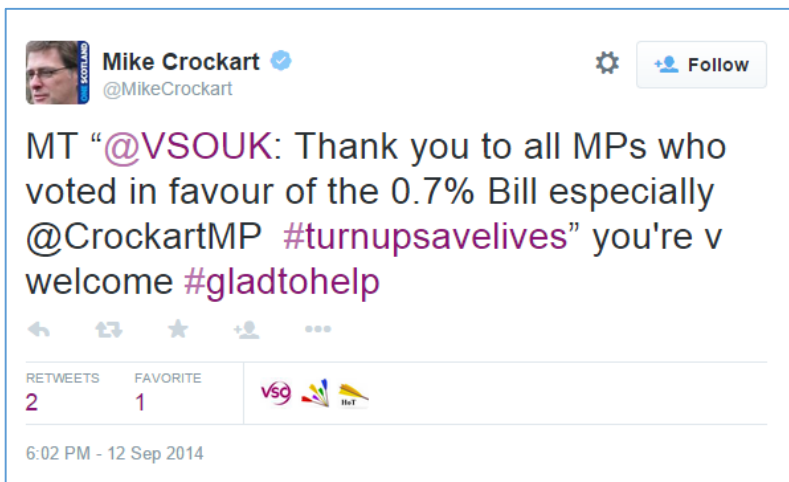
Users are notified of any tweets which are directed specifically to them (known as an @reply) or which mention them (known as an @mention). An @mention is any tweet which includes the name of another Twitter user anywhere within the body of the tweet. For example, the tweet below (Figure 112) from David Cameron mentions the

Twitter names of two potential Conservative candidates. Both @anna_firth and @KellyTolhurst would then be notified that they had been mentioned in another user's tweet. Other users will see any @mentions posted by someone they follow – they do not have to be following the people who are mentioned in the tweet.

Figure 112 - Example of a tweet including mentions
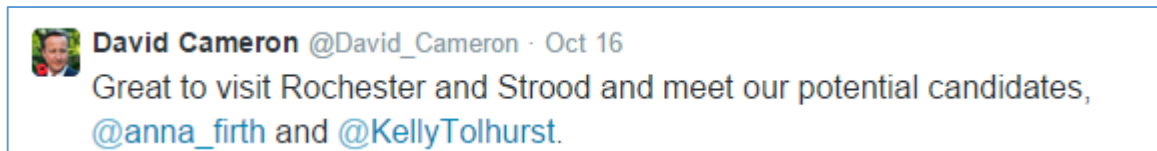
An @reply is when a user replies to a particular tweet sent by someone else. In this instance the name of the user to whom they are replying will appear at the start of the tweet. Figure 113 below shows an example of an @reply.

Figure 113 - Example of an @reply
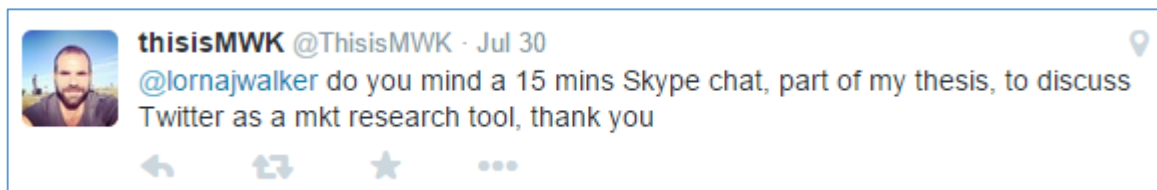
There can be some overlap between these categories. For example, posts, retweets and @replies can also contain a @mentions. However, the same tweet cannot simultaneously be both a retweet and an @reply. Likewise, a post cannot also be a retweet or an @reply. The three most significant categories of tweet for the purposes of this research – posts, retweets and @replies – are mutually exclusive.

An important feature of Twitter is its use of hashtags. Hashtags function as a method of alerting people to the fact that a tweet is about a particular topic. Through the use of hashtags, users can communicate with wider communities of people beyond simply their own followers. It is common practice nowadays for television programmes to alert viewers to the hashtag they should use if they wish to comment on the programme. For example, Question Time, the BBC's current affairs discussion programme, has its own Twitter account (@bbcquestiontime) and encourages viewers to interact with the programme using the #bbcqt hashtag.

People can tweet under the same hashtag without being further connected to each other as followers on Twitter, enabling them to communicate beyond their direct network of connections as part of a looser network of ad hoc connections which might spring up around a particular topic (Bruns and Burgess, 2012). For example, in the tweet shown in Figure 114 Sadiq Khan uses hashtags #forthemany and #Lab14 to signal that his tweet is part of a wider conversation about Labour's campaign. In the run up to the 2015 election, the Labour Party encouraged people to tweet using the hashtags #forthemany, #labourdoorstep and #votelabour in order to be part of a wider campaign-based conversation on Twitter.

Figure 114 - Example of how hashtags can be used in tweets
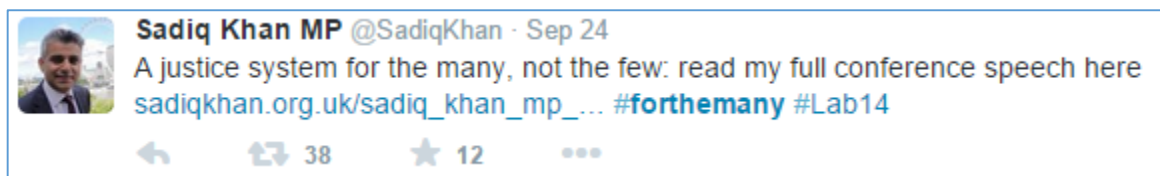
It is important to bear in mind that not everyone who tweets about a particular topic will use a hashtag so tweets which contain particular hashtags do not represent the total Twitter discussion of that topic. Estimates of the number of tweets containing hashtags vary. Liu *et al.* (2014) found that somewhere around 15% of all tweets contain hashtags, with an average of 1.6 hashtags per tweet however more recent research from Buffer (a commercial social media management tool) suggests that around half of tweets contain at least one hashtag (Lee, 2015). In the dataset used for this thesis 35% of the tweets sent by MPs contain at least one hashtag. Thus hashtags represent only a sample of the discussion on a particular topic, and the sample is biased towards more sophisticated Twitter users who better understand how to use its conventions (Bruns and Burgess, 2012). Additionally, when someone replies to a hashtagged tweet they may not include the original hashtag in their reply, meaning that there may be an ongoing conversation about a topic which a simple search of hashtags will not reveal.

One must also be wary of assuming that the use of a particular hashtag indicates agreement with or endorsement of the dominant view of typical users of that

hashtag. Using political hashtags from 'the other side' can be a way of trying to expose ideologically opposed users to one's own point of view. For example, American conservatives use the hashtag #p2 ('progressives 2.0') as a way of exposing liberals to views from the other side of the political spectrum (Conover *et al.*, 2011). Such 'political spamming', even to the extent where dedicated Twitter accounts are set up to spam particular political hashtags in the run up to an election, has also been noted as a phenomenon in Canadian municipal elections (Raynauld and Greenberg, 2014) so use of a particular hashtag definitely cannot be assumed to mean agreement with its core sentiment. Additionally, people often use hashtags ironically rather than as a way of signalling agreement and this can be a particular issue when conducting computer-based sentiment analysis – a tweet including #awesome cannot necessarily be assumed to be a positive tweet, as shown in Figure 115.

Figure 115 - Example of ironic use of hashtags
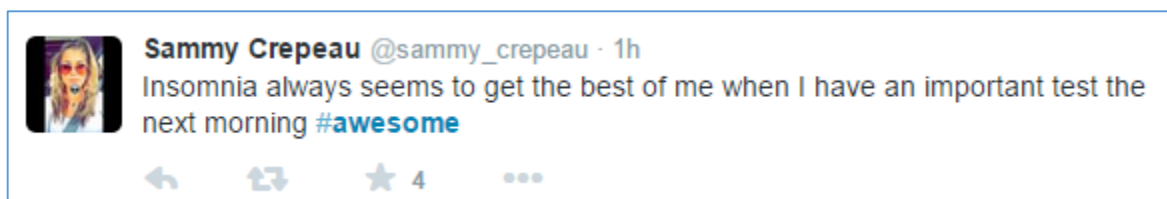
Tweets are often used as a way of passing on content to others that one feels may be of interest to them. This content takes three main forms: videos, pictures and links to other websites. Sadiq Khan's tweet in Figure 114 showed an example of including a link in a tweet – he is using the tweet as a way of directing people to the full text of his conference speech, available on his own website sadiqkhan.org.uk. Ellen DeGeneres' Oscar selfie (Figure 116) is an example of a picture being included in a tweet. Pictures are particularly easy to include in tweets as all phone-based Twitter applications offer the ability to tweet pictures at the click of a button. In 2013 Twitter introduced Vine, enabling users to share short (6 seconds or less) video clips in their tweets and in March 2015 it added additional live streaming video functionality through its

Periscope service, which enables users to stream live videos directly from their mobile phones onto Twitter.

Figure 116 - The most retweeted tweet in history

Twitter facilitates wider communication across its network through the medium of the retweet. Users who deem a particular tweet to be of likely interest to their own followers can retweet it, essentially a method of forwarding it on. Estimates suggest that around 28% of tweets are retweets (Liu, Kliman-Silver and Mislove, 2014), although as with hashtags, how many tweets are retweets depends a lot on who is asking and retweet rates vary widely from user to user, with some users relying almost exclusively on retweets as a way of generating content whilst other users very rarely retweet anything. A retweet acts as a form of endorsement of the content of the original tweet, which can be retweeted with or without additional commentary from the retweeter. Popular tweets can be retweeted many times. The record is currently held by Ellen DeGeneres whose 2014 Oscar night selfie (Figure 116) was retweeted more than three million times.

Another way in which users can indicate some level of engagement with a tweet is to favourite it. When a user sees a tweet they like they can click the 'favourite' button

(this is the star in the bottom right hand corner of the tweet in Figure 116. The tweet is then saved to their favourites list and its sender is notified that their tweet has been favourited[39] (see an example of such a notification in Figure 117). Favouriting can be used as a way of keeping an archive of tweets that one may wish to refer to again at some point. It can also be used as a way of indicating to the original recipient that the tweet was appreciated. Either way, favouriting a tweet implies some kind of engagement with the content although perhaps at a lesser level than implied by a retweet.

Figure 117 - Example of a favourited tweet
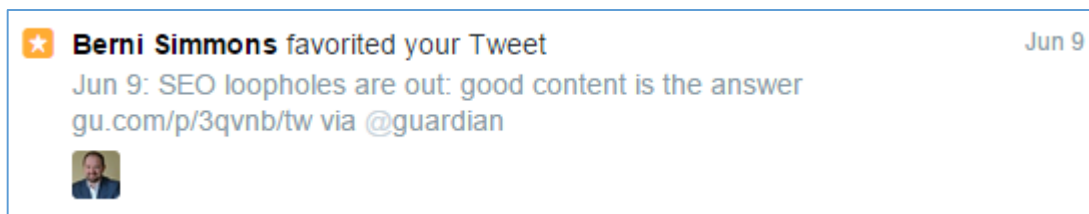
The retweet functionality in Twitter has developed considerably since Twitter was launched. Between 2006 and 2009 there was no official retweet ability provided by Twitter and so conventions around retweeting were developed informally amongst users. This meant that the first researchers to consider retweets had a much harder job of identifying them than is the case now. For example, boyd *et al*. (2010) had to use content analysis to identify retweets, assuming that any tweets including 'RT', 'retweet' or 'via' were retweets, whereas now retweets are flagged as such in Twitter's database and so identifying them is a simple matter. As methods of identifying retweets have changed over the years, it is important to remember that older retweet research will not have identified retweets in the same way as more recent research and so comparing the two will not be comparing like with like.

This research, in line with the approach most commonly used in the literature (e.g. Petrovic *et al.,* 2011; Uysal and Croft, 2011; Kupavskii *et al.,* 2012) counts retweets as

---

[39] Note that in November 2015 Twitter changed 'favourites' to 'likes'. However, since favourites was the term in use at the time of data collection, that is the term that is used throughout this thesis.

those tweets explicitly identified as retweets in the Twitter meta data. The Twitter native app interface offers two methods of retweeting – a simple retweet and a retweet with quote (as shown in Figure 118).

Figure 118 - Retweeting options in the native Twitter app

Twitter only logs a retweet when someone pushes the 'Retweet' button. If someone pushes the 'Quote Tweet' button this is not counted as a retweet. Once the person has made their edits to the original tweet and sent it, Twitter counts this as a new tweet (Compston, 2014). Similarly, if someone cut and pasted the content of a tweet into the Twitter edit window and then tweeted it, this would also be counted as a new tweet rather than as a retweet. One could argue that quoting a tweet whilst adding some comments of one's own indicates a higher level of engagement with a topic than simply retweeting without comment however there is no reliable way of identifying such tweets as shown by the fact that different researchers do it in different ways. For example, Suh *et al*. (2010) identify retweets using an automated

regex search that picks up any tweets that include RT, RT@, rt, retweet, retweeting or via and analyse these alongside what they call 'feature retweets' – those which simply use the retweet button. Achananuparp *et al*. (2012) make a distinction between 'strong retweets' (those which include RT@ or via@) and weak retweets (essentially a form of modified retweet). So *et al*. (2015) build an algorithm which considers how similar tweets are to other tweets and assumes that any tweets with more than 65% similarity to another tweets are in fact retweets. In essence, once one moves away from simply using Twitter's own retweet flag to identify retweets there is no consensus at all on how retweets should be identified, with every group of researchers appearing to favour their own distinct method. Thus, using the Twitter database's count of retweets is the clearest and most unambiguous way of identifying retweets. It is also the method most commonly used in the literature and provides the greatest level of consistency between different pieces of research. For these reasons it is the method used in this research.

It is also important to understand that the retweet count of a tweet is always linked to the original tweet, not to any subsequent retweets of that retweet. For example, if user A sends a tweet and user B retweets it that counts as one retweet for the original tweet. If user C sees the retweet in user B's timeline and hits the retweet button, this then becomes two retweets for the original tweet sent by User A. User B's retweet cannot itself be retweeted. Thus all the retweets in the data collected for this research themselves have a retweet value of zero. This means that if a tweet is retweeted many times it is not possible to identify which of those retweets have been further retweeted as all the retweets will simply be credited back to the original tweet. Therefore, the politicians' retweets (the retweets they themselves sent, passing on other people's tweets) have been excluded from the analysis as they each have a retweet score of zero and hence leaving them in suggests that politicians are generating retweets at a much lower rate than is actually the case. The analysis presented in this research focuses on predicting the extent to which the politicians' original tweets get retweeted.

One final point to note regarding retweets is that the number of retweets for a particular tweet in the dataset used here represents the number of times that tweet had been retweeted *at the point at which the data was collected* so any subsequent retweets will not be included. It may be that the actual retweet numbers are therefore higher than the data presented here suggests, although as most tweets tend to get retweeted fairly quickly and are then forgotten it is unlikely that there would have been many additional retweets after the data was collected, particularly since the data was collected six months after the election which should be enough time for all the retweeting that was going to happen to be done.

4. Summary of key papers relating to politicians on Twitter

Authors	Country	Date	Election?	Sample	Method	Unit of analysis
Adams and McCorkindale, 2013	US	2012	Presidential election	605 tweets of presidential candidates	Basic computerised content analysis	Tweets
Golbeck *et al.*, 2010	US	2009	Outside election time	4,959 – 200 most recent tweets from Members of Congress	Manual coding of tweets – each tweet into as many categories as relevant	Tweets
Graham *et al.*, 2013	UK	2010	General Election	26,282 tweets from 416 candidates of 3 main parties in two weeks before election	Manual coding – not clear how categories determined.	Tweets
Hemphill *et al.*, 2013	US	2012	Outside of election time	31,164 tweets from Members of Congress	Manual coding of 791 tweets to develop coding scheme, computer coding of 30,373 tweets	Tweets
Jackson and Lilleker, 2011	UK	2009	Outside election time	Tweets from 51 MPs on Twitter – sample size not identified	Combination of manual coding, computer content analysis and regression	Tweets
Momoc, 2012	Romania	2009	Presidential election	Tweets sent by 7 out of 12 presidential candidates	Content analysis of tweets – exact method used unclear	Tweets
Vergeer *et al.*, 2010	Holland	2009	European Elections	Tweets sent by 36 Dutch candidates between 1 Feb – 4 June 2009. Sample size unclear.	Coding method unclear. Bivariate analysis.	Politicians
Vergeer *et al.*, 2011	Holland	2009	European Elections	Tweets sent by 36 Dutch candidates between 1 Feb – 13 Oct 2009. Sample size unclear.	Basic stats on calculated variables – no content analysis	Politicians
Plotkowiak and Stanoevska-Slabeva, 2013	Germany	2009	Bundestag elections	240,000+ tweets sent by 599 candidates	Sentiment analysis (combination of manual and computerised) and network analysis	Politicians

Authors	Country	Date	Election?	Sample	Method	Unit of analysis
Enli and Skogerbø, 2013	Norway	2011	Local elections	754 tweets sent by non-random sample of candidates from main parties in 5 areas.	Limited content analysis to determine if tweets were broadcast or not. Manual coding.	Tweets
Larsson and Kalsnes, 2014	Norway and Sweden	2013	General Election in one country compared with non-election time in the other	193 Norwegian politicians and 377 Swedish politicians	Politicians are the unit of analysis rather than tweets.	Politicians
Yoon and Park, 2014	South Korea	2010	Not during election time	189 South Korean politicians	Network analysis and statistical analysis. No content analysis of tweets.	Politicians
Raynauld and Greenberg, 2014	Canada	2010	Local elections	9,409 Tweets using #ottvote and accounts of candidates for range of local positions	Sentiment, topical and quantitative text analysis of tweets using Crimson Hexagon. Network analysis using Gelphi.	Tweets
Grant et al., 2010	Australia	2009 - 2010	Not an election period	118,122 tweets collected outside an election period	Compared tweets of a random sample of 477 Australian Twitter users with 152 Australian politicians.	Politicians
Lawless, 2012	US	2009	Outside an election period	7,668 tweets from 186 members of congress who tweeted	Tweets and Facebook updates of members of US Congress	Tweets and politicians
Adi et al., 2013	UK	2012	Not in election period	4,363 tweets sent by 21 peers sitting on the Labour frontbench	Content analysis (manual and computer)	Tweets
Zamora Medina and Zurutuza Munoz, 2014	Spain	2011	General election 16 day campaign	2,274 tweets sent by 2 candidates for PM during the campaign	Manual content analysis	Tweets

5. Coding schema for content analysis

Code	Tweet content	Example
Achievement	Highlighting some personal achievement or achievement of the party. This includes statement about things that the party has done since the last election or is responsible for and statements about the MPs' own record for their constituency or things that the party has done that have influenced the constituency. The party does not have to be explicitly named – the reference to how things have improved or are better because of the party is the key thing.	So pleased to hear a passport is on its way to a constituents daughter – much of yesterday spent sorting the problem – huge thanks to HMPO (@carolinenokes) There are now 2million more people in work than in 2010. Largest increase in employment has been in the North West (@Andrew4Pendle)
Call to action	Requesting recipients to do something other than vote (e.g. register to vote, sign petition, volunteer, donate money, attend event)	Come and help me win in Cardiff Central! Sign up here: cardiffld.org.uk/volunteer (@JennyWillott)
Info	Providing information (or a link to information elsewhere)	8 questions you should ask about the internet of the things from the @guardian (and me☺ theguardian.com/media-network/… (@ChiOnwurah)
Call to vote	Calling for people to vote	Riverside residents. Have you voted yet? Me+Michelle Corrigan, your Lab candidate – dedicated cllr – working for you #voteLabour #winNW15 (@LouiseEllman)
Position taking	Expressing an opinion on some aspect of own or opposition's policy	I believe Birmingham and Britain as a whole only succeed when working people succeed. #LabourManifesto sets out how labour.org.uk/page/-/Britain… (@RichardBurdenMP)
Meeting people	Referring to individuals or groups of people that has met during campaign	Lovely to meet Sarah whos 92 and been a member of the Party since 1947. #VoteLabour pic.twitter.com/LFGRJFmtNU (@GordonMarsden)
Thanks	Thanking people for help, volunteering, support etc	Our unsung heros @IslingtonSouth: the stuffing team, the data team, the boards men. Thankyou! Pic.twitter.com/Gqp7mblZER (@EmilyThornberry)

Code	Tweet content	Example
Campaign	Mention of some aspect of the day's campaigning – the focus of this category is on door-to-door campaigning and on each MPs own personal campaign rather than the wider party campaign.	We spent this afternoon in Hawes Side for the #12Wards12Days challenge, discussing the cost of living with voters. Pic.twitter.com/3jVNPLdOPv (@GordonMarsden)
Event	Mention of an event that has attended or is planning to attend – mention of hustings events come into this category rather than into the campaign category.	On my way now to Shildon hustings #GE2015 (@HelenGoodmanMP)
Personal	Non-political tweet related to some other aspect of life (e.g. sport, music, humour, restaurant visit)	Congrats to @PtstudioWigan and all at #ptstudio on being crowned Gym of the Year. Still don't like #burpees tho! (@Y_FovargueMP)
Local	Some mention of local constituency	Spring has sprung gloriously in Wycombe pic.twitter.com/fv3VnP6DRW (@SteveBakerHW)
Attack	A negative tweet attacking some aspect of opposing party's policy or individual politicians from the other side	Labour campaign on NHS in Wales takes voters for idiots. Its Labour cuts that got us in this mess. (@GutoBebb)
Fear appeal	A fear appeal is a particular form of negative tweet in which the MP explicitly warns of negative consequences if the other side were to win.	Our economy is recovering. Don't let Labour wreck it! Vote Conservative today! (@chhcalling)
Charity	Mention of some charity that the MP supports or has worked with	Such a fantastic charity Im currently visiting with @MarkHunter – training people to give first aid to children. Pic.twitter.com/5Tzp5FeqTx (@nick_clegg)
Support for others	Expression of support for someone else from own side. Good luck messages. Statements about being out campaigning on someone else's behalf.	Good luck today @JustineGreening. #Putney -#VoteConservative (@S_Hammond)
Support for self	Passing on messages of support which MPs have received for themselves. Mentioning that a senior politician has been campaigning in a particular MP's constituency.	Glad to see my GP Dr Weir will be voting @UKLabour #labourdoorstep (@KarlTurnerMP)

Code	Tweet content	Example
Media response	Tweets in response to something that the MP has seen on the television or heard on the radio.	Boris show his true colours in last few minutes on #marr (@Debbie_abrahams)
Weather	Mentions of the weather	Wow. Reckon its hottest day of year already! Kids dropped off at mega maker holiday club with sun lotion on. Excitement levels very high. (@sbrine)
Other	Tweets which cannot be allocated to another category e.g. links with no context given, tweets sent in error, tweets which appear to be part of a larger conversation	The final countdown has begun in the L (@heidi_mp)

6. CHAID decision tree rules

A full set of decision tree rules for model one are provided in this appendix for illustrative purposes. Decision tree rules for all the other decision trees are available online here or by cutting and pasting or copying the address below into a browser.

https://www.dropbox.com/s/3ua3b9frt2is4ky/full list of decision tree model rules.pdf?dl=0

Model one decision tree rules

Rules for FALSE - contains 4 rule(s)

      Rule 1 for FALSE

            if medialinkyn in ["FALSE"]

            and urlyn in ["FALSE"]

            and mentionyn in ["TRUE"]

            and hashtagyn in ["FALSE"]

            then FALSE

      Rule 2 for FALSE

            if medialinkyn in ["FALSE"]

            and urlyn in ["TRUE"]

            and hashtagyn in ["FALSE"]

            then FALSE

      Rule 3 for FALSE

            if mentionnumber <= 0

            and medialinkyn in ["TRUE"]

            and hashtagyn in ["FALSE"]

            then FALSE

      Rule 4 for FALSE

            if medialinkyn in ["FALSE"]

            and mentionyn in ["FALSE"]

            and urlyn in ["TRUE"]

and hashtagyn in ["TRUE"]

then FALSE

Rules for TRUE - contains 7 rule(s)

Rule 1 for TRUE

if medialinkyn in ["FALSE"]

and urlyn in ["FALSE"]

and mentionyn in ["FALSE"]

and hashtagyn in ["FALSE"]

then TRUE

Rule 2 for TRUE

if mentionnumber > 0

and mentionnumber <= 1

and medialinkyn in ["TRUE"]

and hashtagyn in ["FALSE"]

then TRUE

Rule 3 for TRUE

if mentionnumber > 1

and medialinkyn in ["TRUE"]

and hashtagyn in ["FALSE"]

then TRUE

Rule 4 for TRUE

if mentionyn in ["FALSE"]

and urlyn in ["FALSE"]

and hashtagyn in ["TRUE"]

then TRUE

Rule 5 for TRUE

if medialinkyn in ["TRUE"]

and mentionyn in ["FALSE"]

and urlyn in ["TRUE"]

and hashtagyn in ["TRUE"]

then TRUE

Rule 6 for TRUE

if medialinkyn in ["FALSE"]

and mentionyn in ["TRUE"]

and hashtagyn in ["TRUE"]

then TRUE

Rule 7 for TRUE

if medialinkyn in ["TRUE"]

and mentionyn in ["TRUE"]

and hashtagyn in ["TRUE"]

then TRUE

Default: TRUE

7. Significance of manual content variables

Table 52 - Manual content variables positively associated with retweeting

		Retweeted?		
		No	**Yes**	
Fear appeal	Count	4	35	Chi-square 24.729
	Residuals	-3.5	3.5	df 1
	Percentage	10.3%	89.7%	Sig (two-tailed) < .001
Support for others	Count	15	51	Chi-square 19.900
	Residuals	-3.1	3.1	df 1
	Percentage	22.7%	77.3%	Sig (two-tailed) <.001
Attack	Count	55	153	Chi-square 53.046
	Residuals	-4.7	4.7	df 1
	Percentage	26.4%	73.6%	Sig (two-tailed) <.001
Media response	Count	29	66	Chi-square 14.742
	Residuals	-2.6	2.6	df 1
	Percentage	30.5%	69.5%	Sig (two-tailed) <.001
Position taking	Count	50	79	Chi-square 6.566
	Residuals	-1.7	1.7	df 1
	Percentage	38.8%	61.2%	Sig (two-tailed) .010
Campaign trail	Count	86	133	Chi-square 11.023
	Residuals	-2.1	2.1	df 1
	Percentage	39.3%	60.7%	Sig (two-tailed) .001

Table 53 - Manual content variables negatively associated with retweeting

		Retweeted?		
		No	**Yes**	
Event	Count	40	25	Chi-square 4.034
	Residuals	1.4	-1.4	df 1
	Percentage	61.5%	38.5%	Sig (two-tailed) .045
Personal	Count	103	46	Chi-square 26.389
	Residuals	3.4	-3.4	df 1
	Percentage	69.1%	30.9%	Sig (two-tailed) <.001

Table 54 - Manual content variables not significantly related to retweeting

		Retweeted?		
		No	**Yes**	
Achievement	Count	43	47	Chi-square 0.105
	Residuals	-.2	.2	df 1
	Percentage	47.8%	52.2%	Sig (two-tailed) .746
Call to action (not voting)	Count	22	19	Chi-square 0.305
	Residuals	.4	-.4	df 1
	Percentage	53.7%	46.3%	Sig (two-tailed) .581
Call to vote	Count	21	27	Chi-square .643
	Residuals	-.6	.6	df 1
	Percentage	43.8%	56.3%	Sig (two-tailed) .422
Charity	Count	9	11	Chi-square .159
	Residuals	-.3	.3	df 1
	Percentage	45%	55%	Sig (two-tailed) .690
Information provision	Count	30	36	Chi-square .440
	Residuals	-.5	.5	df 1
	Percentage	45.5%	54.5%	Sig (two-tailed) .507
Local	Count	220	239	Chi-square .658
	Residuals	-.5	.4	df 1
	Percentage	47.9%	52.1%	Sig (two-tailed) .417
Meeting people	Count	23	22	Chi-square .053
	Residuals	.2	-.2	df 1
	Percentage	51.1%	48.1%	Sig (two-tailed) .817
Support for self	Count	40	37	Chi-square .210
	Residuals	.3	-.3	df 1
	Percentage	51.9%	48.1%	Sig (two-tailed) .647
Thanks	Count	77	64	Chi-square 1.718
	Residuals	.9	-.9	df 1
	Percentage	54.6%	45.4%	Sig (two-tailed) .190
Weather	Count	22	22	Chi-square .006
	Residuals	.1	-.1	df 1
	Percentage	50%	50%	Sig (two-tailed) .938
Other	Count	15	15	Chi-square .004
	Residuals	.0	.0	df 1
	Percentage	50%	50%	Sig (two-tailed) .949