King's Research Portal

Document Version
Peer reviewed version

[Link to publication record in King's Research Portal](#)

Fraud, Plot or Collective Delusion? Social Media and Perceptions of Electoral Misconduct in the 2014 Scottish Independence Referendum*

Sarah Birch (King's College London) and Fatma ElSafoury (University of Glasgow)

ABSTRACT: Online discussions of electoral fraud are becoming an increasingly important aspect of the electoral landscape in many contexts, as cyberspace is one of the few places where concerns about electoral conduct can be aired openly and freely. But it is often difficult to assess what this online activity tells us about actual electoral processes. This paper analyzes a surge of tweets about electoral fraud at the time of the Scottish independence referendum of 2014 in order to ascertain whether this online activity reflected: (a) actual offline fraud observed by the social media users, (b) a concerted effort to undermine confidence in electoral administration, or (c) a collective delusion. Data mining and machine learning techniques are deployed in this analysis, which comes out strongly in favour of the collective delusion hypothesis.

KEYWORDS: electoral fraud, Scottish independence referendum, Twitter, text mining, machine learning

Fraud, Plot or Collective Delusion? Social Media and Perceptions of Electoral Misconduct in the 2014 Scottish Independence Referendum

Social media has in recent years come to play a major role in elections. Communication taking place on Twitter, Facebook, YouTube, online forums and other social platforms has become part of election debate, as was evident notably in the US presidential election of 2016 (Allcott and Gentzkow, 2017; Le et al., 2017). Over the course of the past decade we have also witnessed a dramatic rise in the use of social media as a forum for discussing the mechanics of electoral administration, particularly in 'contentious elections' – elections where there is popular mobilization around the conduct of the poll (Diamond, 2010; Howard, 2010; Ikufor, 2010; Norris, Frank and Martínez i Coma, 2015).[1] Yet we have limited understanding of the relationship between the online discussion and the offline reality of electoral integrity. Is vigorous social media debate about electoral misconduct likely to be the 'smoke' that points to the 'fire' of actual irregularities in the electoral process? Or ought such debates to be treated as forms of political communication that have no necessary referent in the real world but instead reflect attitudes and dispositions? Answers to these questions will inform the study of social media and 'fake news'; they will also help students of electoral fraud to assess the relevance of online discussions of fraud.

The Scottish independence referendum of 2014 is an interesting case where electoral administration - which had not in recent times been a subject of extensive popular debate or mobilization - was suddenly thrust into the spotlight largely via social media. Online

---

[1]   This type of social-media mobilization and online activism received considerable international attention at the time of the so-called 'Twitter revolution' in Moldova and the 'Green Revolution' in Iran - both in 2009 - as well as various electoral events that formed part of the Arab Spring (Bennett and Segerberg, 2013; Diamond, 2010; Hänska-Ahy and Shapour, 2013; Margetts et al, 2015).

discussion during the referendum campaign reached unprecedented levels (Tierny, 2015), especially among 'Yes' supporters (Holby, 2014; Shephard et al., 2014), and several studies have analyzed the way in which social media discourse shaped debate about emergent news themes (Brigadir, 2015; Pedersen et al. 201d4; Quinlan, Shephard and Paterson, 2015; Shephard *et al.* 2014). Yet there has been no systematic consideration of the intense online debate about the conduct of the referendum, and this is a lacuna this paper seeks to fill, with a view to shedding light on the relationship between social media-based discussions of electoral fraud and fraud itself.

Specifically, the aim of the paper is to deploy recently-developed data mining and text classification techniques from the burgeoning discipline of data science to assess social-media discussion of electoral misconduct in the referendum campaign, in the aim of determining what offline political phenomena, if any, this discussion reflected. We consider three possibilities. The first possibility is that there was in fact electoral misconduct, and the social media discussion of this phenomenon was based on evidence of wrong-doing uncovered by the initiators of these exchanges. We label this for convenience the 'fraud' hypothesis, and it implies that the users of social media were fulfilling the function of 'social sensors', accurately detecting and reporting offline reality. The second possibility is that the social media discussion of electoral conduct was the result of an orchestrated effort by certain groups to discredit the referendum outcome; the 'plot' hypothesis sees a hard core of social media-savvy political activists shrewdly exploiting online tools to distort popular interpretation of offline electoral events. The third possibility is that there was no actual fraud and no orchestrated campaign to discredit the outcome, but that disappointment expressed online by individual 'yes' supporters led to the development of a form of spontaneous collective delusion in which those involved retroactively reinterpreted their personal

experience of the voting process and projected onto that experience concerns about fraud that had not occurred to them at the time. This we term the 'collective delusion' hypothesis.

The objective of the analysis which follows is to analyze social media traffic before during and after the referendum in order to identify which of these hypotheses best accords with the patterns of interaction we observe. If the 'fraud' hypothesis is correct, we should expect to see concerns mainly being expressed at the time of the voting and counting, tailing off thereafter. If the 'plot' hypothesis is correct, we would in all probability observe discussion of malpractice beginning *before* polling day, as plans to discredit the poll were made and suspicions purposely sowed in the minds of voters. If, on the other hand, the social media discussions of fraud were the result of collective delusion, we would expect to see these discussions grow in volume *after* polling day as the ill-founded concern spread through cyberspace. A final possibility is that a combination of factors was at play; for example, there may have been a limited amount of genuine abuse, which was then magnified in the popular mind via one of the other two mechanisms identified – intentional or unintentional misinformation.

The analysis has both methodological and theoretical ramifications that go well beyond Scottish politics, as it will help us to understand the relationship between the content of social media discussions and offline events, which will further our theoretical understanding of social media as political communication. The investigation will also inform event detection methods that are widely used on social media data (McCreadie et al, 2015; Minhas, Ulfelder and Ward, 2015; D'Orazio et al., 2014).

The paper unfolds as follows: the next section identifies key aspects of the context surrounding the Scottish independence referendum in order to place the hypotheses and the empirical analysis in context. The second section of the paper explores the justification behind the three hypotheses set out briefly above, drawing on the literatures on social media,

social psychology and perceptions of electoral integrity. The third section describes the methods used to collect the data to test these hypotheses, while the fourth section details the results of the analysis. A final brief section concludes.

I The 2014 Scottish context

The Scottish independence referendum of 18 September 2014, which could have led to Scotland being independent of the rest of the United Kingdom, resulted in a vote of 44.7 percent 'yes' and 55.3 percent 'no,' hence maintenance of the Union in its current borders.[2] Perhaps the most interesting statistic is not, however, the overall result (which had long been predicted), but the referendum turnout of 84.6 percent. With electoral participation figures in Scotland and the rest of the UK in decline, the mobilization of over five-sixths of the electorate was a spectacular achievement, resulting in the highest turnout ever recorded in a Scotland-wide vote. Moreover, it meant that many of those voting were citizens who rarely made it to the polls and were largely disengaged from politics-as-usual. The UK Electoral Commission estimated that ten percent of those who took part in the referendum had not before voted (Electoral Commission, 2014b: 63).

The unprecedented popular mobilization occasioned by the referendum can be seen as a positive development from the point of view of political engagement and participation. At the same time, it resulted in a number of unexpected challenges for electoral administrators, challenges that were logistical (the supply of polling venues and materials), educational (explaining the voting process to novice voters) and presentational (maintaining popular confidence in the electoral process).

---

[2]   The question put to the Scottish people was 'Should Scotland be an independent country?'

The UK context is one in which electoral administration had long been taken for granted and was very seldom a topic of political debate (James, 2012). This began to change somewhat following the Labour party's ascent to power in 1997, when electoral reform emerged as one of the new government's strategic goals. Reform gained momentum in the wake of the 2001 General Election, when a slump of turnout to 59 percent caused ripples throughout the political elite, leading to a number of years of hand-wringing about political disengagement. Electoral administrative reform was also given a boost by the formation of the Electoral Commission in 1999, a body that was afforded the capacity to carry out extensive research on electoral administration, including a series of pilots of new voting procedures starting in 2001, together with a number of studies of popular perceptions of the electoral process. Several of the reforms undertaken during this period spurred further debate, including the introduction of postal voting on demand, which is widely believed to have increased the vulnerability of UK electoral procedures to abuse. Following a string of high-profile cases of postal vote fraud in 2003 and 2004, efforts were made to introduce measures to increase the security of elections. Thus the 21st century has been a time of heightened debate about electoral procedures in the UK. Nevertheless, this debate has been conducted largely by a small group of specialists who have a direct interest in the topic, including electoral officials, political elites, public officials and academics (Birch and Watt, 2004; Clarke, 2014a; Electoral Commission, 2014a; James, 2012; Watt, 2006; Wilkes-Heeg, 2008). Prior to 2014, the topic of electoral administration had only episodically engaged large sectors of the general public, and it is possible that this periodic popular interest in electoral administration is linked to the high stakes that have attended a number of recent electoral contests, especially the Scottish referendum on independence.

Electoral misconduct arose as an issue before, during and after the Scottish independence referendum. Prior to the vote, there were a number of reports in the media

about administrative problems in some areas, popular concern that the election might be 'rigged,' and even that the security services could be involved in electoral malfeasance (e.g. McTague, 2014; Maddox, 2014). Following the vote, a widely-reported online petition, which eventually garnered tens of thousands of signatures, demanded a recount of the referendum result on the grounds of fraud.[3] Objects of particular concern included: violence and intimidation during the campaign and on polling day; possible problems with ballot integrity (including postal vote fraud); possible breaches of the chain of custody in the transport of ballot boxes from polling stations to counting centers; and possible miscounting.

With social media playing an integral part in all aspects of the referendum process, from voter mobilization to debate and post mortem analysis, cyberspace was the principal arena in which concerns about electoral integrity were aired. In the end, very few of the allegations raised resulted in formal investigations, and the formal 'yes' camp conceded defeat without protest. Each of the 32 counts was observed by multiple representatives (counting agents) of the two camps, as well as by official observers, elected representatives at various levels and Electoral Commission staff. According to the Electoral Commission, no recounts were requested by those present at any of the counting centers (Electoral Commission, 2014b: 130).

With any event involving millions of people, there are bound to be some problems, however, and the Scottish independence referendum was no exception. In their official report on the referendum, the Electoral Commission listed a number of documented issues: 10 cases of alleged personation in Glasgow; staff involved in some counts having been previously

---

[3]   The petition can be found at: https://www.change.org/p/nicola-sturgeon-we-the-undersigned-demand-a-revote-of-the-scottish-referendum-counted-by-impartial-international-parties.

active in the campaign; voters in some polling stations not folding their ballot papers and not being required to do so by poll workers; a number of instances of attempted buying and selling of votes online; and the 'sampling' (illicit counting) of postal votes during the formal verification process prior to counting (Electoral Commission, 2014b; 132-1). There is also evidence from the Scottish Assessors Association of widespread abuse of the system for applying for proxy votes prior to the referendum (SSA, 2014).

It is unlikely that any of the documented problems could have had more than a miniscule impact on the outcome of the vote. The main effect such problems appear to have had was to raise popular suspicion in a context where many voters were already on the alert. Research conducted by the Election Commission suggested that by and large, the Scottish electorate was satisfied with the conduct of the referendum; a total of 94 percent of survey respondents who voted at polling stations and 98 percent of those who voted by post reported being happy with the way the poll was conducted (Electoral Commission, 2014b: 1). At the same time, the Electoral Commission acknowledged that concerns over electoral integrity were somewhat higher at the time of the referendum than in previous electoral events in Scotland. A total of 34 percent of respondents in the Commission's post-referendum survey thought that fraud had taken place. To put these figures in perspective, 30 percent of those surveyed by the Commission after the 2014 European parliamentary election had believed that fraud had taken place at those elections, and the figure for the Scottish local government elections of 2012 was 28 percent (Electoral Commission, 2014b; 76). Thus, an additional 4-6 percent of the population believed that there had been fraud in the referendum, an increase of approximately 20 percent over previous elections. The Commission also found that 42 percent of 'Yes' voters but only 21 percent of 'No' voters believed there had been fraud in the referendum (Electoral Commission 2014b: 76).

Most who believed there had been fraud claimed to hold this view because of media reports; very few of those who voiced concerns about fraud did so on the basis of personal experience, a fact in which the Commission found consolation: "While electoral fraud is often thought to occur we found no increase in people who had had first-hand experience of fraud. 53% of respondents thought that electoral fraud was not a problem and 26% indicated that they did not know if it was a problem. This reflects the Commission's view that confidence in the safety and security of the voting process can be strongly influenced by the wider context and in particular the reporting of alleged electoral abuse in the media. Just under a third of our survey respondents (32%) reported knowing hardly anything about fraud relating to elections and voting, with 16% saying that they knew nothing at all about the subject" (Electoral Commission, 2014b: 77).

At the same time, the Commission was obliged to recognize that the figure of 73 percent who through voting had been generally safe from fraud was lower than at other recent elections. On the basis of these findings, the Electoral Commission acknowledged the "need to address the issue of how we make the public more aware of count procedures." (Electoral Commission, 2014b: 15, 80), noting that "there continues […] to be a lack of confidence that voting is safe from fraud and abuse when in reality the level of allegations of electoral fraud are low in Scotland" (Electoral Commission, 2014b: 80).

In sum, the evidence available from official bodies suggested very limited actual misconduct, and the problems that did occur were not of such a nature as to have any perceptible impact on the result. The possibility remains that there was wrong-doing, but that it was effectively concealed by the perpetrators. Whatever the truth of the matter, there is no doubt that there developed among a not insignificant minority of the Scottish population serious concerns about the conduct of the referendum.

II Theoretical perspectives on social media and electoral conduct

The role of social media in documenting and discussing electoral misconduct is a nascent area of study that currently lacks a set of accepted theories or robust findings. In developing the hypotheses to be tested here, we draw on the literatures on social psychology and the role of social media in creating and spreading misinformation, together with insights from the growing sub-field of electoral integrity studies.

Three possible links between online discussion of electoral misconduct and offline reality were sketched above. First, it is possible that there was some electoral malpractice in the referendum, and that social media debate was picking up on something that is of genuine concern to electoral governance. If, on the other hand, the claims of electoral misconduct voiced online were entirely or very largely without foundation, then efforts to portray the referendum as having been misconducted were instances of misinformation, which can been intentional or unintentional. The second hypothesis we explore is that of intentional misinformation, namely that concerns about electoral integrity were 'seeded' online in a deliberate attempt to discredit the conduct of the referendum in the eyes of the people, so as to call into question the legitimacy of the result. Third, it is possible that the online frenzy that arose at the time of the referendum around the topic of electoral abuse was an episode of unintentional misinformation, a crisis of mass confidence or collective delusion about electoral conduct. We explore the causal mechanism behind each of these hypotheses in turn, bearing in mind that there could be an element of truth in more than one of them.

These hypotheses can be summarized as follows:

H1: There was more fraud at the referendum than was acknowledged by officials, and social media discussion reflected this reality.

H2: Discussion of fraud on social media was a deliberate attempt by the initiators of the discussion to discredit the referendum.

H3: Discussion of fraud on social media was the result of a collective delusion.

Fraud hypothesis:

The UK is generally recognized as a state with a high degree of electoral integrity, and there is no precedent for large-scale electoral manipulation in mainland Britain (Clarke, 2014a; Electoral Commission, 2014a; James, 2012; Watt, 2006; Wilkes-Heeg, 2008). That said, isolated incidents of electoral abuse have been documented, and there is evidence that the frequency of such incidents has increased in recent years in the United Kingdom (Clarke, 2014a; James, 2012). A report on electoral fraud released by the Electoral Commission in January of 2014 identified 16 local authorities in the UK where elections had repeatedly been beset by problems; none of these is located in Scotland (Electoral Commission, 2014a).

Yet the poor printing of ballot papers in the Scottish Parliament election of 2007 and the subsequent spike in invalid ballots had already raised the salience of electoral administration in the popular mind (Clarke, 2014a: 171). The scathing Gould report on these problems, which found that voters had been treated as an 'afterthought' by electoral administrators, hardly inspired confidence on the part of the electorate, although this episode had also led to a major overhaul of Scottish electoral administration and the establishment of a new Electoral Management Board (Clarke, 2014b). The closeness of the race in 2014 was also undoubtedly a factor in heightening awareness of possible wrong-doing. We know from comparative evidence that fraud is more likely in high-stakes elections (Birch, 2011), and given the stake involved in the independence referendum was the establishment of a new state, they were extremely high in this case. The fact that the referendum was carried out on the basis of new and untested legislation (the Scottish Independence Referendum Act 2013) further increased the possibility that problems would arise, as potential fraudsters were more likely to be able to find loopholes in the law that they could exploit. It is thus entirely

possible that there was at least some wrong-doing in the referendum, and all else being equal, wrong-doing is more likely in this poll than in most electoral events held in the UK.

Serious concerns were raised by a number of groups that analyzed the referendum. Russian observer Georgy Fyodorov, of the Association for the Protection of Electoral Rights "Civil Control" told a Russian news agency that "according to what our observers at the polling offices tell us, there were more Yes votes during the vote count … Those on the UK side campaigning for a No vote resorted to every violation imaginable" (Sputnik News, 2014). Well-know American political commentator Naomi Wolf publicized concerns that many referendum ballots lacked the Unique Identifying Numbers necessary to make them valid (Wolf, 2014). Another analysis focused on suspiciously high postal vote turnout in Argyle and Bute, claiming that known inaccuracies in the electoral register together with natural trends in population movement and mortality mean that the reported postal vote turnout of over 95 percent would have been virtually impossible to achieve without manipulation (Democratic Socialist Federation, n.d.). Unlike the official reports cited above, these reports rated as serious the irregularities that allegedly marred the referendum.

Plot hypothesis:

The use of media to disseminate unwarranted concerns about electoral conduct during the pre-electoral period is regularly practiced by opposition activists in many countries, especially when it is clear in advance of the election that they have little hope of winning (Beaulieu, 2014; Beaulieu and Hyde, 2009; Schedler, 2013: chap. 9). In many electoral contexts, opposition leaders will prime the electorate to expect fraud by means of communication strategies that emphasize risks to electoral integrity and speculate about possible plans by their political rivals (in particular incumbent leaders) to 'steal' the election, as was evident in the 2016 US presidential election campaign. A recent paper on local newspaper coverage of electoral fraud in the US finds the number of stories on electoral

abuse to be highest in states where fraud had been a topic of debate and legislative activity at the elite level, but that coverage was unrelated to the incidence of actual fraud in the state (Fogarty et al, 2015). If effective, such priming strategies could be expected to lead voters to interpret innocuous occurrences or minor logistical snarls as evidence of serious electoral misconduct. It is thus plausible that political elites and activists might seek to use the social media platforms at their disposal to prime voters in advance of the poll to be suspicious about the electoral process.

In the Scottish context, one question that immediately arises is who might seek to discredit the outcome of the independence poll. Much of the debate about electoral malpractice took place in online spaces dominated by the 'Yes' camp. The 'Yes' camp had a significant challenge in that a minority of Scottish voters had historically favored independence. In the final months of the campaign, opinion polls suggested that though the 'No' side remained in the lead, the 'Yes' camp was gaining ground; two polls released in the fortnight leading up to the vote even put 'Yes' ahead of 'No'. Yet on the eve of the poll, most surveys still pointed to a 'No' win (see Figure 1). It is therefore likely that any plan to discredit the referendum would have come from supporters of the 'Yes' option.

Figure 1 about here

There is some evidence of an orchestrated campaign to breed suspicion of the voting process among electors. The Electoral Commission reported having received a number of queries prior to polling day about the use of writing implements to mark ballots (pens v pencils), the folding of ballot papers, and the role of the Unique Identifying Number on the back of the ballot paper (Electoral Commission, 2014b: 134-6). Political scientist and official referendum election observer Alistair Clarke reported that many voters arrived at the polling station already suspicious about the process in which they were taking part (Clarke, 2014c). As noted above, there had been several newspaper reports prior to election day of possible problems,

though it is also possible that many voters' suspicions were derived from other sources. The question that concerns us here is whether there is any evidence that social media was systematically seeded in advance of polling day with messages designed to breed suspicion about the electoral process.

Collective delusion hypothesis:

One of the key characteristics of social media is that traditional forms of political leadership and planning are not required for mass mobilization to occur; the spontaneous formation and expression of group sentiment is enabled by the open and non-hierarchical format of the online world (Bennett and Segerberg, 2013; Margetts et al, 2015). This means that it is perfectly possible for collective delusion to have arisen online organically, without 'plotters' intentionally seeding cyberspace with concerns about electoral fraud. What interests us for the moment is why and how such a phenomenon might arise.

The growing literature on confidence in elections has identified a considerable gap even in established democracies between elite and popular perceptions of the integrity of electoral processes (Anderson *et al*., 2005; Atkeson and Saunders, 2007; Birch, 2010; Hall, Monson and Kelly, 2009). The literature on 'losers' consent' has found that in many states, being on the losing side of an electoral contest is associated with a more critical view of electoral conduct (Anderson et al., 2005), and that voters' estimation of the integrity of the election changes following the outcome (Cantu, 2014).

The literature on conspiracy theories is also instructive in this regard. In their book-length study, Uscinski and Parent identify 50 major conspiracy theories in the United States between 1890 and 2010. Of these, 14 percent (7) were theories about election rigging (Uscinski and Parent, 2014: 59-62). This is a remarkably high proportion, given the many things people might theorize about; it therefore appears that electoral conduct is a topic that is

highly vulnerable to conspiracy theorizing. Uscinski and Parent concur, arguing that "conspiracy theories are for losers" in the sense that those who support losing electoral options are more likely to espouse beliefs in that covert manipulation of elections by malign political agents was responsible for the outcome (Uscinski and Parent, 2014: 22, 91-2, 138-43). In other research, speculative and even fictional accounts of conspiracy theories involving state actors have been found to make people more willing to believe in such theories (Butler, Koopman and Zimbardo, 1995; Douglas and Sutton, 2008; Mulligan and Habel, 2013), and we can anticipate that the same would be true of online discussion.

The 'sour grapes' argument chimes with research on confirmation bias. We know that people are 'motivated reasoners' who rate more highly arguments associated with outcomes of which they approve (Baron and Hershey, 1988; Kunda, 1990), and that ideological position has a strong impact on political belief, with people more likely to convince themselves of theories that confirm their ideological world view (Lodge and Taber, 2013; Pasek et al., 2014; Redlawsk, 2002). In as much as many social media users were hooked into dense networks of other users whose views reinforced each other, social media use during the polarized atmosphere of the referendum campaign may have served to magnify this effect.

Evidence from social psychology suggests that belief corroboration tends to lead to more extreme views (Baron *et al*., 1996), and some scholars have maintained that new media usages leads to online sorting or homophily, generating 'echo chambers' where people tend to consume information that reinforces their pre-existing opinions (Sunstein, 2002; Farrell, 2012). If this theory is correct, the online sorting effect could have served to increase the perceived veracity of claims of fraud. In other words, observation of these claims being repeated strengthened people's impression that they must reflect valid concerns about the conduct of the referendum.

Yet there is also evidence from recent studies that the homophily characteristic of much online activity is less pronounced in Twitter use, which has the effect of reducing issue polarization (Barberá, 2014; 2015; Yardi and Boyd, 2010). People monitoring Twitter streams organized around particular hashtags are likely to be exposed to a wide variety of competing beliefs, including views that differ considerably from their own. There is thus reason to believe that interaction on at least some social media platforms should lessen the 'echo chamber' effect. This could be expected to work against the generation of collective delusion. On Facebook, where people select the contacts from whom they get postings though, homophily is likely to be more marked.

The final possibility that deserves to be explored is that social media served to significantly magnify in the popular mind what was in fact a very limited amount of electoral abuse. When people hear that a particular problem may have affected 'dozens' or 'hundreds' of ballots, this many seem like a large number, whereas in fact it represents a minute proportion of the overall electorate and is highly unlikely to have a material impact on the outcome. The reported issues with proxy voting mentioned above are a case in point, as are reports of polling cards being sent to small children in East Ayrshire (Maddox, 2014). There are two common thought processes which could account for this phenomenon: miscalculation and the propensity to view available evidence as the 'tip of the iceberg.' We know from behavioral economics that many people have a poor understanding of basic mathematical concepts such as proportions and frequencies (Gigerenzer, 2002; Tversky and Kahnemann, 1982). We also know from research on perceptions of politics that British citizens have a tendency to view revealed wrongdoing as the 'tip of the iceberg' (Allen and Birch, 2015). Reports of minor irregularities could thus get magnified in the popular mind if extensively publicized.

In the UK, concern increased following a spate of instances of electoral abuse in the mid 2000s that were widely reported in the media (Wilkes-Heeg, 2008; James, 2012). At the time of the Scottish independence referendum, there were ongoing investigations of several cases of suspected electoral misconduct alleged to have taken place in the 2014 local elections, including a widely-reported and high-profile case in the London borough of Tower Hamlets. Another factor that is likely to have magnified concern is the fact that many of those who took part in the referendum were not habitual voters and were thus not familiar with voting procedures. It is possible that social media communication around the referendum enhanced these concerns in the minds of some voters and caused them to believe that widespread fraud had taken place in the absence of objective evidence of irregularities. Thus conjecture about possible fraudulent acts, however outlandish, might have increased people's propensity to believe in conspiracy theories.

III Modeling and Classifying Twitter Discussion of Fraud During the Independence Referendum

In order to distinguish between the three interpretations empirically, we rely on the temporal distribution of social media activity that we would expect each of these three processes to involve. We select Twitter for empirical testing of our models due to its public nature and availability. We focus in particular on the initiation of Twitter traffic related to fraud, as the onset of such Twitter activity represents a useful observable implication of our theoretical expectations.

If the referendum was beset by substantial actual fraud of the type alleged, involving voter registration, voting and vote counting, we would expect to see online expressions of concern begin to mount gradually in the hours leading up to the opening of the polls, and gain intensity during the period of approximately 24 hours when the votes were cast and counted and the results widely disseminated, as it is during this period that voters would be tuned into

17

the mechanics of the electoral process and it is then that they would have had the opportunity to witness malpractice first-hand.[4]

If, by contrast, the online discussion of fraud was intentionally 'seeded' in order to raise false concerns, we would expect to see the initiation of large amounts of online activity well *before* polling day, as efforts were made to sow doubts in people's minds as to the credibility of the voting and encourage them to look for flaws in the process when they went to vote.

Finally, if the online discussion of electoral irregularities was the result of collective delusion, we would expect to see a spike in social media activity immediately *after* the results were announced, as disillusioned 'Yes' supporters reinterpreted their voting experience and began to believe accounts of abuse.[5]

It is possible that what we will observe empirically is a mix of these three pure models, but the onset of intensive discussion of electoral fraud should provide a good indicator of which hypothesis has the most traction.

---

[4] Whether Twitter activity of this sort dies off gradually thereafter as authorities either address the problem or refuse to countenance allegations of wrong-doing, or whether activity continues to increase as concerns mount and 'informational cascades' (Lohmann, 1994) form is not particularly relevant to this analysis, as our main interest is in when Twitter activity around fraud *begins* in earnest.

[5] In a study of tweets in Arabic about the Boston Marathon bombing of 2013, tweets classed as evidencing beliefs in conspiracy theories peaked over a week after the event itself, whereas tweets that fell into the four other categories employed by the researchers all peaked sooner, mostly within 48 hours of the event (Jamal et al, 2015: 65).

We now detail how we identified, categorized, harvested and analyzed relevant social media discussions in order to test these alternative possibilities. For the sake of convenience, we categorize possible acts of electoral misconduct according to stage of the referendum process: violations of democratic campaign norms prior to the vote; violations of voting procedures; problems with the 'chain of custody' that ensures ballots are handled appropriately and securely from the time they are cast to the time they are counted; and problems with the counting of ballots. During the campaign, there were also concerns revolved around allegations of inappropriate acts of harassment, intimidation and actual or threatened violence; we classify discussion of violence of this type as misconduct also.

As noted above, we rely on analysis of communication on Twitter, due to its public nature, its instantaneous dissemination and its ready availability to researchers. The public nature of Twitter means that tweets were widely disseminated among large groups who were actively involved in online discussions of the referendum. The instantaneous aspect of Twitter makes it possible to undertake very fine-grained analysis of temporal trends. The availability of tweets renders them an amenable data source; unlike the partially-public world of Facebook, the Twitterverse is almost all in the public realm, obviating concerns about the possible confounds associated with personal contacts between users.

We collected a corpus of tweets sent during a period extending from 10 days before to 10 days after the Scottish independence referendum, which took place on 18 September 2014 – in other words, from 8 to 28 September. The corpus collected includes 222 gigabytes of material. Filtration was used to remove non-English language tweets and to get retain only the tweets that contain relevant keywords known from previous research and perusal of news websites to be associated with the independence referendum. These include: Scotland, Scottish, Referendum, Voting, Votes, Indyref, ScottishAnimals, Ballots, and Elections. This filtration resulted in a corpus of 8,120,039 tweets. A 10 per cent sample of this corpus was

drawn from each day to reduce the processing time required for the classification task; this included the first one per cent of each decile of Tweets for that day. The resulting sample for the 21-day period included 999,790 tweets, which was used in the subsequent analysis.

In order to identify the best classification strategy, we tested three classifier models commonly used in the field of machine learning: Naïve Bayes (NB) (Weiss, Indurkhya, Zhang, & Damerau, 2005), Decision Trees (DT) (Manning & Schütze, 1999) and Support Vector Machine (Kubat, 2015). These tests, detailed in the Appendix, demonstrated that the Support Vector Machine approach was superior to the other strategies.

IV Results

The Support Vector Machine model was applied to classify the entire corpus of 999,790 tweets harvested from the period 8 – 28 September 2014. The results of the classification process are shown in the Table 1.

Table 1 about here

We next divide each day to four time periods: the first from midnight to 5.59am, the second from 6.00am to 11.59am, the third from noon to 5.59pm and the last from 6.00pm to 11.59pm. The spread of the fraud tweets over these four periods is shown in Figure 2.

Figure 2 about here

The data in this graph are striking in what they reveal; with the exception of a brief spike in tweets on 13 September, which appears to have been mainly due to multiple retweets of a single rather vague tweet about fears of electoral manipulation,[6] the vast majority of tweets related to fraud occur *after* the announcement of the referendum result. Moreover, there is a

---

[6] The tweet reads: 'Be left in absolutely no doubt Scotland we are being manipulated at the highest level. Take a deep breath; think logically'.

very large spike in such tweets immediately after the result was announced and it became

apparent that the 'no' option had won. The number of tweets documenting fraud on the day

of the referendum (18 September) is miniscule, suggesting that almost no polling fraud was

observed by social media users. It is also noteworthy that there are relatively few tweets

about fraud prior to the vote, undermining the 'plot' hypothesis. The only hypothesis that

receives support from these data is the collective delusion hypothesis; the pattern of tweets is

strongly consistent with the theory that fraud occurred to some voters retrospectively as a

potential explanation for the result. It appears that a portion of 'yes' supporters were

unprepared to accept the result of the referendum, and fraud was their 'sour grapes' account

of what went wrong. This finding is consistent with previous research in the 'loser's consent'

vein; we know from comparative data that those on the losing side in an electoral contest

often cry foul as a way of reconciling their expectations with the electoral outcome

(Anderson et al., 2005; Cantu, 2014). The results of this analysis also support the

interpretation of the online referendum fraud debate by most commentators, including most

who supported a 'yes' vote. In this sense they are not surprising, but it will be of use to

electoral administrators and to students of electoral fraud alike to have an empirical basis for

rejecting the 'plot' and 'fraud' hypotheses.

V Conclusion

UK electoral legislation is based on the assumption that electoral misconduct will rarely be

attempted, but it does still occur. Neither the hypothesis of actual fraud nor the hypothesis of

a concerted effort to seed doubt in voters' minds in the run-up to the voting is supported by

the data presented here. Instead, tweets about fraud spiked immediately after the results were

released, which is consistent with the interpretation that this online activity was a

manifestation of collective delusion borne out of an inability or unwillingness to accept the outcome of the referendum.

At the same time, there is still cause for concern when a section of the citizenry is not confident in the conduct of an electoral process. Aspects of the UK electoral process lack a range of security measures that are common in states where electoral abuse is more common. With confidence in electoral conduct falling, it may well be time to consider introducing some of these measures so as to address concerns about fraud and other forms of abuse. The requirement to show identity cards at polling stations has already been recommended by the Electoral Commission (Electoral Commission, 2014a) and is being piloted by the UK government in 2017; other measures that might usefully be considered include stricter regulation of the chain of custody of ballot papers, more rigorous training of polling station staff, and limitations on the availability of postal and proxy votes. Even if such measures do not markedly improve what is already a very high standard of electoral performance in most parts of the UK (including Scotland), they are likely to go a significant way toward improving confidence in electoral conduct.

These results of this analysis have implications for the burgeoning efforts to employ data mining techniques to identify incidents of fraud in elections around the world (e.g. Yang, Ounis and Macdonald, 2016). Though such techniques may be very useful in many contexts, they will need to be capable of sifting through large amounts of speculative allegations lacking evidence in order to identify reports of fraud that are corroborated. The findings of this paper also hint that such techniques must be used with extreme caution in contexts where fraud is not expected a priori.

A final lesson to be learnt from this paper is the utility of combining the traditional approaches and tools of political science with recent developments in data mining and text

classification. The analysis presented here shows the huge benefits that can be gained by reaching beyond the tools of political science and engaging in multidisciplinary data science.

Future research could usefully examine the pattern identified here in a wider range of contexts, in order to determine whether the timing of 'Twitter storms' can be used a reliable means of interpreting them. If this is the case, users of this method would then need to determine under what conditions this method might *not* work well, an also how easy it might be for Twitter users to manipulate their behavior in order to 'game' such interpretations.

Another future research area that follows from this analysis is the link between online and offline behavior; it would be of interest to examine, for example, whether those who are more concerned about fraud are also more likely to be social media users, or to verify the supposition that those who find themselves on the losing side in an election or a referendum are more likely to tweet about fraud. The role of online allegations of fraud in conditioning actual electoral behavior is also an interesting area of possible investigation; a study of 'fake news' in the US 2016 presidential election demonstrated that false news stories were believed by over half of all those who viewed or read them (Allcott and Matthew Gentzkow, 2017), raising concerns about the ability of Twitter storms such as that studied here to condition the future orientations of their participants toward politics. The advent of social media has opened up a wide range of research possibilities to computer scientists and data scientists alike, and the study of social media content about fraud promises to enrich our understanding of both electoral integrity and online communication.

Appendix: Classification Strategy Selection Procedure

This appendix details the methods used to select the Support Vector classification strategy used on this analysis. Three commonly-used classifier models were tested: Naïve Bayes (NB) (Weiss, Indurkhya, Zhang, & Damerau, 2005), Decision Trees (DT) (Manning & Schütze, 1999) and Support Vector Machine (Kubat, 2015).

The models were implemented using the algorithm packages e107 and C50 in the R programming language. The task of the classifier was to classify the tweets according to whether or not they contain allegations of fraud. Following best practice in machine learning, a two-stage process was employed in order to refine the classification. Two different datasets were created and used sequentially; each was split into a training dataset and a testing dataset. The first dataset (DS1) was generated by filtering the 10 per cent sample using fraud-related keywords (rigged, fraud, claims, investigate and police). This dataset contains 5,413 tweets with about 10 percent 'positive' tweets (those containing allegations of fraud) and about 90 percent 'negative' tweets (those not containing allegations of fraud). DS1 was split into 3,788 tweets as a training dataset and 1,625 tweets as a test dataset. When DS1 was used to train the three models (NB, SVM and DT), DS worked only if the MostFreqTerms (most frequent terms) setting was greater than or equal to 50. The MostFreqTerms setting entails removal of a term from the vector space model if it does not appear in at least 50 documents. This setting is risky; if we used it we sacrifice the words that occur less frequently, and these words could also be important features for the model to learn from. Table A1 show the results of the three models on the testing set using different settings. For fewer 'most frequent words' (found in at least 10 documents) SVM gave better results than NB.

− Table A1 about here −

The second dataset (DS2) was built by applying the NB and SVM models to the first dataset. The tweets classified as referring to fraud were grouped with tweets containing words most

frequently found in the fraud tweets via semi-supervision classification. This method is typically used with very large datasets that cannot be checked manually. DS2 contains 5,782 tweets split into a 4,044-tweet training set and a 1,738-tweet testing set. The dataset contains 570 'positive' tweets and 5,212 'negative' tweets. The results of the two models are shown in Table A2.

– Table A2 about here –

In order to check for over-fitting, a virgin dataset was drawn from the original corpus by filtering on keywords 'fraud' and 'rigging'. Although the results of the above-reported testing set are better for SVM than NB, when the two models were used to classify the new unlabelled tweets (a virgin dataset), the results of the NB classifier were superior when compared to human classification (manual coding) of the same virgin dataset. The human classification data represents the average classifications of three different human coders who classified the same dataset manually. The performance of the two models on the virgin dataset is compared to the human classification in Table A3. The only limitation of NB model is that it is much slower than SVM. This experiment demonstrated that to classify the virgin data NB model took 21 minutes, whereas SVM took only two minutes when the 'most frequent words' setting was set to 1.

- Table A3 about here -

The preliminary results of these experiments thus show that the Naïve Bayes model retrieves more positive tweets than the Support Vector Machine model. This means that NB is higher in recall than SVM. However, when a human went through the positively labeled tweets resulting from the models (the second experiment with MostFreqTerms=5), it was found that 80 tweets out of 2,382 were misclassified as positive by the NB model with error rate 0.033 per cent. In comparison, 64 tweets out of 2,032 were misclassified as positive by SVM with error rate 0.031 per cent. This means that there is no major difference between the

performance of the two models on new unlabelled data. This shows that the Support Vector

Machine model retrieves more accurate positive tweets and is this higher in precision than the

Naïve Bayes mode. In this research precision is more important than recall, so SVM is the

model that is used to classify tweets in this analysis.

References:

Allcott, Hunt and Matthew Gentzkow (2017), "Social Media and Fake News in the 2016 Election", NBER Working Paper 23089, http://www.nber.org/papers/w23089

Allen, Nicholas and Sarah Birch (2015*), Ethics and Integrity in British Politics: How Citizens Judge Their Politicians' Conduct and Why it Matters*, Cambridge: Cambridge University Press.

Anderson, Christopher J., André Blais, Shaun Bowler, Todd Donovan and Ola Listhaug, (2005), *Losers' Consent: Elections and Democratic Legitimacy*, Oxford: Oxford University Press.

Atkeson, Lonna Rae and Kyle L. Saunders (2007), "The Effect of Election Administration on Voter Confidence: A Local Matter?" *PS: Political Science & Politics* 40:4, 655-60.

Barberá, Pablo (2014), "How Social Media Reduces Mass Political Polarization: Evidence from Germany, Spain, and the U.S.", unpublished manuscript.

--- (2015), "Birds of the Same Feather Tweet Together: Bayesian Ideal Point Estimation Using Twitter Data", *Political Analysis* 23, 76-91.

Baron, Jonathan and John C. Hershey (1988), "Outcome Bias in Decision Evaluation", *Journal of Personality and Social Psychology* 54(4): 569–579.

Baron, R., Hoppe, S., Linneweh, B., and Rogers, D. (1996), "Social Corroboration and Opinion Extremity", *Journal of Experimental Social Psychology, 32*, 537-60.

Bartholomew, Robert E. (2001), *Little Green Men, Meowing Nuns and Head-Hunting Panics: A Study of Mass Psychogenic Illness and Social Delusion*, Jefferson, NC: McFarland.

Beaulieu, Emily (2014), *Electoral Protest and Democracy in the Developing World*, Cambridge: Cambridge University Press.

Beaulieu, Emily and Susan D. Hyde (2009), "In the Shadow of Democracy Promotion: Strategic Manipulation, International Observers and Election Boycotts", *Comparative Political Studies* 42:3, 392-415.

Birch, Sarah (2010), "Perceptions of Electoral Fairness and Voter Turnout", *Comparative Political Studies*, 43:12, 1601-1622.

--- (2011), *Electoral Malpractice*, Oxford: Oxford University Press.

Birch, Sarah and Bob Watt (2004), "Remote Electronic Voting: Free, Fair and Secret?", *Political Quarterly*, 75:1, 60-72.

Brigadir, Igor, Derek Greene and Pádraig Cunningham (2015), "Analyzing Discourse Communities with Distributional Semantic Models", *ACM Web Science 2015,* http://websci15.org/.

Butler, Lisa D., Cheryl Koopman and Philip G. Zimbardo (1995), "The Psychological Impact of Viewing the Film "JFK": Emotions, Beliefs, and Political Behavioral Intentions", *Political Psychology* 16:2, 237-57.

Cantu, (2014), "Partisan Losers' Effects: Perceptions of Electoral Integrity in Mexico", working paper.

Clarke, Alistair (2014a), "Investing in Electoral Management", in Pippa Norris, Richard W. Frank and Ferran Martínez i Coma (eds.), *Advancing Electoral Integrity*, Oxford: Oxford University Press, 165-88.

--- (2014b), "The Unsung Heroes? Election Administration in the Scottish Independence Referendum", *Political Insight,* http://www.psa.ac.uk/insight-plus/blog/unsung-heroes-election-administration-scottish-independence-referendum

--- (2014c), "Trouble at the Polls…? Reflections on the Referendum Process", *Political Insight*, http://www.psa.ac.uk/insight-plus/blog/trouble-polls%E2%80%A6-reflections-referendum-process

Democratic Socialist Federation (n.d.), "Defending Democracy: Democratic Socialist Federation Dunoon Unit Report: The Postal Ballot at the Scottish Independence Referendum", https://docs.google.com/document/d/13OPs4c91U4ggD1XrHWGAig8YOoXbehVSEpGw aJJWtpc/pub

Diamond, Larry (2010), "Liberation Technology", *Journal of Democracy* 21.3, pp. 69–84.

Electoral Commission (2014a), *Electoral Fraud in the UK: Final Report and Recommendations*, London: Electoral Commission.

--- (2014b), *Scottish Independence Referendum: Report on the Referendum Held on 18 September 2014*, London: Electoral Commission, http://www.electoralcommission.org.uk/__data/assets/pdf_file/0011/146927/The-Scottish-referendum-Response-to-consultations.pdf

Farrell, Henry (2012) "The Consequences of the Internet for Politics", *Annual Review of Political Science* 15: 35–52.

Fogarty, Brian J., Jessica Curtis, Patricia Frances Gouzien, David C. Kimball, and Eric C. Vorst (2015), "News Attention to Voter Fraud in the 2008 and 2012 US Elections", *Research and Politics*, DOI: 10.1177/2053168015587156, May 2015

Gigerenzer, Gerd (2002), *Reckoning with Risk: Learning to Live with Uncertainty*, London: Penguin.

Hall, Thad, J. Quin Monson and Kelly D. Patterson (2009), "The Human Dimension of Elections: How Poll Workers Shape Public Confidence in Elections" *Political Research Quarterly* 62:3, 507-22.

Holby, Tom (2104), "#Indyref: A Geographical Comparison of Tweets Related to the Scottish Independence Referendum of 2014", MA thesis, Geographical and GeoInformation Science, George Mason University.

Howard, Philip N. (2010), *The Digital Origins of Dictatorship and Democracy: Information, Technology and Political Islam*, Oxford: Oxford University Press.

Ikufor, Presley (2010), "'Elections' or "'Selections'? Blogging and Twittering the Nigerian 2007 General Elections", *Bulletin of Science, Technology & Society* 30:6, 398–414.

Jamal, Amaney A., Robert O. Keohane, David Romney and Dustin Tingley (2015), "Anti-Americanism and Anti-Interventionism in Arabic Twitter Discourses", *Perspectives on Politics* 13(1), 55-73.

James, Toby S. (2012), *Elite Statecraft and Election Administration: Bending the Rules of the Game?*, Basingstoke: Palgrave-Macmillan.

Kubat, M. (2015), "Inter-Class Boundaries: Linear and Polynomial Classifiers". In *An Introduction to Machine Learning* (pp. 65–90). Springer International Publishing, http://dx.doi.org/10.1007/978-3-319-20010-1_4

Kunda, Ziva (1990), "The Case for Motivated Reasoning", *Psychological Bulletin*, 108(3): 480-98.

Le, Huyen T., G.R. Boynton, Yelena Mejova, Zubair Shafiq and Padmini Srinivasan (2017), "Revisiting The American Voter on Twitter", *CHI 2017*, DOI: http://dx.doi.org/10.1145/3025453.3025543

Lodge, M. and C. Taber (2013), *The Rationalizing Voter*, Cambridge: Cambridge University Press.

Lohmann, Susanne (1994), "The Dynamics of Informational Cascades: The Monday Demonstrations in Leipzig, East Germany, 1989-91", *World Politics* 47(1): 42-101.

McCreadie, Richard, Karolin Kappler, Magdalini Kardara, Andreas Kaltenbrunner, Craig Macdonald, John Soldatos and Iadh Ounis (2015), "SUPER: Towards the use of Social Sensors for Security Assessments and Proactive Management of Emergencies",

WWW 2015 Companion, proceedings of the International World Wide Web Conference 2015, http://dx.doi.org/10.1145/2740908.2741723.

McTague, Tom (2014), "One in Four Scots Believe British Spies are Secretly Working against the 'Yes' to Independence Campaign', *Mailonline* 8 September, http://www.dailymail.co.uk/news/article-2747809/One-four-Scots-believe-British-spies-secretly-working-against-Yes-independence-campaign.html

Maddox, David (2014), "Scottish Independence; Vote-Rigging Fears", *Scotsman* 7 September, http://www.scotsman.com/news/politics/top-stories/scottish-independence-vote-rigging-fears-1-3533762

Manning, C. D., & Schütze, H. (1999), *Foundations of Statistical Natural Language Processing*, MIT Press.

Margetts, H, P John, S Hale and T Yasseri (2015), *Political Turbulence: How Social Media Shape Collective Action*, Princeton, NJ: Princeton University Press.

Minhas, Shahryar, Jay Ulfelder and Michael D. Ward (2105), "Mining Texts to Efficiently Generate Global Data on Political Regime Types", *Research and Politics*, July-September, 1–8.

Mulligan, Kenneth and Philip Habel (2013), "The Implications of Fictional Media for Political Beliefs", *American Politics Research* 41:1, 122-46.

Norris, Pippa, Richard Frank and Ferran Martínez i Coma eds. (2015), *Contentious Elections: From Ballots to Barricades*, London and New York: Routledge.

D'Orazio, Vito, Steven T. Landis, Glenn Palmer and Philip Schrodt (2014), "Separating the Wheat from the Chaff: Applications of Automated Document Classification Using Support Vector Machines", *Political Analysis*, doi:10.1093/pan/mpt030

Pasek, Josh, Tobias H. Stark, Jon A. Krosnick and Trevor Tompson (2014), "What Motivates a Conspiracy Theory? Birther Beliefs, Partisanship, Liberal-Conservative Ideology,

and Anti-Black Attitudes", *Electoral Studies*, pre-print published online at
http://dx.doi.org/10.1016/j.electstud.2014.09.009 0261-3794.

Pedersen, Sarah, Graeme Baxter, Simon Burnett, Ayse Goker, David Corney, and Carlos
Martin (2014), "Backchannel Chat: Peaks and Troughs in a Twitter Response to
Three Televised Debates During the Scottish Independence Referendum Campaign
2014", available from OpenAIR@RGU at http://openair.rgu.ac.uk.

Quinlan,Stephen, Mark Shephard and Lindsay Paterson (2015), "Online Discussion and the
2014 Scottish Independence Referendum: Flaming Keyboards or Forums for
Deliberation?", *Electoral Studies* 38, 192-205.

Redlawsk, David P. 2002. "Hot cognition or cool consideration? Testing the Effects of
Motivated Reasoning on Political Decision Making." *The Journal of Politics*, 64(4):
1021-1044.

Scottish Assessors Association (2014), "Scottish Independence Referendum – 18 September
2014 Review: Scottish Assessors Association Electoral Registration Committee",
http://www.scottish.parliament.uk/S4_ScotlandBillCommittee/General%20Document
s/Scottish_Assessors_Association_SIR_Review.pdf.

Schedler, Andreas (2013), *The Politics of Uncertainty: Sustaining and Subverting Electoral
Authoritarianism*, Oxford: Oxford University Press.

Shephard, Mark, Stephen Quinlan, and Stephen Tagg and Lindsay Paterson, (2014), "'Fading
to Grey': Consideration and Opinion Moderation in the 2014 Scottish Independence
Referendum", Fourth annual general meeting of the European Political Science
Association, Edinburgh, June.

Sputnik News (2014), "Russian Observers Say Westminster Put Pressure on Scots on
Referendum Day", 19 September,
http://sputniknews.com/world/20140919/193097294.html

Schedler, Andreas (2013), *The Politics of Uncertainty: Sustaining and Subverting Electoral Authoritarianism*, Oxford: Oxford University Press.

Sunstein, Cass (2002), *Republic.com.* Princeton, NJ: Princeton University Press.

Tierney, Stephen (2015), "Reclaiming Politics: Popular Democracy in Britain after the Scottish Referendum", *Political Quarterly* 86(2), 226-33.

Tversky, Amos and Daniel Kahneman (1982), "Belief in the Law of Small Numbers", in Daniel Kahneman, Paul Slovic and Amos Tversky (eds.), *Judgment under Uncertainty: Heuristics and Biases*, Cambridge: Cambridge University Press, 23-31.

Uscinski, Joseph E. and Joseph M. Parent (2014), *American Conspiracy Theories*, Oxford: Oxford University Press.

Watt, Bob (2006), *UK Election Law: A Critical Examination,* London: Routledge-Cavendish.

Weiss, S., Indurkhya, N., Zhang, T., & Damerau, F. (2005), "Using Text for Prediction", in *Text Mining* (pp. 47−84). Springer New York, http://dx.doi.org/10.1007/978-0-387-34555-0_3

Wilks-Heeg, Stuart (2008), *Purity of Elections in the UK: Causes of Concern*, York: Joseph Rowntree Trust.

Wilson, Andrew (2005), *Virtual Politics: Faking Democracy in the Post-Soviet World*, New Haven, CT and London: Yale University Press.

Wolf, Naomi (2014), Facebook post:

https://www.facebook.com/naomi.wolf.author/posts/10152698425099476.

X. Yang, I Ounis, and C Macdonald, 2016, "Monitoring Electoral Violence through Social Media: A Machine Learning Approach", in International Conference on Computational Social Science, Evanston, IL, USA, 24-26 Jun 2016, http://eprints.gla.ac.uk/120312/

Yardi, Sarita and Danah Boyd (2010), "Dynamic Debates: An Analysis of Group Polarization

Over Time on Twitter", *Bulletin of Science, Technology & Society* 30:5, 316–27.

Figure 1: Opinion poll results in the run-up to the referendum
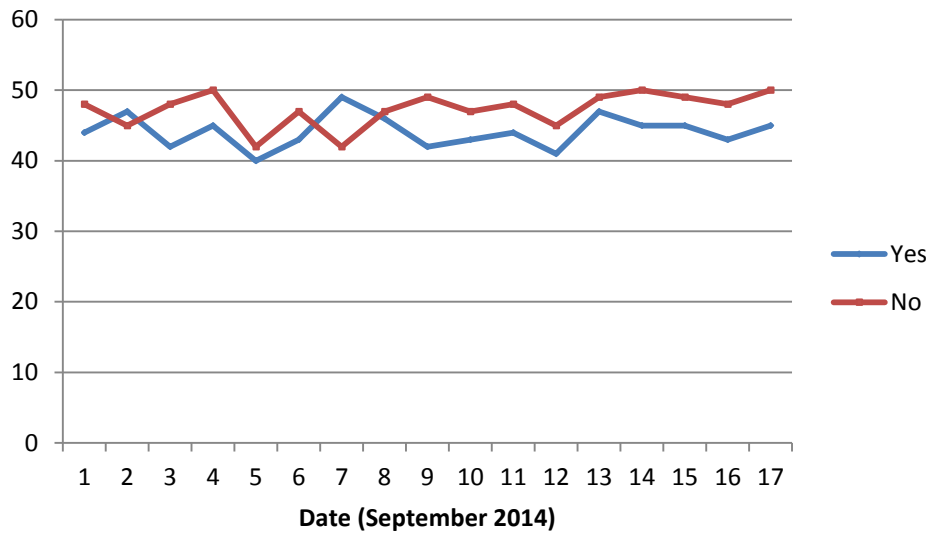
Source: Websites of polling firms Panelbase, YouGov, Survation, ICM, Opinium and Ipsos Mori.

Notes: Data were taken from all polling organisations that are members of the British Polling Council. Figures are for likely voters ('don't know' and 'undecided' figures not shown); dates are the midpoint of each period of fieldwork.

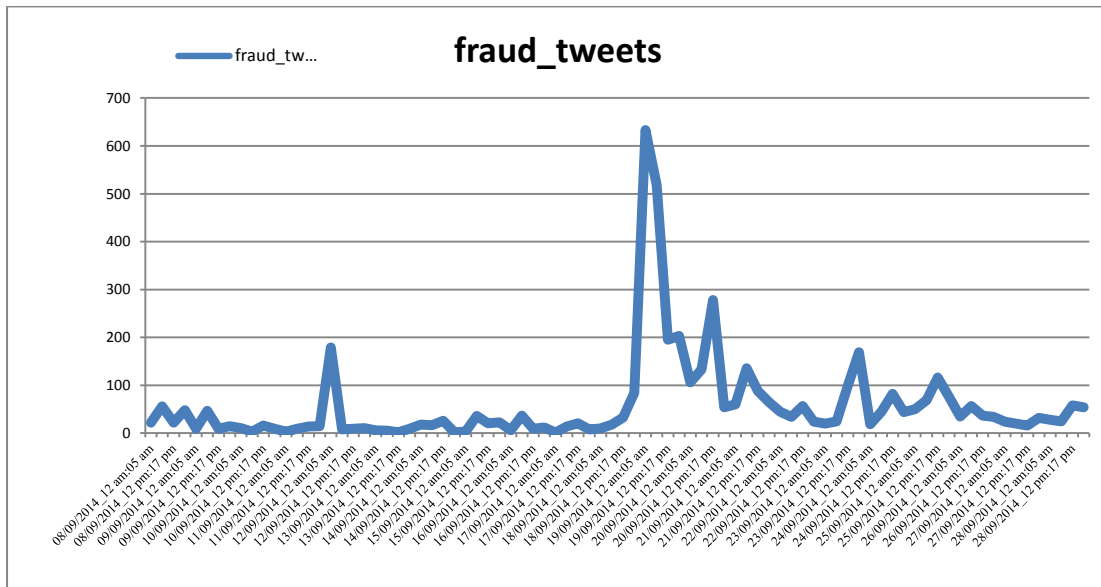Figure 2: The distribution of fraud tweets over the time

Table 1: Distribution of Tweets by Day

Date	Total number of tweets	Fraud tweets	Non-fraud tweets
08/09/14	29440	148	29292
09/09/14	33960	80	33880
10/09/14	37000	39	36961
11/09/14	41070	41	41029
12/09/14	32220	207	32013
13/09/14	33980	22	33958
14/09/14	33730	63	33667
15/09/14	39140	83	39057
16/09/14	47980	64	47916
17/09/14	57550	44	57506
18/09/14	126900	145	126755
19/09/14	168620	1549	167071
20/09/14	31620	571	31049
21/09/14	16980	349	16631
22/09/14	13050	160	12890
23/09/14	21942	313	21629
24/09/14	11270	190	11080
25/09/14	9170	311	8859
26/09/14	9100	163	8937
27/09/14	8700	92	8608
28/09/14	7510	165	7345

Table A1: Results of Three Classification Models on DS1

Model	Error rate	Precision	Recall	F-measure	Settings
SVM	0.02	0.988	0.98	0.98	MostFreqTerms=50
NB	0.04	0.97	0.978	0.97	MostFreqTerms=50
DT	0.02	0.99	0.977	0.98	MostFreqTerms=50
SVM	0.02	0.992	0.98	0.985	MostFreqTerms=10
NB	0.04	0.96	0.98	0.969	MostFreqTerms=10

Table A2: Results of Three Classification Models on DS2

Model	Error rate	Precision	Recall	F-measure	Settings
NB	0.05926352	0.9776643	0.9575	0.95	MostFreqTerms=10
SVM	0.04833142	0.9968092	0.9518586	0.96	MostFreqTerms=10

Table A3: Classification of the Virgin dataset

Human Classification		NB		SVM		Settings
Positive	Negative	Positive	Negative	Positive	Negative	
2943	153	2278	818	1970	1126	MostFreqTerms=10
2943	153	2382	714	2032	1064	MostFreqTerms=5
2943	153	2723	373	2051	1045	MostFreqTerms=1