CORA Cork Open Research Archive
Cartlann Taighde Oscailte Chorcaí

Title	Security information sharing via Twitter: 'Heartbleed' as a case study
Authors	Briggs, Pam;Coventry, Lynne;Jeske, Debora;McNeill, Andrew R.
Publication date	2017-06-08
Original Citation	Jeske, D., McNeill, A. R., Coventry, L. and Briggs, P. (2017) 'Security information sharing via Twitter:'Heartbleed'as a case study', International Journal of Web Based Communities, 13(2), pp.172-192. https://doi.org/10.1504/ijwbc.2017.10005189
Type of publication	Article (peer-reviewed)
Link to publisher's version	10.1504/ijwbc.2017.10005189
Rights	© 2017, Inderscience Enterprises Ltd.
Download date	2026-10-25 23:14:56
Item downloaded from	https://hdl.handle.net/10468/16164

UCC
University College Cork, Ireland
Coláiste na hOllscoile Corcaigh

Security information sharing via Twitter: 'Heartbleed' as a case study

Debora Jeske*

Edinburgh Napier University,
219 Colinton Road, Edinburgh EH14 1DJ, UK
Email: d.jeske@ucc.ie
*Corresponding author

Andrew R. McNeill, Lynne Coventry and Pam Briggs

Psychology and Communication Lab,
Northumbria University,
Northumberland Road,
Newcastle upon Tyne, NE1 8ST, UK
Email: andrew.mcneill@northumbria.ac.uk
Email: lynne.coventry@northumbria.ac.uk
Email: p.briggs@northumbria.ac.uk

Abstract: The current paper outlines an exploratory case study in which we examined the extent to which specific communities of Twitter users engaged with the debate about the security threat known as 'Heartbleed' in the first few days after this threat was exposed. The case study explored which professional groups appeared to lead the debate about Heartbleed, the nature of the communication (tweets and retweets), and evidence about behaviour change. Using keywords from the Twitter user profiles, six occupational groups were identified, each of which were likely to have a direct interest in learning about Heartbleed (including legal, financial, entrepreneurial, press, and IT professionals). The groups participated to different degrees in the debate about Heartbleed. This exploratory case study provides an insight into information sharing, potential communities of influence, and points for future research in the absence of a voice of authority in the field of cybersecurity.

Keywords: Heartbleed; tweet content; influence; behavioural change.

Reference to this paper should be made as follows: Jeske, D., McNeill, A.R., Coventry, L. and Briggs, P. (2017) 'Security information sharing via Twitter: 'heartbleed' as a case study', *Int. J. Web-Based Communities*, Vol. 13, No. 2, pp.172–192.

Biographical notes: Debora Jeske is a Lecturer in Work and Organisational Psychology at University College Cork, Republic of Ireland. The current paper was submitted during her time at Edinburgh Napier University, Scotland, UK. Her research interests include technology and psychology at work.

Andrew R. McNeill is a Senior Research Assistant working at PaCT Lab, Northumbria University, Newcastle upon Tyne. He is currently working on the ACANTO project, part of which involves designing social network technology for older adults. His research focuses on social media and user-centred design.

Lynne Coventry is the Director of PaCT Lab (Psychology and Communication Technology) at the Northumbria University. Her research interests include exploring the role of communication technology in the lives of older adults, trust of online information, and the design of usable security.

Pam Briggs holds a Chair in Applied Psychology (PaCT Lab, Northumbria University), delivering innovative research and consultancy around issues of identity, trust and security in new social media. Her research interests include understanding personal information disclosure, trust in electronic media, and online privacy.

1 Introduction

Twitter, like many other social media, is a frequently used social medium in communication campaigns, particularly in the area of health (e.g., Morris, 2011). The advantage of social media is that it allows rapid dissemination of official advice (e.g., in the form of tweets and retweets (RTs) on Twitter; see also Li and Li, 2014) to various stakeholders (e.g., Pang et al., 2014) via a large network of people who can continue to share the information, thus reaching a broad audience (i.e., a snowball effect). Such awareness campaigns have provided researchers with a wealth of data relating to a wide range of issues such as flu trends, vaccination attitudes and psychological wellbeing (Achrekar et al., 2011; Anger and Kittl, 2011; Byun et al., 2013; De Choudhury, 2014; Kim et al., 2009; Love et al., 2013; see also Vargo et al., 2015).

One area of interest in pandemic research is how people and organisations respond to crises in different ways (see also Pang et al., 2014; Miyabe et al., 2014; Romenti et al., 2014). One example in healthcare is the exploration of how anti-vaccination sentiment can have an impact on discussions and the spread of information through social media networks (e.g., Chew and Eysenbach, 2010). Similar approaches have been used to explore sentiments of customers about products, and access the customer's mind set (Misopoulos et al., 2014).

In order to better understand how and with whom information is shared, researchers may wish to ask several exploratory questions: which specific Twitter groups participate in debates when a crisis or incident becomes public knowledge? Which groups participate and contribute more than others (creating potential disseminators and opinion shapers)? Does the Twitter data provide evidence that users respond to the crisis or incident by changing their behaviours? Some of these questions may already be answered based on existing research. In terms of consequences, the views and behaviours that people will adopt are likely to be in line with the normative beliefs of their social groups (see McNeill and Briggs, 2014; Mols et al., 2014). Such groups (e.g., parents, health professionals or alternative-health advocates) share a common identity and draw on the advice of each other because of within-group trust and willingness to adhere to group norms – also known as 'communities of influence' (Zhou, 2011; see further work on 'users of influence' by Räbiger and Spiliopoulou, 2015). Analysis of these communities can be conducted in a variety of ways (see Ediger et al., 2010; Pagoto et al., 2013; Zappavigna, 2014), each of which seek to explore the role of different stakeholders (groups) on social media.

Cybersecurity threats could be considered as cyber pandemics that affect a large number of people in often unknown ways, creating uncertainty and a search for information and advice. Researching and improving the effectiveness of communication in security 'pandemics' becomes more important as such pandemics are likely to reoccur in the future. The 'Heartbleed' bug in OpenSSL was selected as a case study of a cybersecurity pandemic as it involved many different stakeholders and communities of influence. This bug allowed 'man in the middle' attacks to access memory leaks between clients and servers (see Callegati et al., 2009). Some important differences exist between pandemic communications in health versus security. In the health sector, health authorities offer terminology and advice which is often picked up by the wider public (Chew and Eysenbach, 2010). Communication in the area of cybersecurity is less structured, not necessarily top-down and driven by a multiplicity of sources along with peer information sharing (e.g., via Twitter and other social media).

Several frameworks exist that help us understand the various elements involved in shaping behaviour in response to an event such as Heartbleed. We focus on the Theory of Planned Behaviour (Ajzen, 1991) to better understand the dynamics in the case of Heartbleed. This theory is useful in this context because it accounts for different influencers on end-user decision making. The Theory of Planned Behaviour suggests that attitudes and norms (which we propose are captured by tweets) are precursors to intentions and actions. We outline the components and relationships between them in Figure 1.

Figure 1 The components of the Theory of Planned Behaviour

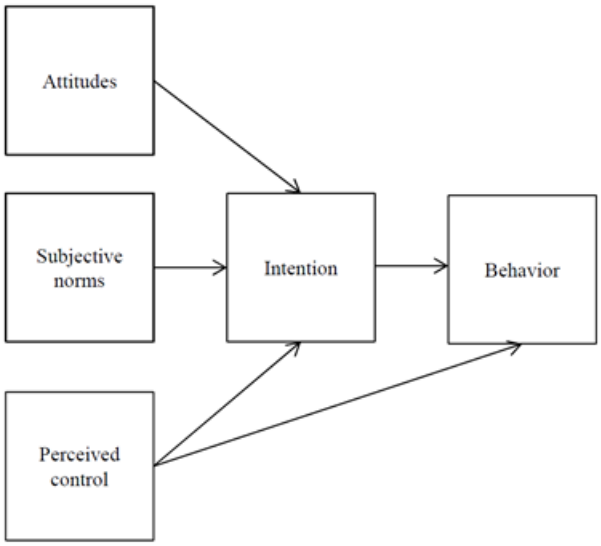

According to this theory, attitudes are expectations or behavioural beliefs about specific outcomes, while subjective norms are shaped by perceived social pressure (individuals generally judge their behaviour against their peers – those who are active on Twitter may feel more inclined to follow the recommendations of similar others), both are expected to shape individuals' intentions, actions and decision making to perform a behaviour (Ajzen, 1991). A third factor, perceived control, relates to the extent to which people

believe they can effect change. Those with little perceived control (which may be more prominent for certain Twitter users than others) may struggle to form and execute an implementation intention (Glanz et al., 2008).

While Twitter data cannot be used to show behaviour change per se (except when the tweets clearly reference behaviours such as password change), this data can be used to reveal the way attitudes and 'norms' differ across the various communities and show how these are associated with specific behavioural intentions. Perceived control may have been a particularly important element in the case of Heartbleed. In the first few days following the first public announcement, it was unclear which advice users should follow. As a result, both experts and non-experts participated in the debate (see also Kelion, 2014). Perceptions of control were compromised in this situation as individuals were essentially unable to act on any advice or reduce the risks to their data until the servers used by their various service providers (e.g., financial, health, educational, social media and governmental institution) had been patched first. This uncertainty made them dependent on the advice and behaviour of others (subjective norms), informed their attitudes about the risk to themselves, and influenced their control perception of whether or not they could secure their data by, for example, changing their passwords.

1.1 Exploratory research questions

This study focused on communications about the Heartbleed cybersecurity incident via Twitter. Twitter data was captured from April to May 2014. As there is very little known about web-based communities of influence in relation to cybersecurity incidents, we proposed the following exploratory research questions (RQ's) to assess whether or not certain groups can be identified in relation to this phenomenon:

RQ 1 Which specific Twitter groups (communities of influence) can be identified amongst the Twitter users? (Who are the participants?)

RQ 2 How and to what degree did these groups participate in the debate about Heartbleed?

RQ 3 Which groups tweet at a similar or different rate across the first few weeks of the Heartbleed debate?

RQ 4 What was the focus (content) of the tweets?

RQ 5 What evidence is there that users discuss actual behaviour change (in this case, password change)?

The next two sections include the method and the results section. The method section outlines the procedure and details regarding data collection. The headings of the results section are organised in line with the research questions:

1 Identification of Twitter groups (participants)

2 Proportion of tweets contributed to the debate

3 Retweet (RT) behaviour

4 Group differences in tweet content – keyword analysis

5 Behavioural change: using the example of password change.

The discussion follows the results and includes a number of areas for future research.

2 Methods

2.1 Data collection approach

In order to learn more about who was involved in the information exchange about Heartbleed, Twitter data was collected for the first few weeks following the first news release about Heartbleed (8 April 8 to 13 May 2014). This time period marked the beginning of the Heartbleed debate (the bug fix was published first on 7 April, see Heartbleed.com, 2014; Nieva, 2014). Data was collected with a Python script using the Twitter public API stream and was subject to limits imposed by Twitter. Tweets were gathered by using the Twitter Streaming API to search for tweets containing the keyword 'Heartbleed'. The final dataset (also considering the streaming limits imposed) included around 100,000 tweets (including RTs).

2.2 General participant description

Once the duplicates had been removed, the dataset included 91,414 tweets from a randomised, English speaking sample. A review of a third of the tweets revealed participants came from more than 100 different countries. This suggests that the tweets captured the views of individuals from around the globe. This study utilised a multi-method approach, using both qualitative and quantitative approaches (see also Misopoulos et al., 2014) to analyse sentiments as well as frequencies of specific trends occurring in the data. The methods used to answer each research question are described in the appropriate sections below. The exact affiliation of participants (and thus Twitter groups) was identified as part of the first research question. Details are outlined in the results Section 3.1.

3 Results

Results are organised into several sections, each addressing one of the exploratory research questions. These analyses focused on the total number of tweets in each group, rather than the total number of tweeters. The decision to focus on the tweets rather than the tweeters was based on two aspects: the focus was to explore how different groups of users rather than individual tweeters participated in the debate. Secondly, social norms may play a role in such groups, thus a focus on groups instead of individuals enabled the researchers to examine group-specific behaviours (such as recommendations of particular actions) rather than individual influencers (based on individual Twitter profiles).

3.1 RQ1 – Identification of Twitter groups (participants)

Past evidence suggests that communities with separate memberships exist within Twitter's network of users which may not necessarily be formal groups (see Ch'ng,

2015). In order to identify communities of influence, the researchers first determined which Twitter groups were most likely to be involved in the debate and/or directly affected by Heartbleed (e.g., insecurities in servers affect finance, legal, marketing professionals, IT and entrepreneurs particularly negatively, with the press having an important role in shaping communication). Knowledge sharing of critical information about Heartbleed in these communities may be influenced by potential benefits resulting from exchanging information, such as access to new information (see also work on rewards and altruism as motivators of knowledge sharing in Lin and Huang, 2013).

Once these occupational groups had been identified, and specific keywords were determined for each group, the profile of each Twitter user was coded accordingly. This approach was based on other work that also utilised information in Twitter profiles. For example, Twitter profile information about contacts has been used to make personality trait attributions (e.g., Quercia et al., 2011), to examine user contributions and contacts (see Zhang and Nasraoui, 2008), and to cluster individuals into groups or derive group profiles for use in simulated learning systems (see Ammari et al., 2012). Our goal was to identify similar groups using keywords in each profile (Fernandez et al., 2014; Sloan et al., 2015; for other methods see Mizzaro et al., 2015; Xie et al., 2014).

To identify these groups, the data was coded in several ways. This was done using R and a set of grouping keywords (the base software only, no extra package was used). R is a statistical programming language with the flexibility of being able to handle both textual and numerical data (see R Core Team, 2014). Tweets were allocated as belonging to users of a specific occupational group when the user profile matched one or more of the keywords (see also Sloan et al., 2015). Allocation was restricted to only one group. If a profile included more than one keyword from different groups (which was only the case for a dozen cases), the tweeter was allocated to the group that matched the profile best.

The keywords were identified in discussion between researchers. These included:

1 IT professionals: keywords ‘IT professional’, ‘computer’, and/or ‘security’ in profile

2 Legal professionals: keywords ‘law’, ‘barrister’, ‘solicitor’ and/or ‘legal’ in profile

3 Finance professionals: keywords ‘bank’, ‘finance’, ‘insurance’

4 Self-employed professionals: keywords ‘entrepreneur’, ‘self-employed’, ‘intellectual property’, and/or ‘own company’ in profile

5 Press professionals: keywords ‘press’, ‘journalist’, ‘writer’, ‘news’, ‘publisher’ and/or ‘blog’ in profile (please note that only .28% of classifications were based on ‘blog’ alone)

6 Those involved in marketing, by selecting them based on the ‘marketing’ phrases in their profiles.

This process allowed a definite group allocation for the users of a total of 44,673 tweets (48.1%). All remaining tweets were included in a group of ‘unclassified’ tweeters (0). These cases featured empty user profiles or profile descriptions that did not match any of our key terms for the other groups (Table 1).

Table 1 Presence of professional groups most likely affected by Heartbleed in present sample

Groups	*n*	*% of N*
IT professional	21,523	23.5
Legal	997	1.1
Finance	1,459	1.6
Entrepreneur	2,488	2.7
Press	16,904	18.5
Marketing	1,302	1.4
Unclassified	46,741	51.1
Total	*91,414*	100.0

The remaining 51.9% of tweets were not matched to any of these groups. This was due to several reasons. First, some did not match any of the keywords used to identify the six occupational groups that were most likely to have been affected by and thus interested in a solution for Heartbleed. This means they may also not have the same stake in the debate. Second, some tweeters in the unclassified group did not include sufficient information for categorisation.

3.2 RQ2 – Proportion of tweets contributed to the debate

Another important element of our analysis was to examine the extent to which different users engaged in the debate across time, using the proportion of tweets from different groups. The participation rates (capturing the tweets from the occupational groups) were converted into percentages (using the subtotal for each group to assess the percentage of the group that participated in the debate on given days). Figure 2 shows the percentage of Twitter users in each group involved in the discussion over the course of each week.

Figure 2 Percentage of groups tweeting about Heartbleed across five weeks (see online version for colours)

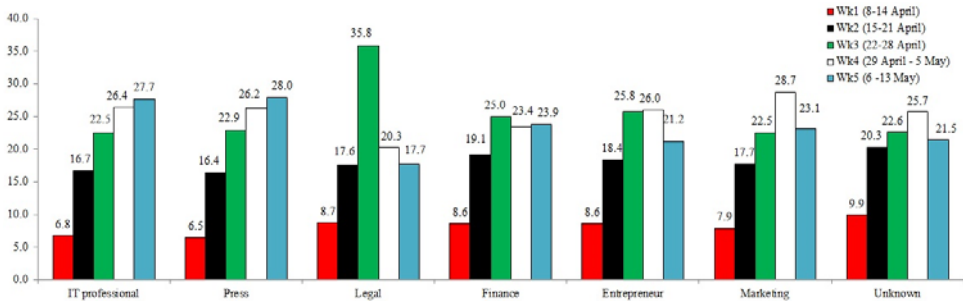

The analysis was conducted using the chi square test of independence as this enables researchers to compare how observed data (tweets per group) relates to what would have been expected by chance (i.e., in proportion to their size). For example, if each of the occupational groups contributed to the debate to the same degree in each week, the proportion of tweets observed should be equal across each of the five weeks (and there should be no significant difference between the observed count of tweets and the

expected number of tweets). In the current study, some occupational groups were expected to be more vocal in some five week than others as the problem and repercussions of Heartbleed became more known over time. The results of the chi square analyses suggested that the proportions of tweets originating from each occupational group are not equal across the five weeks. The frequencies showing participation patterns are outlined in Figure 2.

The patterns for each of the five weeks suggest that the groups participated to different degrees in the debate (as reflected in the number of tweets) as time went on – assuming, of course, that our tweets were collected from relatively homogeneous and comparative samples each week. The results suggest the following. First, both IT professionals and the press group tweeted more and more over the course of the five weeks, while other groups became more vocal at particular but select points over the five weeks only. IT professionals in our dataset increased their observed vs. expected tweets steadily across all five weeks ($\chi^2(4) = 1{,}096.73$, $p < .001$). A similar tendency could be seen in the press group ($\chi^2(4) = 278.25$, $p < .001$). Second, the close relationship between tweeting counts observed for the IT and press suggests that the press may simply be mirroring the activity of the IT groups. Third, the tweet pattern also appeared to differ across the five weeks for the legal group ($\chi^2(4) = 79.05$, $p < .001$), the financial group ($\chi^2(4) = 7.85$, $p = .097$), and the entrepreneurs ($\chi^2(4) = 13.663$, $p = .008$). No significant differences in tweet patterns were observed for the marketing professionals in terms of their tweet rate across all five weeks ($\chi^2(4) = 5.710$, $p = .222$). These different groups are likely to have had different perspectives on the crisis and how to react and the discussion about appropriate responses is likely to have created norms among each group.

3.3 RQ3 – RT behaviour

The RT frequencies observed for different occupational groups were examined next. Such frequencies can help establish which groups appeared to share information more widely. The focus was on which Twitter users responded to, rather than initiated, a tweet (see also research on credibility and favourability of communal versus exchange tweets by Li and Li, 2014). The analysis was conducted with all identifiable user groups, excluding the unclassified group (IT professional, legal, finance, entrepreneur, press, and marketing, $n = 44{,}673$). All RTs (12.2%; total RTs $n = 11{,}150$) were identified using R (base packages).

Table 2 shows the observed and expected values for both tweets and RTs. The actual tweets are listed in the first column to provide a baseline against which to compare the number of RTs within the groups. The second column lists the RTs in each group. Chi square analysis results indicate a difference among different professional groups in their tweeting and retweeting patterns ($\chi^2(5) = 1{,}761.98$, $p < .001$). An example shows how to interpret the contents of Table 2. The frequencies for IT professionals suggest that the number of observed unique (i.e., non-retweeted) tweets was lower than would have been expected statistically by chance (obs./exp. = 14,959/15,605), while the number of observed RTs was higher than the expected number (obs./exp. = 6,838/6,178). This means IT professionals had issued fewer unique tweets compared to what would have been expected for this group. At the same time, they had actually retweeted content more often. Perceived tweet and RT patterns were also examined statistically, in line with the descriptive outlined for the groups in Table 2. More information about the patterns is

provided in the notes for those readers interested in the statistical difference in frequencies.

When comparing tweet against RT for each group, the IT group ($\chi^2(1) = 129.33$, $p < .001$) retweeted materials significantly more often than sending original tweets. The same was observed for the legal group ($\chi^2(1) = 28.07$, $p < .001$). The press send out more original tweets than RTs ($\chi^2(1) = 2{,}250.42$, $p < .001$), similar to marketing ($\chi^2(1) = 28.57$, $p < .001$). In other words, the press and marketing groups appeared to create content more rather than simply sharing it like other groups. The finance group ($\chi^2(1) = .71$, $p = .400$) and the entrepreneurial group ($\chi^2(1) = 2.65$, $p = .104$) were, however, just as likely to tweet or RT.

Table 2 Tweets and non-tweet frequencies

Groups	*Tweet frequencies*			*RT frequencies*		
	% in each group	*Obs (exp)*	*(obs-exp)^2 exp*	*% in each group*	*obs (exp)*	*(obs-exp)^2 exp*
IT profess.	69.5%	14,959 (15,605)	26.74	31.8%	6,838 (6,178)	70.51
Legal	68.0%	678 (755)	7.85	37.7%	376 (299)	19.83
Finance	76.6%	1,118 (1,103)	.20	28.9%	422 (437)	0.513
Entrepreneur	72.0%	1,791 (1,827)	.71	30.5%	760 (723)	1.893
Press	90.0%	15,210 (12,657)	514.96	14.5%	2,458 (5,011)	1,300.703
Marketing	81.2%	1,058 (970)	7.98	22.7%	296 (384)	20.17
∑ (tweets)		34,814 (32,917)	558.449		11,150 (13,032)	1,413.61

Notes: Note that df(5) → $F(558.4451/5) = 111.689$ → 1–111.689 = –110, which is smaller than .05 and subsequently there is a difference in tweet frequency between groups. In terms of tweeting patterns in each group, we conducted a number of additional analyses: using the equation $\frac{(observed - expected)^2}{\exp}$ we obtain 558.449***. To compute the χ^2-value the expected tweets are distracted from the observed tweets divided by the observed tweets. Values are squared so that the sum is unequal 0. The sum of the column represents the χ^2-value for the column. Subsequently there is an in-between group difference as the descriptive values already indicate. Since our χ^2-value is 558.45 and subsequently greater than χ^2crit = 11.07, the null hypothesis is rejected, the groups differ in their tweeting behaviour. Using the probability of success in a Bernoulli experiment significance can also be computed cell wise (Miettinen and Nurminen, 1985; Koutras et al., 2006). Each group significantly differs from the tweeting behaviour that would have been expected. For example, using odd-ratio to describe the difference in tweeting behaviour we find that IT professionals (14,959/34,814 = 0.4296835) are 22 times as likely as legal (678/34,814) to tweet about Heartbleed (0.43/0.02 = 21.5). Please note that the results remain highly significant even when controlling for alpha-value inflation using the false discovery rate (Benjamini and Hochberg, 1995). In terms of retweeting patterns, the same analysis was conducted. Since categories have not changed the critical Chi-Square value remains χ^2crit = 11.07. Subsequently the null hypothesis can be rejected; the groups differ in their retweeting behaviour. Using the probability of success in a Bernoulli experiment it was found that the groups significantly deviate from the retweeting behaviour expected.

3.4 RQ4 – Group differences in tweet content – keyword analysis

Expanding on the previous analyses, R was used to identify the main keywords occurring in tweets in each of the six groups (including unique tweets as well as RTs). The keyword frequency provides a means to assess the most prominent themes in tweets. As expected, the content of tweets generally listed technology-related words (including web, internet, InfoSec, update, tech, source, software, threat, servers, risk, protect and risk). Some keywords were specific to Heartbleed (bug, VPN, security, password, data, open SSL, cybersecurity and hackers). In addition, many of the groups referred to organisations or products (apps, android, Amazon, Facebook, Apple, Microsoft, and NSA). Authentication was also a topic discussed by the individuals in the IT, entrepreneurs and finance group. Patches were discussed by the IT, marketing and press group. The service 'logmeonce' was predominantly mentioned by marketing, entrepreneurs and finance Twitter users. 'logmeonce' is 'a multi-platform security management solution supporting PC, Mac, Android, iPhone, and iPad' (logmeonce.com) Finally, both marketing and finance discussed Heartbleed in relation to mobile technology.

Some keywords were also group-specific. For IT, these included banking, code and Linux. In the legal group, tweets made references to keywords such as government, power, impact, law and privacy. The finance group discussed tax, filing, takes and tips. The entrepreneurial group discussed ideas, initiative and start-ups. The press tweeted about networks, targets, Windows and switches (e.g., services). Lastly, the marketing group was concerned with companies, consumers, encryption, and (social) media. These findings suggest that while the debate about Heartbleed focused on many of the same issues, the groups also had specific concerns that were unique to their industry/sector and expertise.

Figure 3 Network graph of all RTs (see online version for colours)

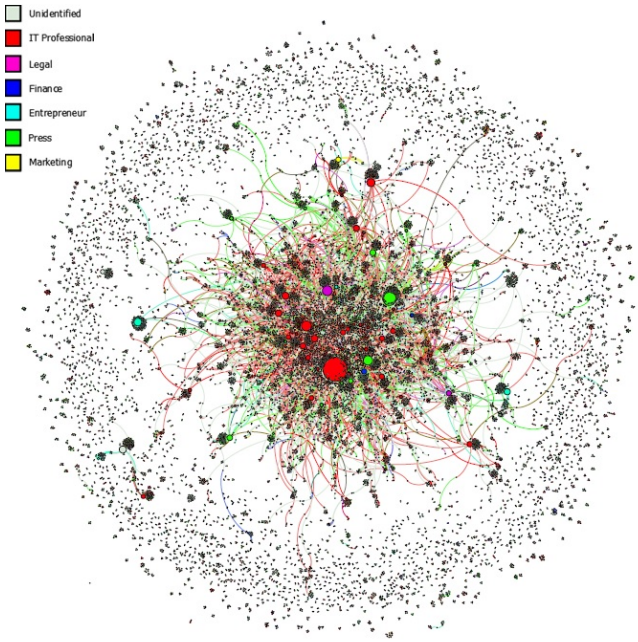

The results were verified using various visualisation tools such as Gephi (Bastian et al., 2009), ForceAtlas2 (see Jacomy et al., 2014) and iGraph (Csardi and Nepusz, 2006). Some researchers have developed specific visualisation tools for Twitter (see Aragón et al., 2011). The focus of these analyses was on examining the relationship between users specifically, rather than the actual tweets. As anticipated and illustrated in Figure 3, the users of the 'unclassified group' were located at the periphery of the graph when visualising their influence in the debate (in other words, they were less influential). Separate analyses revealed that RTs were particularly prevalent for groups associated with IT and the press.

3.5 RQ5 – Behavioural change: using the example of password change

The last research question explored the extent to which users discussed changing passwords (implying potential behaviour change) In order to prepare the data; the entire dataset was recoded using several different methods. The work frequency of 'password' was explored first using NVivo 10 (Richards, 1999). NVivo is a program that enables analysis of unstructured data by providing various search, query and visualisation tools. Next, the exact phrases used by Twitter users were examined. These were coded as:

1 recommending password changes in general

2 indicating password change has been made, and finally

3 no password change or against password change.

Specifically, using NVivo and Excel key word searches, these identified phrases were then selected as search terms.

1 In terms of advice: 23 different word groupings were used to identify advice (n = 2,680). Examples include 'passwords you need to change', 'change your passwords' and 'passwords you should change'.

2 In terms of expressed behaviour change, seven search phrases were located, such as 'changed passwords' or 'changing all of my passwords' (n = 498).

3 Those tweets advocating no behaviour change were identified using seven word phrases such as 'have not changed my passwords' (n = 64).

Table 3 Discussion of passwords in tweets amongst different professional groups

Groups	*General advice (1)*	*Change indication (2)*	*No Change (3)*
IT professional (n = 21,523)	464 (2.15%)	115 (.53%)	16 (.07%)
Legal (n = 997)	37 (3.71%)	5 (.50%)	2 (.20%)
Finance (n = 1,459)	63 (4.32%)	8 (.55%)	1 (.07%)
Entrepreneur (n = 2,488)	86 (3.46%)	11 (.44%)	2 (.08%)
Press (n = 16,904)	387 (2.29%)	97 (.57%)	14 (.08%)
Marketing (n = 1,302)	50 (3.84%)	10 (.77%)	1 (.08%)
Unclassified (n = 46,741)	1,593 (3.41%)	252 (.54%)	28 (.06%)
Total	2,680	498	64

The word 'login' was used in other contexts than password change, so these tweets were not included in the analysis of tweet content. Counts for each coding are listed in Table 3. The total number of tweets coded for in this analysis was 3,242 (3.54% of tweets). The low count is perhaps indicative of:

a ambivalence over what to do

b the predominantly information-sharing nature of Twitter rather than a forum to declare personal intentions/behaviours (cf., Chew and Eysenbach, 2010).

Despite this, the level of behaviourally-relevant information is high enough in absolute (rather than relative) terms to be of interest and thus can provide insight into behavioural intentions.

When examining which tweets (associated with specific occupational groups and including the unclassified group as well, $n = 7$) advocated any of the three positions, only a statistically significant group effect was observed for who recommended password change in general ($\chi^2(6) = 125.70$, $p < .001$, $N = 91{,}391$). Tweets from the legal (obs./exp. = 37/29), finance (obs./exp. = 63/43), entrepreneurial groups (obs./exp. = 86/73) and the users in the unclassified group (obs./exp. = 1,593/1,370) promoted password change more heavily. In contrast, IT professionals (obs./exp. = 464/631) and the press promoted password change less often than expected (obs./exp. = 387/495).

Several explanations may be offered for why password change was not as much of a topic as expected amongst IT professionals and the press. On the one hand, the IT professionals were aware that changing the passwords was not the immediate solution. First, the bug had to be fixed on the servers. Second, once the servers had been patched, only then should the users change their password. This may explain the lower than expected frequency observed in terms of tweets advising password changes. In addition, by recent accounts up to 300,000 servers still remain at risk from Heartbleed (see Hamilton, 2014). As a result, IT professionals may have been more aware of Heartbleed being a security risk.

In response to this, IT professionals did not appear to take the lead in the discussion of password change although they would have been the more knowledgeable Twitter users to educate others. They appeared to focus on disseminating new information – without recommending behavioural change to the same degree as observed in the groups identified as financial, legal and entrepreneurial. The latter groups may have had a greater vested interest in educating others about the risks due to their potential role as service providers (in banks, legal settings, and business). They may have been required to be seen to be doing something and so reduce the perception of risk by increasing self-efficacy through the prescription of a definite action that could be performed to maintain user safety. That said there is no guarantee that the advice would be appropriate.

4 Discussion

The wide use of technological means to save, transmit and share data has created numerous new opportunities for new threats to arise; threats that may result in concerns and vulnerabilities for many thousands or even millions of people. Exploring how people respond to crises is an important concern in pandemic research, similar to other crisis-focused research in other areas than health or security (see also Miyabe et al., 2014;

Pang et al., 2014; Romenti et al., 2014). Cybersecurity threats represent 'cyber pandemics.' In this paper, Heartbleed was examined as a case study of a security 'pandemic', a situation similar to a health pandemic in which a need for action and further information arises.

In this exploratory analysis of Twitter data, several different occupational groups were engaged in the debate about Heartbleed (RQ1), in the absence of any clear national or international voice of authority (RQ2). Additional research questions focused on the activity of these groups over the first few weeks (RQ3), the content of RTs (RQ4), and any evidence of Twitter users responding to the threat by discussing and changing their behaviours (specifically passwords, RQ5). A summary of findings is presented next.

The debate as captured in our dataset did not appear to be led by the most knowledgeable group (the IT professionals). That is, going on their activity while also controlling for the number of IT professionals identified and the number of their tweets captured in the dataset, their activity in terms of tweets is relatively low or on par with only the legal group in terms of the proportion of tweets and RTs issued by the members of the other groups (Table 2). The press and marketing group produced – given the group size – more tweets as a group about Heartbleed than the IT group.

These trends may be interpreted in two possible ways. While it is entirely possible that the IT group retweeted content could be viewed as evidence that they played a supporting role in the debate (see also work on message strategies by Li and Li, 2014). However, another interpretation focuses on expertise rather than leadership in this debate. That is, it is also possible that the Twitter users captured in the IT group may have realised that the potential for peer-to-peer problem solving (see also work by Chen et al., 2013) was not an option in the case of Heartbleed. This may have shaped how they decided to share knowledge. Future research into similar cyber-incidents may provide more clarity about the motivations, attitudes, and knowledge contribution strategies of IT professionals depending on the nature of the cybersecurity incident.

The impact of messages from the IT group may have been significantly bolstered by the press group whose tweeting patterns closely mirrored those of the IT group. Visualisation of RTs also demonstrated that the majority of information sharing in the form of RTs was from IT professionals and the press group that functioned as sources of information for other users. These two groups may have felt more involved in the issue (see also Wang et al., 2012) and hence more willing to contribute their knowledge and resources to sharing information with the online community. Nevertheless, the results suggest that no particular group emerged as a voice of influence in the debate captured in our dataset.

Further quantitative analyses were conducted with tweets to examine the extent to which password change was discussed by users. The results suggested that changing one's password was not a common topic amongst the tweets that had been collected (3.54% of tweets made references to passwords). IT professionals were significantly less likely to promote password change, probably because they were aware that this would not fix the problem overall. Instead, legal, finance and entrepreneurs discussed or recommended password change more often. In the case of Heartbleed, the average user had no control over the situation (only if the site is fixed would it make sense to change one's password).

Our exploratory results support three conclusions. First, according to the data, there was no evidence that any particular group took on the role of authority. Secondly, the tendency of less knowledgeable groups to give advice on behaviour (promote password

change when this was not the first step to a solution) suggests that establishing a voice of authority, representing the most knowledgeable experts, may be helpful to raise awareness and suggest action (see recent work on identifying authoritative actors by online communities by Bouguessa and Romdhane, 2015). And third, given the connections that marketing and entrepreneurs appear to foster (at least in this paper's dataset), these may also be suitable target groups to disseminate information, in addition to the press.

4.1 Theory reflections

The Theory of Planned Behaviour (Ajzen, 1991) recognises the role of attitudes, subjective norms and perceived control in predicting intentions, and subsequently, behaviour change. Attitudes played out in the communication of risk, with some tweets suggested very high levels of anxiety around the kinds of impact that Heartbleed may have. Perceived control and self-efficacy may also have been critical factors. Future research with Twitter users may be able to better determine the link between affect in tweets published in response to a critical incident and users stated beliefs about their ability to deal with the critical incident. The relative lack of password change advice from IT professionals and the press may have only added to this problem as there was little available or consistent guidance that addressed the way that individuals should take control.

It is important to acknowledge that tweets may also be shaped by the presence of influential others on Twitter. Research suggests that when 'important others' on social media networks recommend actions on privacy, people are more likely to follow their advice (Saeri et al., 2014). Future research may wish to explore whether or not the presence of 'favourite' contacts (important others in the contact list) can influence the formation of attitudes and norms on new security issues. Such important others can be identified using eigenvector centrality measures, which measure the importance of a node (user) based on the number of important nodes (users) that link to that node (user). Users with higher centrality measures would then be expected to exert greater influence on the attitudes, norms and perceived control of those who follow that user. In addition, this research on social influence may also consider the influence of specific groups on followers when the group's perceived expertise appears to be particularly important (e.g., IT expertise in the case of cybersecurity incidents).

4.2 Contributions to research

This research on influence makes three contributions to the literature. First, the paper provides an example of a recent case study in which the responsibility for informing the public was less clear, compared to regular (health) pandemics. In health pandemics, the drivers of information are news sources and health authorities (Love et al., 2013). At present it is not clear who drives information in the context of a security threat. Heartbleed represented a unique situation: uncertainty about the best course of action was high (particularly at the beginning), as was the ability of the individual Twitter user to affect change or improve the situation on their own. This uncertainty is common of security alerts where it is often unclear which course of action should be advocated.

Second, the present findings also add to the work on message propagation and social influence (Anger and Kittl, 2011; Ye and Wu, 2010). For example, Ye and Wu (2010)

examined message propagation following the death of a celebrity, focusing on stabilities, assessments, and correlations of social influence in their data. They proposed that social influence is wielded via followers (more followers suggest more influence on others), reply influence (greater replies received by a user suggests greater influence) and RT influence (with greater retweeting suggesting greater influence). Similar indicators were utilised in the present dataset.

Third, this work adds to the literature on Twitter influence in cybersecurity, another domain similar to health that usually requires a rapid response to an emergency state. This paper outlined important communication patterns around Heartbleed, specifically how and to what extent different occupational groups would tweet and contribute to the same degree week in and week out to the debate about Heartbleed. Exploring this is important as information shared on Twitter is not equally influential. Some users are very influential and are widely followed and shared whereas others simply comment on the ideas of others (Tinati et al., 2012). This analysis enables us to identify potentially useful disseminator groups and provides an insight into the dynamics of how different groups communicate with each other.

Finally, this paper makes a contribution to the work on risk communication (see work by Burns and Slovic, 2012). One area of interest in this area of risk and uncertainty management focuses on how to improve communication strategies to reach as many potentially affected individuals as possible. Identifying different groups' responses to crises can reveal normative behaviours within each group and targeting these groups separately can lead to more persuasive appeals for behaviour change (Mols et al., 2014). The current results suggest that in the absence of a voice of authority, it is important to recognise which other groups may lead, shape and influence the debate. In addition, by observing and evaluating content of messages on a topic over time, it may be possible to study longitudinally how different groups respond to threat revelations, how perceptions change, and the extent to which risk-reducing strategies can be implemented with the help of the most trusted but also best connected communicators (Burns and Slovic, 2012).

4.3 Limitations of exploratory study

As is the case in many exploratory studies, a number of limitations are worth noting. First of all, we used very simple features to explore how Twitter users make sense of a security event. We employed a restrictive filtering approach, based on a single keyword. This may not be suitable for a 'representative' study of information sharing on Heartbleed as many tweets may also have discussed this incident without referring specifically to the precise keyword 'Heartbleed'. However, this would be a criticism that can be employed to many Twitter searches that rely on specific keywords.

Secondly, it is possible that the sample of tweets was biased towards one group due to the sampling method that was used. One alternative approach to using a top-down method (selecting classification codes) would have been to extract key words from data, and compare the results of the auto-classification subsequently with a small sample of manually coded groups to get the accuracy rate. However, in this case we believe that the classification rate of 48.1% using a limited number of keywords was promising.

The decision to focus on specific occupational groups (IT, press, legal, financial, entrepreneurial, and marketing) was based on an a priori decision to consider all those with the highest stake in learning about Heartbleed as they may be more directly impacted. Of course, this list may not have been sufficiently exhaustive (e.g., we did not

include very specific keywords such as 'developer' or 'software engineer' in the hope that the more general labels would suffice). However, it is also worthwhile to keep in mind that every Twitter user could comment on Heartbleed. Our interest was in certain key interest groups, not to identify every single group. Using specific user information (profile, tags, or comment based) to identify user groups sharing a common denominator (often a common interest, topics, emotions, or health conditions) is a shared element in news recommendations (see also Jonnalagedda and Gauch, 2013; Li et al., 2010) and e-health initiatives involving personalised health profiling (see Batool and Khan, 2012). Another approach to analysing data is to look at individual influencers. However, this requires data that captures the activities of the same number of tweeters over time (this was not an option in the current dataset). Future research may wish to explore this approach in combination with professional grouping.

In addition, behavioural change could only be inferred from the content of tweets. Twitter users may use tweets to communicate, receive and exchange information – but not necessarily report on personal behaviour to the same degree that these users might on other, alternative social platforms. Furthermore, it was not possible to test whether password change intentions were driven by Twitter (or were derived from an external source) due to the ostensible expiry of the shortened URLs collected.

The question of whether Twitter is predominantly a social network or a news medium has already been raised by others (e.g., Anger and Kittl, 2011; Kwak et al., 2010). This suggests that the utility of using Twitter data to examine behavioural change is an issue that certainly warrants further attention. It has also been suggested that Twitter may simply be a self-affirmation device. That is, Twitter may be a place where users can post their thoughts and reported behaviours to gain the approval of their peers. This may certainly be the case for more active Twitter users (see also Wang et al., 2012). In which case, it would be an effective mechanism for establishing normative attitudes. This also suggests that there are communities of influence on Twitter who share attitudes in accordance with what they perceive to be group norms. Such a perception would presumably be gained from other posts on Twitter.

4.4 Future research

There are a number of other avenues that deserve more attention in the future. To date there is little research that looks at public responses to cybersecurity threats on a large scale, which is one of the reasons for the relatively small number of studies that are cited in this paper. Several areas for future research could be considered. First, it may be useful to examine the degree to which Twitter influences behavioural change in situations where individual Twitter users can exert some kind of control over their security or health (a limitation in Heartbleed).

Second, the utility of analysing threat discussions over time with the help of Twitter and other data sources should be explored further and extended to consider the role of emotion; in addition to the cognitive and behavioural content in messages (see also Slovic et al., 2004; Larsen et al., 2015). It was not possible to examine the persuasiveness of certain sentiments and to explore reactions to emotional content in this study. The discussion of perceived or actual expressed feelings may be helpful in understanding why some users are more or less likely to act on advice (an example here would be Kwak et al., 2010). This may be particularly important in the area of cybersecurity where the

user often does not have the means to influence the situation, and acting on advice may require significant effort on behalf of the novice users.

A less technical Twitter debate may present a better opportunity to explore the role of affect. Several methods for content and sentiment analyses already exist (e.g., 'topic modelling', see Blei (2012); 'opinion mining' and 'sentiment analysis', Pang and Lee, 2008; see applications of other tools in Aragón et al., 2011; Ediger et al., 2010). Several such methods have already been implemented in libraries for languages such as Python and R. The examination of emotional content (see also Batool and Khan, 2012) in addition to a groups' potential to support information dissemination and behavioural change may increase the effectiveness of risk communication strategies. Messages could be potentially tailored to different groups of respondents and stakeholders (Pang et al., 2014; see also the literature review by Veil et al. (2011); this work discusses the use of social media in risk and crisis communication). Developing such guidelines to support risk-reducing strategies is one of our future research goals.

4.5 Conclusions

The findings from our case study analysis of Heartbleed may provide starting points for how to identify Twitter groups that may need to be targeted in future communication campaigns aimed at dispersing information about security. Further work in this area may help to pinpoint differences in the discourse of topics related to cybersecurity or health. This may subsequently present a useful means to develop more tailored or personalised interventions to shape risk communication and strategies in the health and cybersecurity domain.

Acknowledgements

The authors would like to thank Joy Backhaus and the anonymous reviewers for their support and feedback.

References

Achrekar, H., Gandhe, A., Lazarus, R., Yu, S-H. and Liu, B. (2011) 'Predicting flu trends using Twitter data', *Proceedings of the IEEE Conference Computer Communications Workshops (INFOCOM)*, IEEE Xplore Digital Library, Shanghai, China, 1–15 April 2007, pp.702–707.

Ajzen, I. (1991) 'The theory of planned behavior', *Organizational Behavior and Human Decision Processes*, Vol. 50, No. 2, pp.179–211.

Ammari, A., Lau, L. and Dimitrova, V. (2012) 'Deriving group profiles from social media to facilitate the design of simulated environments for learning', *Proceedings on 2nd International Conference on Learning Analytics and Knowledge (LAK)*, ACM Digital Library, Vancouver, BC, Canada, 29 April to 2 May, pp.198–207.

Anger, I. and Kittl, C. (2011) 'Measuring influence on Twitter', *Proceedings of the 11th International Conference on Knowledge Management and Knowledge Technologies (i-KNOW)*, Graz, Austria, 7–9 September, No. 31, pp.1–4 [online] http://www.l2f.inesc-id.pt/~fmmb/wiki/uploads/Work/misnis.ref07.pdf (accessed 12 May 2015).

Aragón, P., García, Í. and García, A. (2011) 'Graph visualization tool for Twittersphere users based on a high-scalable extract, transform and load system', *Proceedings of the International Conference on International Conference*, ACM Digital Library, Sogndal, Norway, 25–27 May, No. 46, pp.1–4.

Bastian, M., Heymann, S. and Jacomy, M. (2009) 'Gephi: an open source software for exploring and manipulating networks', *Proceedings of the Third International AAAI Conference on Weblogs and Social Media*, AAAI Publications, San Jose, CA, USA, 17–20 May, pp.361–362 [online] https://www.aaai.org/ocs/index.php/ICWSM/09/paper/view/154 (accessed 12 May 2015).

Batool, R. and Khan, W. (2012) 'Towards personalized health profiling in social network', *Proceedings of the 6th International Conference on New Trends in Information Science, Service Science and Data Mining*, IEEE Xplore Digital Library, Taipei, Taiwan, 23–25 October, pp.760–765.

Benjamini, Y. and Hochberg, Y. (1995) 'Controlling the false discovery rate: a practical and powerful approach to multiple testing', *Journal of the Royal Statistical Society. Series B (Methodological)*, Vol. 57, No. 1, pp.289–300.

Blei, D.M. (2012) 'Probabilistic topic models', *Communications of the ACM*, Vol. 55, No. 4, pp.77–84.

Bouguessa, M. and Romdhane, L.B. (2015) 'Identifying authorities in online communities', *ACM Transactions on Intelligent Systems and Technology (TIST)*, Vol. 6, No. 3, pp.1–23.

Burns, W.J. and Slovic, P. (2012) 'Risk perception and behaviors: anticipating and responding to crises', *Risk Analysis*, Vol. 32, No. 4, pp. 579-582.

Byun, C., Lee, H., Kim, Y. and Kim, K.K. (2013) 'Twitter data collecting tool with rule-based filtering and analysis module', *International Journal of Web Information Systems*, Vol. 9, No. 3, pp.184–203.

Callegati, F., Cerroni, W. and Ramilli, M. (2009) 'Man-in-the-middle attack to the HTTPS protocol', *IEEE Security and Privacy*, Vol. 7, No. 1, pp.78–81.

Ch'ng, E. (2015) 'The bottom-up formation and maintenance of a Twitter community: analysis of the #FreeJahar Twitter community', *Industrial Management & Data Systems*, Vol. 115, No. 4, pp.612–624.

Chen, G-L., Yang, S-C. and Tang, S-M. (2013) 'Sense of virtual community and knowledge contribution in a P3 virtual community: motivation and experience', *Internet Research*, Vol. 23, No. 1, pp.4–26.

Chew, C. and Eysenbach, G. (2010) 'Pandemics in the age of Twitter: content analysis of Tweets during the 2009 H1N1 outbreak', *PloS ONE*, Vol. 5, No. 11, pp.1–13.

Csardi, G. and Nepusz, T. (2006) 'The igraph software package for complex network research', *InterJournal, Complex Systems*, Vol. 1695, No. 5, pp.1–9.

De Choudhury, M., Monroy-Hernández, A. and Mark, G. (2014) 'Narco emotions: affect and desensitization in social media during the Mexican drug war', *Proceedings of the 32nd Annual Conference on CHI*, ACM Digital Library, Toronto, ON, Canada, 26 April 26 to 1 May, pp.3563–3572.

Ediger, D., Jiang, K., Riedy, J., Bader, D.A., Corley, C., Farber, R. and Reynolds, W.N. (2010) 'Massive social network analysis: mining Twitter for social good', *Proceedings of the 39th International Conference on Parallel Processing (ICPP)*, ACM Digital Library, San Diego, CA, USA, 13–16 September, pp.583–593.

Fernandez, M., Scharl, A., Bontcheva, K. and Alani, H. (2014) 'User profile modelling in online communities', *Proceedings of the 3rd International Workshop on Semantic Web Collaborative Spaces, 13th International Semantic Web Conference (ISWC-2014)*, Riva del Garda, Italy [online] http://eprints.weblyzard.com/84/1/swsc-2014-user-profile-modelling.pdf (accessed 12 May 2015).

Glanz, K., Rimer, B.K. and Viswanath, V. (Eds.) (2008) *Health Behavior and Health Education: Theory, Research, and Practice*, Jossey-Bass, San Francisco.

Hamilton, A. (2014) *Heartbleed Bug Still a Risk for 300,000 Unpatched Servers* [online] http://www.itpro.co.uk/security/22538/heartbleed-bug-still-a-risk-for-300000-unpatched-servers (accessed 12 May 2015).

Heartbleed.com (2014) *The Heartbleed Bug* [online] http://heartbleed.com/ (accessed 12 May 2015).

Jacomy, M., Venturini, T., Heymann, S. and Bastian, M. (2014) 'ForceAtlas2, a continuous graph layout algorithm for handy network visualization designed for the Gephi software', *PloS ONE*, Vol. 9, No. 6, p.e98679.

Jonnalagedda, S.S.N. and Gauch, S. (2013) 'Personalized news recommender system using Twitter', *Proceedings of International Joint Conferences on Web Intelligence (WI) and Intelligent Agent Technologies (IAT) on IEEE/WIC/ACM*, IEE Xplore Digital Library, Atlanta, GA, USA, 17–20 November, Vol. 3, pp.21–25.

Kelion, L. (2014) *US Government Warns of Heartbleed Bug Danger* [online] http://www.bbc.co.uk/news/technology-26985818 (accessed 12 May 2015).

Kim, E., Gilbert, S., Edwards, M. and Graeff, E. (2009) 'Detecting sadness in 140 characters: sentiment analysis and mourning Michael Jackson on Twitter', *Web Ecology Project*, No. 3, pp.1–15 [online] http://www.webecologyproject.org/wp-content/uploads/2009/08/Detecting_Sadness_in_140_Characters2.pdf (accessed 12 May 2015).

Koutras, M.V., Bersimis, S. and Antzoulakos, D.L. (2006) 'Improving the performance of the chi-square control chart via runs rules', *Methodology and Computing in Applied Probability*, Vol. 8, No. 3, pp.409–426.

Kwak, H., Lee, C., Park, H. and Moon, S. (2010) 'What is Twitter, a social network or a news media?', *Proceedings of the 19th international Conference on World Wide Web*, ACM Digital Library, Raleigh, NC, USA, 26–30 April, pp.591–600.

Larsen, M., Boonstra, T., Batterham, P., O'Dea, B., Paris, C. and Christensen, H. (2015) 'We feel: mapping emotion on Twitter', *Journal of Biomedical and Health Informatics*, Vol. 19, No. 4, pp.1246–1252.

Li, Q., Wang, J., Chen, Y.P. and Lin, Z. (2010) 'User comments for news recommendation in forum-based social media', *Information Sciences*, Vol. 180, No. 24, pp.4929–4939.

Li, Z. and Li, C. (2014) 'Tweet or 're-tweet'? An experiment of message strategy and interactivity on Twitter', *Internet Research*, Vol. 24, No. 5, pp.648–667.

Lin, F. and Huang, H. (2013) 'Why people share knowledge in virtual communities?: the use of Yahoo! Kimo Knowledge+ as an example', *Internet Research*, Vol. 23, No. 2, pp.133–159.

Love, B., Himelboim, I., Holton, A. and Stewart, K. (2013) 'Twitter as a source of vaccination information: content drivers and what they are saying', *American Journal of Infection Control*, Vol. 41, No. 6, pp.568–570.

McNeill, A.R. and Briggs, P. (2014) 'Understanding twitter influence in the health domain: a social-psychological contribution', *Proceedings of the 23rd International Conference in Seoul on World Wide Web (WWW)*, ACM Digital Library, Republic of Korea, 7–11 April, pp.673–678.

Miettinen, O. and Nurminen, M. (1985) 'Comparative analysis of two rates', *Stat Med*, Vol. 4, No. 2, pp.213–226.

Misopoulos, F., Mitic, M., Kapoulas, A. and Karapiperis, C. (2014) 'Uncovering customer service experiences with Twitter: the case of airline industry', *Management Decision*, Vol. 52, No. 4, pp.705–723.

Miyabe, M., Nadamoto, A. and Aramaki, E. (2014) 'How do rumors spread during a crisis?: analysis of rumor expansion and disaffirmation on Twitter after 3.11 in Japan', *International Journal of Web Information Systems*, Vol. 10, No. 4, pp.394–412.

Mizzaro, S., Pavan, M. and Scagnetto, I. (2015) 'Content-based similarity of Twitter users', in Hanbury, A., Kazai, G., Rauber, A. and Fuhr, N. (Eds.): *Advances in Information Retrieval. Lecture Notes in Computer Science*, Vol. 9022, pp.507–512, Springer International Publishing, Switzerland.

Mols, F., Haslam, S.A., Jetten, J. and Steffens, N. (2014) 'Why a nudge is not enough: a social identity critique of governance by stealth', *European Journal of Political Research*, Vol. 54, No. 1, pp.81–98.

Morris, K. (2011) 'Tweet, post, share – a new school of health communication', *The Lancet Infectious Diseases*, Vol. 11, No. 7, pp.500–501.

Nieva, R. (2014) *How to Protect Yourself from the 'Heartbleed' Bug* [online] http://www.cnet.com/uk/news/how-to-protect-yourself-from-the-heartbleed-bug/ (accessed 12 May 2015).

Pagoto, S., Evans, M., Whited, M., Bauman, M., Vickey, T.A. and Schneider, K.L. (2013) 'Hashtag your way to health? The use of hashtags to create healthy communities and spread health behavior on Twitter', *Proceedings on Medicine 2.0*, London, UK, 23 September [online] http://www.medicine20congress.com/ocs/index.php/med/med2013/paper/view/1796 (accessed 12 May 2015).

Pang, A., Hassan, N.B.B.A. and Chong, A.C.Y. (2014) 'Negotiating crisis in the social media environment: evolution of crises online, gaining credibility offline', *Corporate Communications: An International Journal*, Vol. 19, No. 1, pp.96–118.

Pang, B. and Lee, L. (2008) 'Opinion mining and sentiment analysis', *Foundations and Trends in Information Retrieval*, Vol. 2, Nos. 1–2, pp.1–135.

Quercia, D., Kosinski, M., Stillwell, D. and Crowcroft, J. (2011) 'Our Twitter profiles, our selves: predicting personality with Twitter', *Proceedings of the Third International Conference on Privacy, Security, Risk and Trust (PASSAT) and IEEE Social Computing (SocialCom)*, IEEE Xplore Digital Library, Boston, MA, USA, 9–11 October, pp.180–185.

R Core Team (2014) *R: A Language and Environment for Statistical Computing*, R Foundation for Statistical Computing, Vienna, Austria.

Räbiger, S. and Spiliopoulou, M. (2015) 'A framework for validating the merit of properties that predict the influence of a Twitter user', *Expert Systems with Applications*, Vol. 42, No. 5, pp.2824–2834.

Richards, L. (1999) *Using NVivo in Qualitative Research*, Sage, London and Los Angeles.

Romenti, S., Murtarelli, G. and Valentini, C. (2014) 'Organisations' conversations in social media: applying dialogue strategies in times of crises', *Corporate Communications: An International Journal*, Vol. 19, No. 1, pp.10–33.

Saeri, A.K., Ogilvie, C., La Macchia, S.T., Smith, J.R. and Louis, W.R. (2014) 'Predicting Facebook users' online privacy protection: risk, trust, norm focus theory, and the theory of planned behavior', *Journal of Social Psychology*, Vol. 154, No. 4, pp.352–369.

Sloan, L., Morgan, J., Burnap, P. and Williams, M. (2015) 'Who tweets? Deriving the demographic characteristics of age, occupation and social class from twitter user meta-data', *PLoS ONE*, Vol. 10, No. 3, p.e0115545.

Slovic, P., Finucane, M.L., Peters, E. and MacGregor, D.G. (2004) 'Risk as analysis and risk as feelings: some thoughts about affect, reason, risk, and rationality', *Risk Analysis*, Vol. 24, No. 2, pp.311–322.

Tinati, R., Carr, L., Hall, W. and Bentwood, J. (2012) 'Identifying communicator roles in Twitter', *Proceedings on Mining Social Network Dynamics (MSND)*, ACM Digital Library, Lyon, FR, 16–20 April, pp.1161–1168.

Vargo, C.J., Basilaia, E. and Shaw, D.L. (2015) 'Event versus issue: Twitter reflections of major news, a case study', in Robinson, L., Cotten, S.R. and Schulz, J. (Eds.): *Communication and Information Annual (Studies in Media and Communications*), Vol. 9, pp.215–239, Emerald Group Publishing.

Veil, S.R., Buehner, T. and Palenchar, M.J. (2011) 'Work-in-process literature review: incorporating social media in risk and crisis communication', *Journal of Contingencies and Crisis Management*, Vol. 19, No. 2, pp.110–122.

Wang, E.S-T., Chen, L.S-L. and Tsai, B-K. (2012) 'Investigating member commitment to virtual communities using an integrated perspective', *Internet Research*, Vol. 22, No. 2, pp.199–210.

Xie, H., Li, X., Wang, J., Li, Q. and Cai, Y. (2014) 'The collaborative search by tag-based user profile in social media', *The Scientific World Journal*, Article ID 608326, pp.1–8.

Ye, S. and Wu, S. (2010) 'Measuring message propagation and social influence on Twitter.Com', *Lecture Notes in Computer Science on Social Informatics*, Vol. 6430, pp.216–231.

Zappavigna, M. (2014) 'Enacting identity in microblogging through ambient affiliation', *Discourse & Communication*, Vol. 8, No. 2, pp.209–228.

Zhang, Z. and Nasraoui, O. (2008) 'Profile-based focused crawler for social media-sharing websites', in *Proceedings of the20th IEEE International Conference on Tools with Artificial Intelligence (ICTAI '08)*, IEEE Xplore Digital Library, Dayton, OH, USA, 3–5 November, pp.317–324.

Zhou, T. (2011) 'Understanding online community user participation: a social influence perspective', *Internet Research*, Vol. 21, No. 1, pp.67–81.