

The Social News Gap: How news reading and news sharing diverge

Abstract

This article seeks to explain variation in news sharing patterns on social media. It finds that news editors have considerable power to shape the social media agenda through the use of “story importance cues”, but also shows that there are certain types of articles (such as those relating to crime and disasters) where this power does not apply. This highlights the existence of a social “news gap” where social media filters out certain types of news, hence producing a social media news agenda which has important differences to its traditional counterpart. The discussion suggests this may be consequential for perceptions of crime and engagement with politics; it might even stimulate a partial reversal of the tabloidization of news outlets.

Keywords: social media, news, news sharing, news gap, social networks

The Social News Gap: How news reading and news sharing diverge

This article addresses the question of why some news articles are shared more than others on social media platforms such as Twitter and Facebook. With their huge user bases, these platforms have become key venues for the dissemination of news content: as Newman puts it, “mainstream media content is the lifeblood of topical social media conversations” (2011, p. 6). However, while there is considerable research on the subject of why individual people choose to share news, hardly any work has addressed what type of aggregate sharing outcomes result from these micro motivations. Hence we know little about what drives overall sharing patterns. In particular, we do not know whether the types of news that gain prominence on social media differ from those emphasised by the traditional news agenda.

This article seeks to tackle this deficit. It is structured in the following way. Section one addresses the importance of news sharing on social media, and discusses existing literature on the micro-motivations for news sharing, connecting these motivations to hypotheses about overall sharing outcomes for individual articles. Section two describes the methodology employed in the article, including the collection of a novel dataset of around 2,400 news articles and the operationalisation of the concepts described in section one through a variety of automatic and manual techniques. Section three describes the results of studying the sharing patterns of these articles on four different social media platforms. The article finds that editorial promotion of the article on the front page and the topic of the article are both significant, even whilst controlling for readership, though considerable differences across platforms are also noted. Furthermore, it finds important divergences in terms of readership and sharing patterns for different types of article: news about technology and social welfare are well shared but not very well read, whilst the reverse is true for news about politics, crimes, accidents and disasters.

The article concludes by exploring the theoretical importance of these findings for the agenda setting power of the media. Two main arguments are presented. First, it is claimed that the significance of story importance cues for driving sharing mean that news editors have considerable power to shape the news agenda on social media. In this respect, social media enhances the power of the traditional media to set the overall news agenda, a finding which contradicts other work which has suggested that the news agenda will be undermined in the internet era. However, it is also shown that this power only applies for certain topics. This indicates a second argument, which is that we are seeing the emergence of a “social news gap” between news reading and news sharing preferences that will make news consumed through social media qualitatively different to news consumed directly through online platforms (equivalent to the “news gap” between news reader and news editor preferences identified by Boczkowski and Mitchelstein, 2013). When sharing news, social media users seem to filter out certain types of news content, such as political news and news which focusses on shocking crimes or accidental deaths. This has significant potential consequences. For the media agenda, it might mean that those reading the news on social media perceive less of a “mean world” than those reading its traditional counterpart; it might also contribute to further the type of media driven political disengagement described by Prior (2007). In terms of news production, meanwhile, it might also reverse some of the tendency towards tabloidization which has been encouraged by the emergence of real time audience statistics.

1) Literature Review: Explaining Aggregate News Sharing on Social Media

One of the most striking recent changes in the way the contemporary news media operate has been the emergence of high penetration social media sites such as Twitter and Facebook as venues for the distribution of news content. Material from online news websites makes up a significant proportion of all information shared in social media (Kwak, Lee, Park,

& Moon, 2010, p. 595), and a significant proportion of social media users share links to news stories (Purcell, Rainie, Mitchell, Rosenstiel, & Olmstead, 2010, p. 2). News media organisations are also making considerable efforts to distribute their content through social media (Ju, Jeong, & Chyi, 2013), with the ability to share at the click of a button increasingly integrated into many news websites (Singer, 2014; Stroud, Scacco, & Curry, 2014); an option that users of online news sites appreciate (Purcell et al., 2010).

The widespread nature of news sharing has meant that major social media sites such as Twitter and Facebook have become very significant sources of traffic for online news sites (for a review see Napoli, 2015), with a recent report claiming that they outrank even search engines as means of directing readership (Ingram, 2015). Their importance is even greater for local news or smaller news outlets (Newman, Dutton, & Blank, 2012; Sonderman, 2012). Increasing quantities of people also list social media as a major venue where they find the news (Mitchell & Guskin 2013; Nielsen & Schroder 2014). On this basis, some have theorised the emergence of a new type of news consumer who, instead of actively reading the news, waits for friends and social relations to deliver it to them (see e.g. Baresch, Knight, Harp, & Yaschur, 2011, p. 2). But even for those who continue to consume news in a more traditional way, social media platforms are developing into an important secondary distribution channel.

The rise of social media as a news distribution channel is significant because of the potential it has to challenge and unsettle the agenda-setting power of the mass media (McCombs & Shaw, 1972). Agenda-setting, of course, refers to the power major news outlets have to shape what topics people are thinking about (as well as shaping how important they think they are, what they think about them, and how they connect specific stories to more general issues such as their evaluation of the competence of political parties – for a good

summary see Scheufele & Tewksbury, 2007). This power emerges from the ability that news media editors have to control which stories are published and hence read.

The potential of the internet to challenge the news media agenda has of course long been a subject of interest in studies of online communication. Several authors have suggested that by significantly increasing choice, and allowing ordinary people to both publish and communicate relatively cheaply, the power of major news outlets to produce a unified news agenda appears under threat through “the prospect of citizens themselves participating in the agenda-setting process” (Goode 2009, p. 1293; see also Bruns 2003; Delli Carpini & Williams 2005; Bennett & Iyengar, 2008). However, the potential impact of social media is different. By widely sharing news from mainstream outlets, social media users do not directly undermine these outlets: indeed, they enhance their influence. But in selectively deciding which types of news to share, these users nevertheless also make editorial decisions that may have an impact on which news becomes more or less visible (Singer, 2014, p. 58). Hence rather than undermining the news agenda, they have the potential to filter and reformat it: perhaps marginalising news that mainstream editors thought was important, whilst promoting other stories that mainstream news editors thought were minor. Furthermore, news sharing may act as an important signal to journalists of the popularity of their content (Holton, Coddington, Lewis, & de Zúñiga, 2015), influencing decisions about which stories to promote on the front page of a news website or improving the chances of commissioning a follow up piece, in the same way that automatically updated readership statistics have already been shown to do (Bright & Nicholls 2014; Lee, Lewis, & Powers, 2012). However, despite the potential importance of news sharing, as yet we know little about the extent to which social media actually reshapes the news agenda, as opposed to simply reproducing what editors already think is important. The aim of this article is to fill this gap, by exploring why some articles are shared more than others.

Micro-motivations for news sharing

Existing literature on news sharing has largely focussed on the micro-motivations behind the decision to share news (that is, why one person decides to pass on news to another), which Singer has called “secondary gatekeeping” (2014, see also Shoemaker & Vos 2009, p. 6). Such literature actually has a long history, which predates social media to a considerable degree (for an early example see Deutschmann & Danielson, 1960). Broadly speaking, two main motivations for sharing news have been identified. First, social connections and social status can be enhanced through the sharing of news (see Gantz & Trenholm, 1979; McCombs, Holbert, Kioussis, & Wanta, 2011, pp. 135-137). Sharing information can make the person passing it on appear well informed and intelligent, and the shared knowledge may also prove useful for the person receiving it, both things that can contribute to improving social standing. This in turn can be a way of either establishing new social relationships or maintaining existing ones. These motivations will be referred to here as “status led” sharing.

Second, researchers have highlighted the importance of news sharing as a way of helping to understand and hence form opinions on the news. As Neuman and Guggenheim put it, “individuals rely on social cues and interpersonal conversation to interpret and contextualize complex media messages” (2011, p. 177). For particularly surprising or dramatic events, people often report sharing news with complete strangers, in order to allow them to have someone to talk to (see e.g. Rogers & Seidel, 2002, p. 215): hence a variety of pieces of research have documented how in cases of dramatic events such as assassinations the majority of people discover the news through friends (see e.g. Hill & Bonjean, 1964; Mayer, Gudykunst, Perrill, & Merrill, 1990). In these cases, sharing news has a certain cathartic value (Ibrahim, Ye & Hoffner, 2008): and studies have documented that people are more likely to share news that they felt emotionally connected to (see e.g. Basil & Brown,

1990; Harber & Cohen, 2005). This connection can defy the absence of any actual personal link in the case, for example, of the death of a celebrity (something often conceptualised as a “parasocial” relation – see Horton & Wohl, 1956). These motivations will be referred to as “emotion led” sharing.

This literature has recently been updated with specific reference to social media, where research has found that socializing and status seeking continue to be significant drivers (Lee & Ma, 2013; Baek, Holton, Harp, & Yaschur, 2011). The self-perception of a user as an opinion leader, and their estimation of the strength of social connections, is also important (Ma, Lee, & Goh, 2014). In other words, findings have largely confirmed motivations for news sharing discovered in other contexts, though they place more focus on status led sharing than emotion led sharing (it may be that social media do not offer a good venue for detailed discussion of the news that emotion led sharing seems to imply).

However, whilst much is known about micro motivations for news sharing, much less work has been done on explaining overall aggregate outcomes (one exception to this is Boczkowski and Mitchelstein’s work on email sharing patterns, which will be discussed below). In other words, even though we may have a good idea why people share news, we have much less of an idea about which pieces of news gets shared the most. Inferences cannot easily be drawn from the micro to the macro level. It may be true that people want to pass on news to enhance their social status, but it is hard to define which *types* of news will enhance this status, something that is also likely to change from one person to the next. Furthermore, status cannot be enhanced by sharing old news, which means that sharing decisions are also determined by what has already been shared within an immediate social circle. Equally, it may be true that news about important events is shared more widely, but *defining* importance in abstract is a very difficult task (for a discussion see Rosengren, 1987, pp. 230-232), and the

perception of importance may also be affected to a significant degree by the way the news itself is presented (Rogers, 2000, p. 563).

Explaining Aggregate Sharing Outcomes

The main aim of this article is to explain variation in aggregate social sharing outcomes for news articles: that is, the total amount of users who decide to share an individual article on their social networks. This section will examine factors that might be used to explain the aggregate amount of sharing an article receives, based on the above discussion of micro level motivations.

A first factor in this regard is the extent to which the news organisation that published the article actively promotes it on its front page. Editorial promotion could be significant because research has demonstrated that it affects how people interpret the story as being more or less important. For example, Iyengar and Kinder find that lead stories in television broadcasts are more likely to be recalled as an important part of the news agenda than stories that appear later in the programme (2010, pp. 34-46), whilst in the online context specifically, Althaus and Tewksbury have shown that readers of online and offline versions of newspapers perceive the news agenda differently, highlighting in particular differing types of “story importance cues” available in the different contexts (2002, p. 183). Perceiving a story as important may be relevant for both the status led and emotion led news sharing motivations described above: more important news is likely to be more significant in terms of enhancing social status, whilst the type of news that provokes emotion led sharing (surprising and dramatic events) is also typically perceived as more important. Editorial promotion is also theoretically consequential because, if it is effective, it will show that news editors have the ability to shape the social media agenda in much the same way as they do the traditional agenda, by choosing which articles to place emphasis on in their publications.

A second factor is the subject matter of the article itself: some topics may lend themselves more to sharing than others. For example, research has indicated that people prefer to pass on bad news, especially news that is exaggeratedly bad (Heath, 1996): hence news about disasters, accidents and emergencies might be shared more. This relates to the emotion led sharing described above: people may feel more personally touched by disastrous news, and they may also feel more need to discuss and contextualise it, both things which might provoke sharing. Stories about celebrities may also be more likely to be shared: first because they might be stories to which people feel an emotional connection (Horton & Wohl, 1956), second because they have been identified as useful ways of starting conversations amongst social connections, an important part of status led sharing (Feasey, 2008). Finally, Boczkowski and Mitchelstein also identify stories which are either bizarre and unusual or useful as particularly likely to be shared via email (see Boczkowski & Mitchelstein, 2013, p.114). This also fits in with the status led type of news sharing described above: hence we might also expect these stories to be similarly shared on social media.

Furthermore, we might also expect that there are some article subjects which are particularly poor candidates for social media sharing. For example, research has suggested people will avoid political stories in conversation (Eliasoph, 1998), especially if they are known to be controversial or they think others hold diverging views (Martin, 2008, p. 82; Jensen, 1990, p. 67; Boczkowski, 2010, p. 135). Another example is the “guilty pleasures” type of media use that people like to engage in privately but not advertise (see e.g. Baruh, 2010; Panek, 2014). It may be that stories about crimes and sex scandals fall into this category. Such stories could have an actively negative effect on social status because of what they suggest about a person’s interests, meaning people avoid sharing them.

A final factor concerns the actual context within which the sharing takes place. Existing work on interpersonal news diffusion has shown that different social and cultural

contexts will mean the same story is considered more or less important in different areas (see e.g. Gantz and Tokinoya, 1987, Rogers, 2000, p. 568), and hence more worthy for sharing and discussion. This raises a further question: does the same apply to different social media platforms, which have different functionalities and aim to attract different user bases (Kietzmann, Hermkens, McCarthy, & Silvestre, 2011)? For example, one important area of difference concerns the extent to which the social network in question focusses on making professional connections (such as networks like LinkedIn) as opposed to more personal connections (like Facebook). The professionalization of a social network should heighten the status led sharing motivations discussed above: when building professional contacts, the need to appear informed and in the know is even more important. Another difference concerns the type of messages that can be exchanged, with Twitter for example focussing on very short, rapid communication, whilst other platforms allow for longer messages and comparatively slower communication (Kaplan & Haenlein, 2011). Hence we might expect social sharing to take place at different speeds on different platforms. The speed of communication may also have an impact on some of the “story importance cues” listed above: for example, fast moving social networks may be even more likely to share top headlines.

This discussion of three factors which might affect news sharing (editorial promotion, article topic and the context of the social media platform) also points to a potential confounding variable in this analysis, which is the amount of readership a news article receives. Readership is of course bound to influence sharing as an article cannot be shared if it has not been readⁱ, and hence more readers implies more potential sharers (of course, this relationship is also likely to be self-reinforcing, as more sharing is also likely to stimulate further readership). However readership is also likely to be correlated with the factors described above. Editorial promotion will drive readership: Graber, for instance, has demonstrated that the positioning of a story on a newspaper front page, combined with other

cues such as the size of the headline and the presence of any pictures, improved the likelihood of the piece being read (1988). Subject matter will also drive readership: for example, Graber also found that crime articles were likely to generate attention regardless of the degree of editorial importance attached (recently Boczkowski and Mitchelstein formalised this disconnect in the idea of a “news gap”, showing how readers and editors diverge in terms of preferences – 2013).

However, as well as being a control variable, readership is also theoretically significant. If it is the most important factor determining sharing, then one might argue that social news distribution is in a certain sense epiphenomenal, simply a reflection of what is already being read, and hence unlikely to alter the news agenda in any important sense. However if it is unimportant then this will point to differences between the traditional news agenda and the social media agenda. Hence identification of the extent to which news articles can become widely shared without being very well read (or vice versa) is an important secondary question for the article.

2) Methodology

To summarise so far, this article tackles a simple question: what explains variation in amounts of social media sharing for different news articles? Three main lines of investigation are pursued: variation in the amount of emphasis editors place on stories (“story importance cues”); different article topics, with some more suitable for sharing than others; and the different context provided by different social media platforms. A second question concerns the extent to which readership levels matter for social sharing: whether sharing is simply a reflection of what is being read, or whether differences emerge between the reading and sharing priorities of online news consumers.

This section describes the methodology used to tackle these questions. The empirical basis of this study is a dataset of articles drawn from the BBC News website.ⁱⁱ BBC news was

chosen because it is a highly visible news website that generates considerable sharing patterns on all of the content it produces, something that facilitates the systematic assessment of variance in sharing. Data collection for this article took place during a two week periodⁱⁱⁱ, during which time the front page of BBC News was downloaded once every 30 minutes by a piece of software developed specifically for the research project. The two week period was selected to give a large enough sample of articles to make generalisations, whilst nevertheless leaving it small enough so that all articles could be coded^{iv}. After each download, every hyperlink found on the front page was extracted. The resulting list was then pared down to include only links related to individual news articles (i.e. links to other sections of the website such as the sports sub section were discarded, as were links to other websites). Following this process, four different social media platforms were then accessed to see how many times the link in question had been shared^v. This querying took place once every half an hour for the 24-hour period after the link was initially observed. In total, over 400,000 observations were made of sharing counts for over 2,400 articles. The dependent variable for this study is the total amount of shares that an article received at the end of this 24-hour period.

The independent variables for this study are designed to capture the major lines of investigation suggested by the discussion above: story importance cues, article topic, context and readership levels. They are operationalised through a mixture of automatic and manual techniques, which are described in turn below.

Story Importance Cues

The presence of story importance cues was operationalised automatically. Four different variables were created to measure this. First, the amount of time the story spent on the front page of the website was recorded, with articles that stay around longer likely to be those editors deem more important. Second, the positioning of the article on the front page was recorded, which allowed the creation of a positioning score that is relative to both the

location of the article (nearer the top of the central part of the screen giving a greater score) and the amount of time the article remained there. Precise calculation details for this score are provided in the appendix. Third, a dichotomous variable recorded whether an article had an image attached to it, on the basis that larger and more important stories often are presented alongside an image. Finally, the time of day the article was published was recorded as a categorical control variable.

Article Topic Classification

The basis of the topic classification was the policy agendas codebook, which has previously been used as a way of breaking down news stories in recent research (Boydston, Hardy & Walgrave, 2014). In particular, the UK media codebook developed by Jennings and Bevan was employed (2010, pp. 50-51), which extends the policy agendas categorisation to fit a news media context (adding in categories such as “sports and recreation” and “human interest”). The topic of the article was coded through a mixture of manual and automatic methods. The BBC provides some limited classification information in the structure of its links (for example, all “sports” news stories have a link that indicates that they belong to the sports category). These were used where possible to classify the link (60% of links could be classified in this way). Where this was not possible (the remaining 40% of links), the story itself was read by the author and assigned to the appropriate category. The results of the classification (shown in table two below) showed some categories to be of very low frequency: these were added into a default “other” category. More details on the classification, including notes on intercoder reliability, are provided in the appendix.

Context

Differences in context are provided by tracking sharing counts on different social media platforms. Four platforms were chosen: were Facebook, GooglePlus, Twitter and LinkedIn. These four platforms were chosen to provide a range of different contexts within

which social sharing takes place, given the remarks on the potential importance of context above. While of course all similar in terms of the basic design and functionality of a social network, these four networks provide variation in terms of the extent to which the network is professionally oriented. LinkedIn is a network designed purely for connections between professionals, largely with the aim of advancing career goals. Twitter is more of a mixture, combining many social elements but also heavily used by professionals. Facebook and GooglePlus, by contrast, are more oriented towards social functions. They also provide variation in terms of speed, with Twitter as a fast micro-blogging service providing a comparison to the other three sites.

Readership Scores

It was not possible to directly measure the level of readership for a given article. However, like many news websites, the BBC maintains a list of “most read” articles on its front page: such lists have previously been used as proxies for readership levels in other research (e.g. in Boczcowski & Mitchelstein, 2013). The software developed checked each article link for presence on this most read list. This allowed the development of a most read “score” that is weighted to take into account the amount of time an article spends on the most read list and the position it appears in. This score is used as a proxy for readership levels. The precise calculation details for the score can be found in the appendix.

3) Results

Descriptive statistics for the sharing patterns observed are shown in table one, whilst descriptive statistics for the independent variables are provided in table two. Almost all links were observed to have been shared at least once on one of the social networks, and the average link receives 1,221 shares in total when all networks are summed together. However, there is wide variation in the absolute amounts of sharing activity across the networks, with

LinkedIn and GooglePlus showing relatively modest amounts of sharing (an average of 10 shares per link), whilst Twitter and especially Facebook show much higher absolute amounts.

Table one about here

To some extent this would be expected based on user numbers alone: Facebook has a user base considerably larger than any of the other three social media sites^{vi}. However, considering that the user bases of Twitter, LinkedIn and GooglePlus are comparable, it also highlights the different types of usage being made of each network. It is also worth highlighting the skewed nature of the distribution of sharing on all sites, with most links receiving a relatively small number of shares, whilst a few receive an amount many standard deviations above the mean.

Table two about here

The major research question addressed in this study concerns why some news articles are shared more than others. This question is addressed in a series of analytical models, presented in table three, which seek to explain variance in aggregate social sharing outcomes for different news articles. Model one is a full model, which analyses social sharing on all four networks combined. Models two to five are specific to individual networks: they hence address the second question of how sharing patterns diverge on different social networks.

The dependent variable for model one is the sum total amount of shares received on all social media sites. Before aggregation, the amount of shares received on each site was divided by the standard deviation of shares received on that platform. This correction was introduced because the absolute number of shares on each platform is very different. The dependent variable for models two to five is simply the number of shares received on the platform in question. This number is also divided by the standard deviation, which makes the results of all models directly comparable. All dependent variables are furthermore log transformed, which changes their distribution from highly skewed to a good approximation of

normality. This transform meant that links with zero shares had to be discarded: hence the N varies between models. It also means the effects are interpreted in terms of percentage changes in sharing^{vii}.

Table three about here

The analytical models presented in table three highlight a number of findings about social sharing. In terms of story importance cues, the article positioning score is found to be important (increasing sharing by around 8% for every point increase in the score in the full model), though it has a smaller effect on Twitter and Google than on the other platforms. Furthermore, the amount of time that a link spends on site has a positive effect in all models. However, the presence of an image link is only significant on Twitter (where it leads to around 30% more shares), whilst publication time does not have too much of an effect, though morning publications are more favoured on Google and especially LinkedIn, which might be expected given LinkedIn's heavy bias towards morning sharing highlighted in figure one. This evidence shows that story importance cues are important, but not all of them matter on all platforms.

There are also considerable effects in the topic categories^{viii}. Some of these effects run across all platforms. Stories about accidents and disasters, and stories about law, crime and family issues (of which the vast majority were about crime) have considerably lower sharing patterns than average (48% less shares for both cases in the full model) contradicting the idea that people might prefer to share bad news. Stories about politics and government are also associated with lower sharing rates (41% less in the full model), with the effect duplicated on all platforms apart from Twitter.^{ix} This lends support to the idea that controversial political topics are generally avoided when sharing news. Sports stories were also associated with negative sharing rates (the positive effect for sports in the Twitter model was no longer observed in the more robust Twitter model presented in the appendix).

By contrast, stories about space, science and technology were shared at a much greater than average rate on all platforms (around 150% more); stories about social welfare also attracted positive sharing in the full model and on LinkedIn. The increased sharing patterns for technology news might be because they fit into the kind of unusual or bizarre news which was identified by Bozowski and Mitchelstein (2013, p. 114), as they typically deal with new discoveries or advances in technology. Some topic effects diverged: news about economics was well shared on LinkedIn, but poorly shared on Facebook. Human interest news was less popular on GooglePlus and LinkedIn, though had no effect on the other models (largely contradicting the idea that news about celebrities might be good for sharing). This evidence shows, in other words, that the topic of a story does have a considerable effect on sharing patterns: and that some topics have an impact independent of platform whilst others have an impact that is platform specific. This again we might expect from what is known about the divergence in platforms, particularly their professionalization. Sharing stories about banking might improve your social status on LinkedIn but worsen it on Facebook, whilst sharing stories about crime is negatively perceived everywhere.

Finally, as expected readership is an important parameter in all models, though the strength of effect does vary, with readership especially important on Facebook but less important on LinkedIn. It should be highlighted again here that the relationship between readership and sharing is expected to be self-reinforcing, and the readership variable in the model is provided primarily as a control variable. However, the divergence of strength between social networks is intriguing, and might point to the possibility highlighted above that some networks facilitate re-sharing more than others (and re-sharing might be achieved without necessarily reading the article).

The data collected also allows description of the temporal evolution of social sharing of news articles on different platforms. This provides a further way of addressing the impact

of different contexts on social sharing. Figure one plots the average cumulative number of shares each link receives against the amount of hours after initial publication on the front page of BBC News. To aid comparison, red lines have been added to mark the time where average cumulative shares cross the 50% and 90% marks.

Figure one about here

This figure shows that social media is a space where sharing happens very quickly following publication decisions. In all cases, 50% of the total sharing observed in the day occurs less than three hours after publication, and 90% occurs less than 16 hours after publication. In the last third of the 24-hour observation period, rates of sharing slow considerably. However as expected there are also some marked differences in sharing speed between networks. Facebook and GooglePlus are comparatively slower, with the average link crossing 50% of cumulative shares around 2 hours after publication and 90% after around 15 hours. LinkedIn and especially Twitter, by contrast, are much quicker to pass these barriers. On Twitter, the average link receives 50% of its shares in the first hour after publication and crosses the 90% barrier before 10 hours are passed. It seems that both the orientation of a site towards microblogging and its professionalization tend to increase the speed at which sharing happens.

Figure one also looks at how the average rate of shares fluctuates during the working day (hence data from the weekends is dropped). All of the sites show a dip in sharing activity in the early hours of the morning, however beyond that the patterns are somewhat different. Twitter and especially LinkedIn show a clear peak during working hours, skewed especially towards the morning. Facebook and GooglePlus show a more consistent pattern of usage, with sharing not declining until much later in the evening. Again, this is to be expected given the relative levels of professionalization of each network.

Figure two about here

As described above, a secondary aim of the article is to address the question of where news readership and news sharing patterns diverge. This is important because it highlights how news shared on social media diverges from what people read when they consume news through traditional channels. In order to address this question, figure two plots the average readership and sharing scores found in different topic categories. This figure shows the readership score on the X axis and the sharing score on the Y axis. Both scores have been scaled to fall between 0 and 100 (where 100 is the maximum readership or sharing score observed for any topic category). Note that category labels have been shortened to improve readability. The red line represents a perfect relationship between readership and sharing: hence categories near to the red line have similar amounts of reading and sharing. Below the red line are the categories that are read more than they are shared, and vice versa above the line.

The results highlight some interesting points. Some of the categories that have the lowest levels of sharing (law, crime and family issues and accident and disaster especially) are also those that are amongst the highest levels of readership. These categories, which seem to fit into the “guilty pleasure” type of news, constitute an important amount of all news published (around 17% based on table two above), yet are barely shared at all. Human interest stories are also notable for having very high levels of readership yet only moderate sharing. Conversely, categories such as social welfare and science and technology have levels of sharing that are disproportionately high compared to their readership.

These results point to an extension of model 1. Readership seems to have a different effect on different article topics: hence an interaction term could be included between the readership score and topic categories. This model (which is not fully reported here to conserve space) confirms the intuition of figure two: the size of the coefficient for the readership score is increased when compared to model 1 (0.40 instead of 0.28), as the model

now correctly accounts for the fact that readership has a strong effect but not a consistent one across topics. Meanwhile, the interaction terms for economics, human interest, and accidents and disasters are all statistically significant and negative. This shows that the correlation between readership and sharing is much weaker for articles in these categories.

Discussion and Conclusion

This article has tackled the question of why some news articles are shared more than others. Four findings were presented. First, story importance cues were shown to be significant in driving social sharing, even when controlling for readership levels (that such cues are also likely to drive). This was attributed to the desire of individuals to share news which they perceive to be important, thus enhancing their social status. Second, different topics were shown to have different outcomes when it comes to sharing. Stories about politics, accidents, disasters and crime were poorly shared, whilst the reverse was true for stories about social welfare stories and science and technology. These differences were attributed again to the status led component of news sharing (there was no real evidence for the idea that emotion led sharing might be important on social media). A person's news reading habits provide an indication of their interests, which means that some topics appear to be specifically avoided when it comes to sharing because they might inflict reputational damage (for example, by suggesting a morbid interest in crime or disasters). By contrast, some topics appear to be especially good candidates for sharing because they can enhance reputations (perhaps by suggesting a compassionate interest in social welfare or showing knowledge of new advances in technology).

A third finding was that there were significant differences between social networks. Sometimes this was just in terms of the size of effect: for example, news about crime was penalized more heavily on Facebook than it was on Twitter. But sometimes news had different effects on different platforms: news about economics, for instance, was popular on

LinkedIn but unpopular on Facebook. The dynamics of sharing were also shown to be different on different platforms, with some sharing much faster than others. This appeared to be connected to the use of the platform for professional reasons, though further work would be needed to confirm this. Finally, readership patterns and sharing patterns were compared, with evidence produced that some news topics are popular for reading but not for sharing, and vice versa. For example, it was showed that people do like to read news about crime but they do not like to share it. This shows evidence of a “social news gap” between what people like to read and what they share, equivalent to the “news gap” identified by Boczkowski and Mitchelstein (2013) between what people like to read and what editors like to publish.

What are the theoretical implications of these findings? The fact that story importance cues matter is significant because it contradicts the idea that editorial agenda-setting power might be in decline in the digital era (Goode, 2009; Bruns, 2003; Delli Carpini & Williams, 2005; Bennett & Iyengar, 2008). Even though the audience has the ability to decide what news to share (Singer, 2014), their decisions are also shaped by the presentational choices made by news editors. Hence in some respects social media enhances the power of the news media to set the news agenda. However the importance of article topic, and the divergence between reading and sharing, shows that this power is not complete. There are some areas where the social media news agenda differs significantly from its traditional counterpart (and hence where the power of news editors themselves is limited): crime stories are unlikely to be widely shared (even if they are presented as important), whilst science and technology stories are likely to gain significant sharing (even if news editors perceive them to be of minor importance). Furthermore, it was also shown that not only will the social media agenda differ from the news agenda, but this difference is distinct on different platforms.

These results have two major implications. First, in the long term, this might mean that those who receive news by social media end up having a significantly different

perception of the news agenda than those who get their news straight from a traditional outlet. Social media amplifies the reach and importance of certain types of news whilst ignoring and marginalising other types. It is possible to speculate about what the impact of this might be. Social media news consumers might perceive less of a “mean world” (Gerbner, Gross, Morgan, & Signorielli, 1980), because they receive less news about crime and disasters. This could have ramifications for the extent to which the public perceive themselves at risk of being a victim of crime. Another impact could be that social news consumers become disconnected from the political process, as their friends shy away from sharing controversial political topics. This could increase the kind of disengagement described by Prior (2007): though instead of disengagement coming through the active choice of individuals to ignore political media, it would come as a result of the choices of their friends and social connections.

A second implication concerns the incentives for news editors created by social media. A variety of scholarship has worried that the editorial motivation to improve online readership, combined with the ready availability of audience metrics, is having the effect of promoting a kind of tabloidization in online news, whereby popular content is favoured over more important “hard news” (see *inter alia* Cohen, 2002; Brossard & Scheufele, 2013; Shoemaker & Vos, 2009, p.7). The need to improve social sharing may however offer an opposite incentive to editors, as classic tabloid style news stories about crimes, disasters and sex scandals do not lend themselves to social sharing even if they are very popular amongst readers. From this perspective, the power of social media to shape the news agenda may actually be considerable, as social media sharers start to drive not only what is read but also the type of content that editors produce.

It is worth concluding by highlighting the limitations of this study, and thus setting out areas for future research. Though the study encompasses multiple social media platforms,

there are many more left unstudied. Further work could address whether the dynamics found here are repeated elsewhere. More importantly, it is limited to one online news source. It might be, for example, that readers of BBC News in particular are disproportionately unlikely to share news about crime (even if they do like reading about it): hence it would be worth exploring if these dynamics are repeated with respect to other online news sources.

Finally, it might be argued that the focus on aggregate sharing outcomes is misplaced: as everyone's experience of social media is different, and based on their social connections, overall social sharing patterns are less important than the sharing conducted by one's individual circle of social connections. While there is certainly some merit to this argument, it is worth noting that overall outcomes are important, for two reasons. First, the overall popularity of news articles is being increasingly integrated into the way social media platforms recommend new content: hence an article that is shared more may also come to the attention of more people even if their friends have not shared it. Second, and more importantly, whilst it might be true that news agendas are personalised, this does not mean that there are infinite varieties of these agendas, or that they are completely distinct from one another. Rather, it is more likely that each personal agenda is a subtle twist on one of a few overarching main themes. In this respect, average sharing outcomes can provide us with a way of measuring the "average" social media agenda. However, further research comparing personalised agendas with average outcomes would be needed to confirm this.

References

- Althaus, S., & Tewksbury, D. (2002). Agenda setting and the “new” news patterns of issue importance among readers of the paper and online versions of the New York Times. *Communication Reserarch*, 29(2), 180-207.
- Baek, K., Holton, A., Harp, D., & Yaschur, C. (2011). The links that bind: Uncovering novel motivations for linking on Facebook. *Computers in Human Behaviour*, 27, 2243-2248.
- Baresch, B., K, L., Harp, D., & Yaschur, C. (2011). Friends who choose your news: An analysis of links on Facebook. *#ISOJ*, 1(2), 65-85.
- Baruh, L. (2010). Mediated Voyeurism and the Guilty Pleasure of Consuming Reality Television. *Media Psychology*, 13(3), 201-221.
- Basil, M., & Brown, W. (1990). Interpersonal communication in news diffusion: a study of “Magic” Johnson’s announcement. *Journalism Quarterly* 71(2): 305-320.
- Bennet, W. L., & Iyengar, S. (2008). A New Era of Minimal Effects? The Changing Foundations of Political Communication. *Journal of Communication*, 58(4), 707-731.
- Boczkowski, P.J. (2010). *News at Work*. Chicago: University of Chicago Press.
- Boczkowski, P.J., & Mitchelstein, E. (2013). *The News Gap*. Cambridge: MIT Press.
- Boydston, A., Hardy, A., & Walgrave, S. (2014). Two Faces of Media Attention : media Storm Versus Non-Storm Coverage. *Political Communication*, 31(4), 509-531.
- Brenner, J., & Smith, A. (2013). *72% of Online Adults are Social Networking Site Users*. Washington: Pew Research Centre.
- Bright, J., & Nicholls, T. (2014). The Life and Death of Political News: Measuring the Impact of the Audience Agenda Using Online Data. *Social Science Computer Review*, 32(2), 170-181.

- Brossard, D., & Scheufele, D. A. (2013). Science, New Media and the Public. *Science*, 339, 40-41.
- Bruns, A. (2003). Gatewatching, not gatekeeping: Collaborative online news. *Media International Australia Incorporating Culture and Policy: quarterly journal of media research and resources*, 107, pp. 31-44
- Cohen, E. L. (2002). Online journalism as market driven journalism. *Journal of Broadcasting and Electronic Media*, 46(4), 532-548.
- Delli Carpini, M. X., & Williams, B. A. (2005). Monica and Bill All the Time and Everywhere: The Collapse of Gatekeeping and Agenda Setting in the New Media Environment. *American Behavioural Scientist*, 47(9), 1208-1230.
- Deutschmann, P. J., & Danielson, W. A. (1960). Diffusion of Knowledge of the Major News Story. *Journalism Quarterly*, 37, 345-55.
- Eliasoph, N. (1998) *Avoiding Politics*. Cambridge: Cambridge University Press.
- Feasey, R. (2008). Reading heat: The meanings and pleasures of star fashions and celebrity gossip. *Continuum: Journal of Media & Cultural Studies*. 22(5), 687-699.
- Gantz, W., & Tokinoya, H. (1987). Diffusion of News About the Assassination of Olof Palme: A Trans-Continental, Two-City Comparison of the Process. *European Journal of Communication*, 2, 197-210.
- Gantz, W., & Trenholm, S. (1979). Why People Pass on News: Motivations for Diffusion. *Journalism Quarterly*, 2, 365-370.
- George, L., & Hogendorn, C. (2012). Aggregators, search and the economics of new media institutions. *The Economics of Digital Media Markets*. 24(1), 40-51.
- Gerbner, G., Gross, L., Morgan, M., & Signorielli, N. (1980). Living with television: the dynamics of the cultivation process. In J. Bryant & D. Zillman (Eds.), *Perspectives on Media Effects* (pp.17-40). Hillsdale: Lawrence Erlbaum

- Goode, L. (2009). Social news, citizen journalism and democracy. *New Media and Society*, 11(8): 1287-1305.
- Graber, D. (1988). *Processing the News: How People Tame the Information Tide*. London: Longman Group.
- Harber, K., & Cohen, D. (2005). The Emotional Broadcaster Theory of Social Sharing. *Journal of Language and Social Psychology*, 24, 382-400.
- Heath, C. (1996). Do People Prefer to Pass Along Good or Bad News? Valence and Relevance of News as Predictors of Transmission Propensity. *Organisational Behaviour and Human Decision Processes*, 68(2), 79-94.
- Hill, R.J., & Bonjean, C.M. (1964). News Diffusion: A Test of the Regularity Hypothesis. *Journalism Quarterly*, 41, 336-42.
- Holton, A., Coddington, M., Lewis, S., & de Zúñiga, H. G. (2015). Reciprocity and the News: The Role of Personal and Social Media Reciprocity in News Creation and Consumption. *International Journal of Communication*, 9, 2526-2547.
- Horton, D., & Wohl, R. (1956). Mass Communication and Para-social Interaction. *Psychiatry*, 19, 215-29.
- Ibrahim, A., Ye, J., & Hoffner, C. (2008). Diffusion of News of the Shuttle Columbia Disaster: The Role of Emotional Responses and Motives for Interpersonal Communication. *Communication Research Reports*, 25(2), 91-101.
- Ingram, M. (2015, August 18) Facebook has taken over from Google as a traffic source for news. *Fortune*. Retrieve from: <http://fortune.com/2015/08/18/facebook-google/>
- Iyengar, S., and Kinder., D. (1987). *News That Matters: Television and American Opinion*. Chicago: University of Chicago Press.
- Jennings, W., & Bevan, S. (2010). UK Topics Codebook. Available from: <https://policyagendasuk.files.wordpress.com/> [Accessed 05/08/2015]

- Jensen, K. B. (1990) The politics of polysemy: television news, everyday consciousness and political action. *Media, Culture & Society*, 12, 57-77.
- Ju, A., Jeong, S. H., & Chyi, H. I. (2014). Will social media save newspapers ? Examining the effectiveness of Facebook and Twitter as news platforms. *Journalism Practice*, 8(1), 1-17.
- Kaplan, A., & Haenlein, M. (2011). The early bird catches the news: Nine things you should know about micro-blogging. *Business Horizons*, 54(2), 105-113
- Kietzmann, J., Hermkens, K., McCarthy, I., & Silvestre, B. (2011). Social media? Get serious! Understanding the functional building blocks of social media. *Business Horizons*, 54(2), 241-251.
- Krippendorff, K. (2004) Reliability in content analysis: Some common misconceptions and recommendations. *Human Communication Research*, 30, 411-433.
- Kwak, H., Lee, C., Park, H., & Moon, S. (2010). What is Twitter, a Social Network or a News Media? *Proceedings of the 19th international conference on world wide web*, 591-600.
- Lee, A. M., Lewis, S. C., & Powers, M. (2012). Audience Clicks and News Placement: A Study of Time-Lagged Influence in Online Journalism. *Communication Research*.
- Lee, C.S., and Ma, L. (2013). News sharing in social media: The effect of gratifications and prior experience. *Computers in Human Behaviour*, 28(2), 331-339.
- Lowry, D., Nio, T., & Leitner, D. (2003). Setting the Public Fear Agenda: A Longitudinal Analysis of Network TV Crime Reporting, Public Perceptions of Crime, and FBI Crime Statistics. *Journal of Communication*, 53(1), 61-73.
- Ma, L., Lee, C.S., Goh, D. H-L. (2014). Understanding news sharing in social media : An explanation from the diffusion of innovations theory. *Online Information Review*, (38)5, 598-615.

- Martin, V. (2008). Attending the news. *Journalism*, 9(1), 76-94.
- Mayer, M., Gudykunst, W., Perrill, N., & Merrill, B. (1990). A comparison of competing models of the news diffusion process, *Western Journal of Speech Communication*, 54(1), 113-123.
- McCombs, M., Holbert, L., Kioussis, S., & Wanta, W. (2011). *The News and Public Opinion: Media Effects on Civic Life*. Cambridge: Polity Press.
- McCombs, M., & Shaw, D. L. (1972) The agenda-setting function of mass media. *Public Opinion Quarterly*, 36(2), 176-187.
- Mitchell, A., & Guskin, E. (2013). Twitter news consumers: Young, mobile, and educated. Washington, DC: Pew Research Center. Available from <http://www.journalism.org/2013/11/04/twitter-newsconsumers-young-mobile-and-educated/> [Accessed 03/08/2015]
- Napoli, P. M. (2015). Social media and the public interest: Governance of news platforms in the realm of individual and algorithmic gatekeepers. *Telecommunications Policy*. doi:10.1016/j.telpol.2014.12.003
- Neuman, R., & Guggenheim, L. (2011). The Evolution of Media Effects Theory: A Six-Stage Model of Cumulative Research. *Communication Theory*, 21, 169-196.
- Newman, N. (2011). *Mainstream media and the distribution of news in the age of social discovery*. Oxford: Reuters Institute for the Study of Journalism.
- Newman, N., Dutton, W., & Blank, G. (2012). Social Media in the Changing Ecology of News: The Fourth and Fifth Estates in Britain. *International Journal of Internet Science*, 7(1), 6-22.
- Nielsen, R.K., & Schrøder, K.C. (2014). The Relative Importance of Social Media for Accessing, Finding, and Engaging with News. An eight-country cross-media comparison. *Digital Journalism*, 2(4), 472-489.

- Panek, E. (2014). Left to Their Own Devices College Students' "Guilty Pleasure" Media Use and Time Management. *Communication Research*, 41(4), 561-577.
- Prior, M. (2007). *Post-Broadcast Democracy*. Cambridge: Cambridge University Press.
- Purcell, K., Rainie, L., Mitchell, A., Rosenstiel, T., & Olmstead, K. 2010. *Understanding the participatory news consumer: How internet and cell phone users have turned news into a social experience*. Washington: Pew Research Centre.
- Rogers, E. (2000). Reflections on News Event Diffusion Research. *Journalism and Mass Communication Quarterly*, 77(3), 561-576.
- Rogers, E., & Seidel, N. (2002). Diffusion of News of the Terrorist Attacks of September 11, 2001. *Prometheus: Critical Studies in Innovation*, 20(3), 209-219.
- Rosengren, K. (1987) Conclusion: The Comparative Study of News Diffusion. *European Journal of Communication*, 2, 227-255.
- Scheufele, D., & Tewksbury, D. (2007). Framing, Agenda Setting, and Priming: The Evolution of Three Media Effects Models. *Journal of Communication*, 57, 9-20.
- Shoemaker, P., & Vos, T. (2009). *Gatekeeping Theory*. London: Routledge.
- Singer, J. (2014). User-generated visibility: Secondary gatekeeping in a shared media space. *New Media and Society*. 16(1), 55-73.
- Stroud, N., Scacco, J. M., & Curry, A. L. (2014). Interactive features in online news. Report from the Engaging News Project. Available from: <http://engagingnewsproject.org/>
[Accessed 05/08/2015]
- Sonderman, J. (2012, September 13). Study: Smaller news websites depend more on social media for traffic than larger sites. *Poynter*. Retrieved from: www.poynter.org

Appendix

Readership Score

The readership score was created using the list of “Most Read” news items that appears on the BBC website’s front page. Every time the front page was downloaded (once every thirty minutes) the articles in this list were compared to the articles already held in the dataset. Any article found to be in the list had its readership score incremented by the following amount:

$$((1 / \text{Position}) * 10) / 2$$

The position is inversed and multiplied by 10 to take into account the descending nature of the most read list (meaning that position 1 is worth 10 times as much as position 10). It is divided by 2 to reflect the 30-minute time window (allowing the readership score to be presented in an hourly format). Hence, for example, an article that spent one hour in position 10 of the most read list and 30 minutes in position 5 of the readership list would finish with a score of 2.

Positioning Score

The positioning score is based on the same logic as the readership score, but with the article’s position on the front page being the important factor. Examination of the HTML of the BBC web page showed that the top of its central column (which is what can be seen by a reader immediately upon opening the page in a typical web browser) is divided into 5 different sections in decreasing order of height on the central page: “top story”, “second story”, “third story”, “other top stories” and “also in the news”. These divisions typically contain one or two major story links, with a number of other minor stories (that may be related to the major piece). In total each section contains 2-10 stories at any one time. These sections were used as the basis of the positioning score. Again, this score is weighted to the

amount of time an article spends in front page positions, as well as the height of the overall position. The top section was worth 5 positioning points every hour, the second 4, and so on. Any articles falling outside of these top 5 sections received 0 positioning points.

Coding

Each article was assigned to one distinct topic category. The topic categories themselves were taken from the media codes section of the UK policy agendas codebook (see Jennings and Bevan 2010, 50-51). The assignment was done in two ways. First, for certain articles, information can be found within the structure of the hyperlink of the article itself about the topic. For example, this link <http://news.bbc.co.uk/sport/0/formula1/23476191> is for an article about sports (specifically about Formula 1 racing). This link information was used where possible to code the category of the article. All links that could not be assigned in this way were assigned by hand by the author. In total 1,467 links out of 2,431 were classified automatically (60%), with the remaining 964 (40%) coded by hand.

A second coder was asked to code a random selection of 10% (243) of the articles. This coder was provided with a codebook and a general description of how the coding was performed. Calculation of intercoder reliability [ICR] statistics showed that this coder agreed with the author's coding 68.7% of the time, with a Krippendorff's alpha of 0.652. In cases when alpha is lower than desired, Krippendorff recommends combining categories and removing distinctions from the data in order to improve agreement amongst coders (2004, 430). After examining areas of disagreement between the coders, a variety of categories were selected for combination because they were thematically similar. These combinations did not in fact imply large changes in the coding of the dataset: only 74 observations were recoded. Using this amended coding schema, ICR statistics were recalculated, and a percentage agreement of 79.8% was produced, with an alpha of 0.773. This amended coding schema is the one which is used in all the results and analysis in the main paper.

Tables

	Total	Facebook	Twitter	GooglePlus	LinkedIn
Mean Shares Per Link (SD)	1221 (4024)	828 (3029)	367 (2103)	16 (85)	10 (24)
% of links with shares	99.5	99	98	64	77

Table one: Link sharing patterns across different social networks

Variable	Summary	Variable	Summary
Mean readership score (SD)	14.4 (46)	% morning publication	39%
Mean positioning score (SD)	12.3 (26)	% daytime publication	49%
Mean time in hours spent on website (SD)	8.8 (8)	% evening publication	12%
		% of links with an image	40%
Category	Count	Category	Count
International Affairs and Foreign Aid	662	Politics and Government	70
Law, Crime, and Family Issues	291	Education	34
Sports and Recreation	263	Transportation	32
Economics	247	Civil Rights, Minority Issues, Immigration and Civil Liberties	17
Other	207	Community Development, Planning and Housing Issues	15
Science, Technology and Environment	199	Defence	12
Human Interest	157	Social Welfare	10
Accidents and Disasters	113	Energy	5
Health	97		

Table two: Descriptive statistics

		Model 1: Full Model	Model 2: Facebook	Model 3: Twitter	Model 4: Google Plus	Model 5: LinkedIn
Story Importance Cues	Positioning Score (log transformed)	0.08***	0.11***	0.05**	0.05*	0.12***
	Image	0.24**	0.14	0.26***	-0.02	0.03
	Time on Site (square root transformed)	0.24***	0.19***	0.20***	0.13***	0.08***
	Evening Publication	-0.19*	-0.05	-0.11	-0.05	-0.09
	Morning Publication	0.13*	-0.03	-0.03	0.15*	0.34***
	Adjusted R²	0.16	0.15	0.16	0.13	0.14
Topics	Accident and Disaster	-0.65***	-0.77***	-0.24	-1.17***	-0.58***
	Economics	0.55***	-0.52***	0.13	-0.18	0.93***
	Civil Rights, Minority Issues, Immigration and Civil Liberties	-0.31	-0.06	-0.42	-1.02**	0.07
	Education	-0.16	-0.49	-0.32	-1.22***	0.36
	Human Interest	-0.05	0.04	-0.11	-0.42**	-0.43***
	Health	0.11	-0.04	0.17	-0.25	0.38**
	Law, Crime, and Family Issues	-0.65***	-0.84***	-0.38***	-0.99***	-0.44***
	Politics and Government	-0.53**	-0.48*	-0.13	-0.96***	-0.53***
	Social Welfare	1.11*	0.92	0.7	0.14	0.90*
	Science, Technology and Environment	0.91***	0.43***	0.40***	0.90***	0.98***
	Sports and Recreation	-0.24*	-0.39***	0.26**	-1.23***	-0.57***
	Transportation	-0.45	-0.78**	-0.45*	-0.58*	-0.25
	Adjusted R²	0.07	0.07	0.07	0.11	0.12
Readership	Readership Score (log transformed)	0.28***	0.39***	0.22***	0.26***	0.11***
	Adjusted R²	0.13	0.13	0.13	0.14	0.16
Full Model	Adjusted R²	0.27	0.29	0.24	0.28	0.31
	N	2420	2410	2388	1550	1878

Table three: Regression models explaining news sharing in different social networks.

Figures

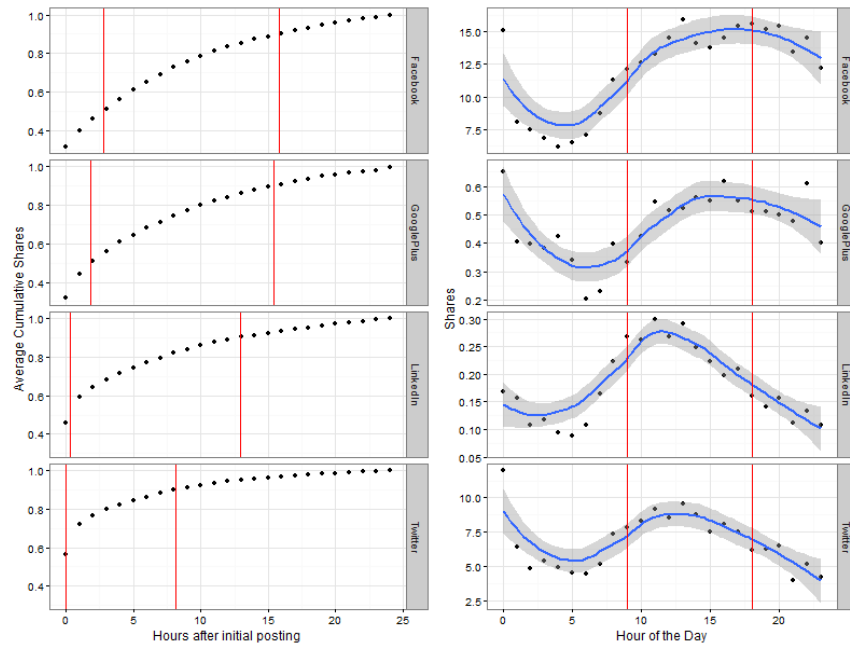


Figure one: (Left Panel) Average cumulative share dynamics in different social networks. Red lines indicate the hour of crossing points for 50% and 90% of cumulative shares (Right Panel) Temporal patterns in share dynamics. The blue line is a LOESS curve with 95% confidence intervals^x




Figure two: Sharing and Readership Scores for Different Topic Categories. Note that topic labels have been shortened

Endnotes

ⁱ An exception to this rule is the fact that on social networks such as Twitter people can re-share content that they have received through their social networks, without necessarily reading it first.

ⁱⁱ The site can be found at: <http://www.bbc.co.uk/news>

ⁱⁱⁱ The sampling period ran from 9.30am on the 24th of July 2013 to 9.30am on the 7th of August 2013: any articles published on the BBC News Online website during this period form part of the dataset. In order to correctly identify which articles were genuinely new, regular samples of the front page were also taken in the two days leading up to the data collection period: any articles published during this time period were excluded from the dataset. As shares were collected up to 24 hours after initial observation samples of the front page were also collected up to the 8th of August to fully track any links that had been created on the 7th (however no new links were added to the dataset during this time).

^{iv} Actual weeks were used, rather than constructed ones, in order to facilitate the capture of data that needed to span multiple consecutive days (for example, fluctuations in position on the most read list). The potential impact of using such a sampling approach is addressed further in the appendix.

^v This tracking was conducted using the API provided by the website “sharedcount.com”. This API offers a service that queries the APIs of the four social networks in question (as well as several others used) and returns sharing counts for any given link. These sharing count APIs are the basis of the web “buttons”, which can be found on many websites, that record how many times a certain article has been shared on a given social media platform.

^{vi} At the time of writing, Facebook had well over a billion users, whilst Twitter, LinkedIn and GooglePlus had comparatively similar user bases at around the 300 million mark. Statistics obtained from <http://www.statista.com>. It should be noted that these statistics reflect created accounts, rather than actual user activity.

^{vii} Diagnostic tests for model one indicated a number of further variable transformations were appropriate: each one is listed in the model. The same transformations were used in models two to five as well, again to improve comparability. Separate diagnostic tests for each of these models indicated however that different transformations would improve the fit in some cases. There are slight differences between these more valid models and the ones presented in the paper; these are mentioned in the text where appropriate.

^{viii} The reference category for the topic variable is “International Affairs and Foreign Aid”, which was chosen because it was a numerous category situated close to the mean value for topics. Hence the effect of each topic variable can also be interpreted in terms of increasing or decreasing sharing relative to the “average” topic. In order to simplify presentation, all topic variables that produced no statistically significant results in any model were removed.

^{ix} In the more robust model presented in the appendix, a negative relationship can also be seen for Twitter.

^x This graphic does not include a small number of outlier observations, where large bursts of sharing happened in very short spaces of time, as they make it difficult to observe the underlying pattern. The cut-off point was set at amounts of sharing more than three standard deviations above the norm.