Rumoring During Extreme Events: A Case Study of Deepwater Horizon 2010

Emma S. Spiro
Department of Sociology
University of California
Irvine, California
espiro@uci.edu

Jeannette Sutton
Trauma Health and Hazards Center
University of Colorado
Colorado Springs, Colorado

Matt Greczek
Trauma Health and Hazards Center
University of Colorado
Colorado Springs, Colorado

Sean Fitzhugh
Department of Sociology
University of California
Irvine, California

Nicole Pierski
Department of Sociology
University of California
Irvine, California

Carter T. Butts
Department of Sociology
Institute for Mathematical Behavioral Sciences
University of California
Irvine, California

ABSTRACT

Social scientists have proposed many different factors thought to influence rumoring behavior. Classical rumor theory points to the perceived importance, the level of uncertainty or ambiguity, and the potential to impact decision making as influential in determining the extent of rumoring. In this work, we test some of these proposed rumor determinants in the context of the the 2010 Deepwater Horizon oil spill, using data on communication dynamics from the popular microblogging service Twitter. Using a latent factor model, we measure rates of hazard-related conversation by exploiting joint variation in multiple conversation streams. Time series analysis of the resulting rates suggests that media coverage of the event is a major driver of rumoring behavior, supporting importance/saliency theories and disconfirming theories of information substitution for this event. Relevance of the event to decision making behavior also turns out to be an influential predictor in this case. Since information diffusion via serial transmission is a fundamental process by which rumors spread, we compare rates of serial transmission between control and hazard-related communication. Twitter posts are much more likely to be retweeted when they contain hazard-related keywords (versus control words). Implications of these findings for disaster response are discussed.

Author Keywords

informal online communication, rumoring, disaster response, microblogging

ACM Classification Keywords

H.5.2 Information Interfaces and Presentation: Miscellaneous

WebSci 2012, June 22–24, 2012, Evanston, Illinois, USA.

General Terms

Experimentation, Human Factors

INTRODUCTION

The nearly continuous, informal exchange of information – including such mundane activities as gossip, rumor, and casual conversation – is a characteristic human behavior [11, 13]. Individuals utilize social ties to obtain information, whether it be about the weather tomorrow or to find possible job leads [16]. Recent developments in social media technologies and mobile devices have transformed informal communication channels allowing individuals to reach a larger number of contacts across much greater distances than previously possible. Such patterns of informal communication have important implications during extreme events, such as those associated with natural or anthropogenic hazards.

Research has shown that social ties are also pathways for information exchange in the event of a disaster or other emergency situation [24, 12, 27]. In many cases informal social ties are actually the *primary* means by which people obtain time-sensitive information, especially when official sources are slow to release updates or are unavailable. With the rise of social media technologies, such informal communication is increasingly conducted in the online environment. Within the emergency management community, there is a growing appreciation for the potential value of these online communication networks for emergency warnings, alerts, etc. However, despite this recognition, researchers know very little about the detailed dynamics of informal online conversation surrounding exogenous events and the underlying networks that structure it, especially in the context of emergency response.

In this work, we explore the dynamics of informal online communication (or rumoring) in the context of the 2010 Deepwater Horizon oil spill. We utilize data from a popular microblogging service to evaluate alternative theories about the determinants of rumoring [6]. Our methods utilize multiple conversation streams, defined by usage of specific keywords, in order to accurately estimate aggregate trends in conversation. We compare informal online communication to proposed rumor determinants in order to evaluate their applicability to this case. In addition, we explore serial transmission behavior during the disaster (a central concern of classical rumor theory, e.g. [1]). Transmission of information from one individual to another is vital for information propagation and diffusion. Though specific to the Deepwater Horizon case, our results offer insight into long-term, large-scale environmental disasters in general. They also have important implications for communication in the disaster response context, as well as research on online informal communication dynamics more generally.

BACKGROUND

Rumoring and Informal Communication

Rumoring, as employed in this paper, refers to communication surrounding facts or events of topical interest that does not occur as part of a formal, institutionalized process (e.g., news broadcasts, warnings and alerts, press briefings, etc.). Classical theories of rumor transmission identify three important factors thought to affect the extent of rumoring and the propensity that individuals will pass along information obtained from others regarding a salient event [1, 7, 5, 26]. These factors include the perceived importance of the event, the degree of cognitive unclarity surrounding the event, and the relevance to behavior of the event.

Drawing from their classic demonstrations of rumoring, Allport and Postman [1] posit that a rumor's intensity is a multiplicative function of its importance and ambiguity. Since this foundational work, there has been much debate about the mechanisms involved in rumor transmission. Many scholars have recognized that rumor experiments rarely replicate the popular image of rumors spreading like wildfire in a population [25, 2, 17], with most showing very little diffusion. Massive diffusion does occasionally occur [31], however, making rumor dynamics a complex process. In particular, effective studies of rumoring behavior should ideally consider not only "successful" rumors that diffuse widely through the population, but also those disseminated to very few parties. This informs our approach to the measurement of rumoring behavior, as discussed below.

Since these early studies, there has emerged a growing body of work that explores different determinants of rumoring behavior. Research indicates that anxiety increases the probability of passing on information [2, 32]. In addition, individual characteristics may affect the paths along which information is exchanged [24], with a tendency for information to be passed to others with similar demographic characteristics to the sender. Another important factor may be the content or form of the message itself, [18, 34, 4]. Environmental contexts, social relationships, and cognitive limitations are all thought to play a role. Together, all of these factors offer direct mechanisms through which the environment and characteristics of an external event may systematically produce changes in information exchange.

Yet another body of work in this area explores the relationship between official or authoritative (as contrasted with informal sources), and rumoring [9]. In his classic work, Caplow [7] argues that official sources of information should *suppress* rumoring behavior: authoritative sources will replace informal channels, making it unnecessary for individuals to seek out information through social ties and casual communications. The exact influence of authoritative sources, however, is still an open question. Fragale and Heath [14], for example, demonstrate that belief in message content is directly related to perceptions of source credibility. Not only do individuals mis-attribute rumors to credible sources when they believe them, but they are also prone to believe information more readily when it comes from credible sources.

Rumoring might also be considered a special case of general information diffusion. In recent years, studies of information diffusion via social ties have offered insight into the structural forms that facilitate or inhibit the exchange of information [21, 8, 19]. These studies focus on tracing information flow in large-scale social networks and characterizing the patters of diffusion [15, 35]. Both the structure of the underlying social network and the context of exchange have important consequences for rumor spread. In the context of social media, rumoring behaviors has been considered, however, less of this work looks at how classic rumor theories operate [20, 3].

Theories of rumor posit many different factors that influence the extent of information exchange and transmission. Yet, there are still many unanswered questions. In particular, differences in face-to-face communication and online information exchange suggest interesting consequences for rumoring and the proposed factor that influence a rumor's prevalence.

Online Communication During Extreme Events

Early researchers studied rumoring characteristics in wartime [26], finding that in the absence of authoritative information, individuals will often share unsubstantiated claims in an attempt to make sense of the event and its surrounding circumstances. Recognized as "improvised news" such activities adapt available information into plausible conclusions that are passed among affected individuals.

This literature also suggests, actors will use their social ties in order to obtain factual information regarding important matters when "official" channels fail [26, 30], particularly information regarding negative or anxiety-provoking events [23, 18, 33, 29]. This notion is backed up by findings from the disaster research literature, which reveal a strong tendency for actors in crisis settings to use social networks to obtain factual information regarding imminent hazards [10, 22]. In fact, in cases where official sources are unavailable (or insufficiently timely), such social ties will serve as the primary conduits of information. Indeed, such conduits generally outpace official sources [12, 28].

While crisis situations evoke increased frequency of interpersonal communication, much of this communication takes place through existing channels: a consistent finding in the above studies is that existing, frequently used network ties were overwhelmingly employed in passing on crisis information. Thus, the first source of notification for many individuals in crisis situations is information diffusion through existing social ties and modes of communication. In the online context, then, everyday tools for interpersonal communication are likely to be those first employed for both seeking and disseminating hazard information.

When crises occur, available social media are "appropriated" for the purpose of collecting and disseminating disaster-relevant information, and new disaster-related content is rapidly created and shared. The same collective behavior processes that have been observed over time during disasters, such as mass convergence, rumoring, and the formation of emergent groups occur within an online environment. What differs in this form of collective behavior response to disasters is the use of new technology to enable communication and information sharing, as well as the distributed nature of voluntary participation.

Indications of significant social and technological change in crisis-related behaviors were first notably apparent in the December 2004 Indian Ocean tsunami, in which globalization and the rise of pervasive ICT expanded the size of publics that can learn about physical and social disaster impacts and participate in response, relief, and restoration activities. Additional evidence of the importance of informal online communication has come from studies of subsequent events. Following the 2007 Southern California wildfires, for instance, researchers investigated information seeking and sharing practices online and found that individuals participated in these "back-channel" communications due to a perception that there was a dearth of accurate information, that public officials and major media outlets were too slow to provide relevant information to communities at risk [29].

Among the obstacles to utilization of social media as information sources, particularly for emergency managers and official government entities, is our current lack of knowledge regarding the factors that shape informal communication patterns during events. Here, we examine these factors in the context of the 2010 Deepwater Horizon oil spill.

DEEPWATER HORIZON 2010 OIL SPILL

On 20 April 2010, an explosion on the Deepwater Horizon offshore drilling rig killed 11 crewmen and resulted in a massive oil spill. It is estimated that the oil spill released more than 5 million barrels, or 206 million gallons, of crude oil into the Gulf of Mexico, directly affecting coastal communities in five states. Fishing areas were closed, devastating the livelihoods of local fishermen and resulting in a declaration of disaster for fisheries off the coast of Louisiana where fishing, shrimping, and oyster bed production was severely affected. It is the largest accidental marine oil spill in the history of the petroleum industry. The disaster has had far-reaching consequences sufficient to impact global economies, marketplaces and policies.

DATA

Informal Online Communication

There are many venues for informal online communication related to the 2010 Deepwater Horizon oil spill. In this study we utilize a large sample of publicly available micro-blog posts from the widely used service Twitter.com. Twitter has become a popular venue for informal communication; conversation topics are widely diverse but not structured as part of a formal, institutionalized process. Twitter has over 300 million users across the world[1]. The premise of Twitter is relatively simple; messages, known as *tweets*, are posted by individual users and then delivered to that person's subscribers, known as followers. This message-passing structure allows users to exchange informal information. Users may also search the set of publicly available tweets in search of information or users they might be interested in. The scale of this micro-blogging community makes collection of all traffic infeasible. Instead, we employ *keywords* as indicators of underlying discussion frequency, inferring the presence of hazard-related conversation from changes in the overall rate with which particular terms are used. Using a combination of observational and statistical controls, we will assess the impact of hazard events on the volume and structure of informal online communication.

The current dataset contains tweets dating from two different observation periods. The first ranges from May 8, 2010 to July 14, 2010 and captures it initial response period. The second from September 1, 2010 to October 23, 2010; here we capture some of the long term effects and more of the recovery period in the response[2]. The dataset was collected by [6]. The data collection strategy was designed around a list of keywords of interest; these keywords, specified a priori, were designed to capture oil spill related conversation. The list of the keywords used to capture oil spill related discussion can be found in Table 1[3]. In addition we also have a control topic where words are chosen at random from Odgen's English Word List, also seen in Table 1. The Twitter SEARCH API was then used to collect tweets containing these keywords over the two observation periods. API queries were made adaptively, based on the previously observed rate of posting, to avoid rate limiting by Twitter. While missing data does occur, this system allows censored points to be identified. In addition, exploiting the joint variation among the multiple keyword streams ensures relatively good estimates of aggregate trends in conversation. More details on the dataset are found in subsequent sections.

Determinants of Rumoring

One of the motivations for this research is to explore the determinants of rumoring behavior. As described previously, many theories of rumor have been proposed in the literature.

[1]http://en.wikipedia.org/wiki/Twitter

[2]Though we would have ideally liked to capture the entire period, data collection tools were still being tested. However, the current dataset captures important points in the response periods and it has well understood data quality.

[3]A greater number of control keywords were chosen a prioi in order to better approximate general levels of conversation on Twitter across time

Keyword	Num. of Tweets	Avg. Hourly Rate
BP	3,830,081	878
deepwater horizon	49,779	14
environmental disaster	12,759	2
gulf coast	151,136	40
oil spill	1,817,480	438
wildlife	265,594	70
chalk	420,597	19
cloth	624,796	38
cloud	5,892,112	455
collar	1,077,849	59
control	12,309,660	642
form	8,272,310	408
heat	6,936,386	502
secretary	1,404,791	124
trouble	8,778,176	416

Table 1. Dataset descriptives: keywords are shown for event-related oil spill conversation as well as control-related conversation. Keywords in the control sample are taken from Ogden's English Word List.

We collect a series of covariates designed to approximate many of the proposed influences on conservation levels. In particular we are interested in the effect of national media coverage, official sources of information, measure of impact and importance, and levels of uncertainty or anxiety in the population. We briefly describe the set of covariates used in this work.

Official Sources of Information

Turning to formal information sources, the most prominent for both experts and laypersons are the traditional news media. In order to measure the extent of event-relatd news coverage by traditional media outlets during the study period, mentions of the Deepwater Horizon oil spill in United States newspapers and newswires was captured using Lexis-Nexis for the same time period that Twitter data was collected (May 8, 2010 to October 23, 2010). Articles that mentioned the Deepwater Horizon by name, articles with keywords "gulf," "gulf coast," "oil spill," and/or "'oil leak" were captured and collected if they pertained to the Deepwater Horizon spill. These keyword are similar to those used in the collect of data from Twitter, however, they are adjusted to fit the two very different styles of discussion in these two different environments. From this corpus of published news articles we count the number of pertinent articles published per day, giving an estimate of the level of national news coverage over the period of observation.

In addition to media coverage we consider offical press releases issued by `restorethegulf.org`, the official site for the Federal response, during the event period. This potential predictor speaks directly to theories of official news sources, such as those proposed by Caplow [7]. Again we consider the number of press releases per day over the observation period. It is important to note that official information sources are typically very controlled and slow. We observe only a few, typically one, official press release per day.

Public Interest

Another estimate of the information availability during the oil spill comes from the online collaborative information archive Wikipedia. We obtain counts of the number of edits and views of the Deepwater Horizon oil spill 2010 page over the observation period. The relationship between long-term information collected and organization with short-term informal communication is an interesting area for future research. Here, however, we are simple concerned with how page views and edits might proxy for interest or uncertainty about the event. It may also speak towards the perceived importance of the event for the online community.

Environmental Impact

Since the primary effect of the oil spill was on the environment of the Gulf coast, wildlife impact served as a natural measure of the importance of the event for individual and group decision making behavior. Daily wildlife updates by the Deepwater Horizon Response Consolidated Fish and Wildlife Collection Report included information on counts of contaminated birds, sea turtles, mammals, and other reptiles, collectively a proxy for observable environmental impact. Oil spill impact on wildlife could have a significant effect on the levels of conversation among concerned members of the public, as an expression of anxiety about the health of native species along the gulf coast. We consider both the total number of contaminated animals collected and released back into the wild.

Another area of significant concern during the event was the impact of the oil spill on fisheries within the region. To capture this, we collected data on federal water areas closed to fishing, as reported by the National Oceanic and Atmospheric Administration (NOAA, a U.S. federal agency focused on the condition of the oceans and the atmosphere) in press releases. Information on water closure serves both as a measure of wildlife impact and as an economic indicator of the severity of the oil spill on commercial fishing and aquaculture (especially along the coastline of Louisiana). Likewise, water area closures are also indicative of concerns regarding food contamination, a problem of relevance not only to those involved in food production, but also to consumers throughout the region.

MODELING RUMOR DYNAMICS

As noted above, a key indicator of the prevalence of discussion or rumoring surrounding a given topic is the frequency with which posts using topic-relevant terms appear within the broader stream of micro-blog posts: while the presence of a given keyword may be neither necessary nor sufficient to conclude that a specific post speaks to a specific topic, aggregate changes in keyword frequency for topic-related keywords relative to control terms provide a strong indication of topic-related discussion prevalence.

To this end, we consider a latent factor model for the topic related keywords. The data used in this research consist of a large sample of tweets containing a given set of predefined keywords collected over a pre-specified period. Consider the series of hourly rates of posting for each tweet

	Model 1			Model 2			Null Model		
Covariate	Coefficient	s.e.		Coefficient	s.e.		Coefficient	s.e.	
AR(1)	0.8145	0.0570	***	0.8014	0.0594	***	0.8229	0.0580	***
Intercept	-0.9749	0.3863	*	-0.9174	0.3883	*	1.5468	0.2053	***
Trend	-0.0049	0.0014	***	-0.0049	0.0014	***	-0.0104	0.0019	***
Log News Coverage	0.3087	0.0608	***	0.2966	0.0617	***			
Wikipedia Edits	-0.0622	0.0526		-0.0557	0.0535				
Water Area Closed	0.0051	0.0024	*	0.0051	0.0024	*			
Saturday	0.0944	0.0435	*	0.0887	0.0443	*			
Monday	0.0464	0.0284		0.0456	0.0290				
Wednesday	-0.0487	0.0300							
Control PCA Scores							-0.3643	0.0996	***
Loglikelihood	51.3911			50.1067			17.2512		
AICC	-81.8765			-81.5133			-24.3560		

Table 2. Parameter Estimates for Top AR(1) Models

streams. Rate estimates are straightforward given the sampling scheme. Twitter maintains a history, for each keyword stream, of 1500 tweets or three months time, whichever is reached first. However, they also enforce restrictions of the number of times per hour one may query. Thus each time the Search API is queried, if we obtain the maximum value of 1500 new (previously unseen) tweets then we know that the observations have been censored; adjusting for this effect is thus important for accurate rate estimation.

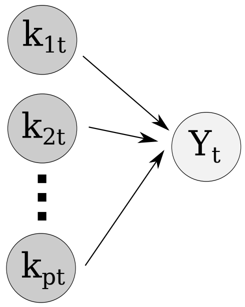

Figure 1. Latent factor model. Joint variation in multiple keyword rate time series.

To estimate rates within the above framework, we employ the following strategy. Let z_{it} be the number of tweets observed for the tth hour of data collection on keyword i, and let δ_{it} be the length of the time interval for keyword i within the tth hour prior to censoring (if any). We take the number of tweets on each keyword within each hour as independently Poisson distributed, with keyword by hour rate parameters λ_{it}. To estimate λ_{it}, we employ the corresponding posterior mean estimate under a Jeffreys prior, which reduces to $\hat{\lambda_{it}} = (z_{it} + 0.5)/\delta_{it}$. In the rare cases for which an entire hourly interval was censored, we interpolate $\hat{\lambda_{it}}$ as the mean of the rate estimates immediate adjacent to it. The resulting series of rate estimates, $\hat{\lambda}$ are then employed for the analyses which follow.

Given the set of multiple time series of posting rates, our aim is the capture the joint movement of these keywords. Factor analysis allows one to reflect the variability of many different observed variables, in our case the different keyword rate estimates, in fewer observed variables, or factors. We use confirmatory factor analysis to determine the subset of words that best represent the joint movement, that is, load most strongly on the first latent factor of the rate estimates ($\hat{lambda}$) for the keywords associated with the oil spill.

Estimating a one factor model for the oil spill related keywords is a simple means of capturing the joint tendency for movement in the level of conversation about the event of interest. Intuitively, we assume that overall propensity to post oil spill related tweets varies over time. Given that one posts such a tweet, one has a certain (unknown) chance of using one or more of the keywords in Table 1. Not all on-topic tweets will contain such keywords, however, and by turns not all uses of a given keyword relate to the oil spill topic. Thus, we model the traffic on each keyword as arising partially from traffic related to the oil spill per se, and partially from other ideosyncratic topics. We show the form of this model in Figure 1. $k_{it} = \log \hat{\lambda i}t$ is here the logged hazard estimate for the ith keyword, while Y_t is the (latent) log frequency of discussion on the oil spill topic. Under the factor model, we take $k_{it} \sim \mathrm{Normal}(\mu_i + \theta_i Y_t, \sigma_i^2)$ for each i and t, with μ_i being a keyword-specific base rate, θ_i a real parameter reflecting the loading of the keyword in question on the general topic, and σ_i^2 a keyword-specific parameter reflecting the idiosyncratic variation in rates (i.e., variance not explained by variation in discussion of the underlying topic). Given the observed sequence of k values, we estimate μ, σ, and θ via maximum likelihood. We then use the time series of resulting factor scores (Y_t) as our measure of the underlying oil spill conversation rate. (Note that this latent rate can be estimated only up to an affine transformation; as we are interested here only in the dynamics of this rate relative to other factors, this is not a problem for our analysis.)

Just as we control for ideosyncratic variation in specific spill-related keywords, we would like to control for overall changes in the use of Twitter itself. These could include, for example, aggregate increases in use or frequency of posting, seasonality effects, etc. Here, we utilize our control sample of keywords for this purpose. These words, chosen randomly from Ogden's Basic English word list, allow us to estimate a joint movement in general conversation (independent of any particular topic). We again estimate the joint variability of these multiple rate estimates, in this case using the scores on the first principle component of the log rate estimates as a control for conversation levels.

Given the above preliminaries, our primary interest is in the evolution of topic-related conversation over time, controlling for the baseline conversation level. Given that we have hourly estimates for the log conversation rate, Y_t, it is natural to approach this problem within the general framework of time series analysis. We find that a basic autoregressive model (AR) fits well in this context. Specifically, we model the oil spill related conversation rate as a function of itself at a lag one, along with a set of covariates. We also explicitly model a trend component, to control for secular change in topic frequency over time. Our model is as follows:

$$Y_t = \alpha_0 + \alpha_1 Y_{t-1} + \beta_1 X_1 \ldots \beta_p X_p + \gamma T_t + \epsilon_t \qquad (1)$$

The response Y_t here represents the latent factor scores for hazard-related conversation, as described above. $\vec{\alpha}$ is our parameter of auto-regressive parameters. $\mathbf{X}$ represents a covariate matrix, with corresponding parameters $\vec{\beta}$. We explicitly model the trend component, T_t, to assess changes in rates of communication over time. We estimate these models using standard estimation techniques. To utilize the `arima` function in the standard `stats` package in the **R** statistical software environment. Results are presented below.

RESULTS

Rumoring Dynamics

We use a basic, auto-regressive time series framework to model the changes in conversation levels. The response of interest here is the time series of factor scores for the oil spill related keywords. We consider each of the previously discussed rumor determinant proxies as predictive of conversation levels, including controls for general trends and seasonality in conversation using a set of English keywords chosen at random from the Ogden Basic English word list. Performing model selection using AICC as our criterion results in the top models seen in Table 2. [4]

As seen in Table 2 all five models indicate significant autocorrelation at a single lag, evidenced by the auto-regressive term. In addition we see a negative overall trend over the observation period. Intuitively this supports general norms that conversation levels decay with the time since the event. In a few of the models seasonality terms influence the overall conversation rates. We find that rates of hazard-related communication increase on Saturdays compared with other days of the week. We believe that this results from higher time availability during the weekend (versus the work week).

We find that both national news coverage and the water area closed in the Gulf have significant positive effects on hazard-relation conversation on Twitter. In the context of the Deepwater Horizon response, we find that news coverage has a positive association with informal online communication or rumoring – the opposite of what Caplow's theory would have predicted. Combined with the fact that offical press releases seems to have no significant effect on online conversation (it is not present in the best model), this result disconfirms Caplow's theory [7] that official sources of news suppress rumoring activity.Rather, our findings seem to suggest that saliency of the event is important in driving communication dynamics. Both the news coverage and the water area predictors speak to the importance of saliency. Further, theories positing that relevance to decision making behavior (here, captured by water area closures) is positively associated with rumoring are supported.

[4]Due to concerns about spuriousness in the time series used in this analysis, especially the news coverage, we perform a series of checks on the observed relationship. Numerous techniques for transforming the news coverage time series to remove seasonality including differencing and decomposition were tried. In each case the results did not change drastically from those presented here. From these procedures and the models presented we conclude that the news coverage drives a large portion of the online conversation.

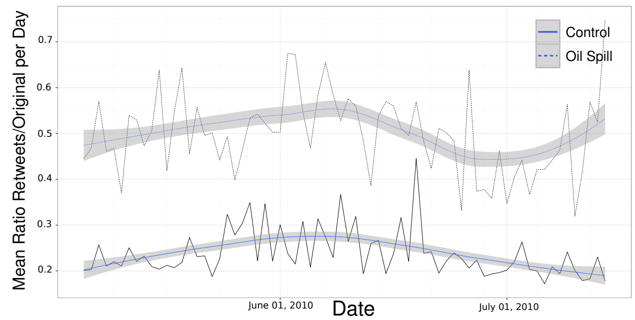

Figure 2. Ratio of original posts to retweets over time for oil spill and control post. Ratio for oil spill content shown as the dashed line (top) and control related content as the sold line (bottom).

Serial Transmission

Throughout literature on rumor, substantial attention is paid to the process by which information is passed from person to person. This *serial transmission* of information is a vital element of rumoring behavior. As such, scholars have been very interested in the factors that influence a person's willingness to transmit rumors [5, 17, 18, 25]. Research suggests that individual characteristics, as well as contextual factors both play a role. In addition, the content of the message itself is important [14]. One of the interesting aspect of the Twitter.com infrastructure is the promotion of serial transmission of information. The built in ease of passing messages along from a friend to one's followers allows for the rapid diffusion of information through the social network. Here we explore the relationship between message content and serial transmission.

In the context of Twitter, the most basic form of serial transmission is *retweeting,* in which one individual reposts a message originally posted by another user [4]. To the extent that serial transmission is present, we will see higher levels of retweeting within the data. To assess this, we consider how the ratio of original content to retweets changes over time. Computing this ratio for each day in the first period of data collection results in the time series of Figure 2. The figure depicts both the raw ratio and a mean smoother for the ratio series (with 95% confidence bands), the latter showing the average trend over time. Comparing the oil spill posts with those from the control keywords, we find that the

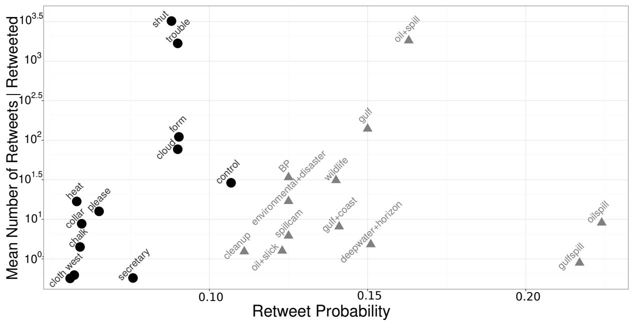

Figure 3. Tweet propagation: probability of being retweeted for any given message is shown on the x-axis. On the y-axis, the mean number of retweets given that a post is retweeted. Posts with oil spill related keywords, in gray, are more likely to be retweeted on average that those containing control keywords, in black.

oil spill posts have a slightly higher proportion of retweets. These results indicate that serial transmission encompasses a greater proportion of overall content in the case of the oil spill, compared to the control sample. This suggests that, in fact, rumor-passing does occur in this environment, and that it occurs more readily in hazard-related communication.

As a further measure, we also take a random sample of 1000 original (i.e., non-retweet) posts for each keyword. We are able to match these posts with all of their retweets, because the retweets necessarily contain the keyword of interest and are thus captured in our sample of tweets. This procedure allows us to estimate two quantities for interest. First we consider the probability that an average post containing keyword X will be retweeted. Additionally, we consider the number of times a post will be retweeted on average, given that it is retweeted at all. A comparison of these two quantities for oil spill and control keywords is seen in Figure 3. Results indicate a clear difference, with oil spill related content having a consistently higher proportion of retweets than control-related posts.

DISCUSSION

Classical rumor theory points to the perceived importance, the level of uncertainty or ambiguity, and the potential to impact decision making behavior as influential in determining the extent of rumoring. As we have demonstrated in this work, in the context of 2010 Deepwater Horizon oil spill media coverage of the event is a major driver of rumoring behavior, supporting importance/saliency theories and disconfirming the presence of information substitution in this case. Relevance to decision making behavior also turns out to be an influential predictor in this case.

We find clear differences in rates of serial transmission, or retweeting, between posts that contain hazard-related keywords and those that contain control keywords. Specifically, posts related to the oil spill are substantially more likely to be passed on to others, although the number of times that they are passed on (given that any passing occurs) is no greater. These findings support previous work suggesting that online communication media are readily used in the context of extreme events, and show moreover that rumor-related behavior online is driven by mechanisms not dissimilar to those governing such behavior in other settings.

While the results presented here look specifically at the Deepwater Horizon oil spill as a case study, the characteristics of this event suggest commonalities to other situations. The Deepwater Horizon disaster was a large-scale, long-term environmental disaster. While the spill itself was localized the impact was far reaching. The response effort brought together organizations and officials from different emergency management sectors as well as different levels of authority. Understanding the dynamics of information exchange in this case could aid in future recovery efforts.

Our work has practical implications for disaster response. Traditionally, emergency managers have been particularly concerned with the *suppression* of rumoring activity [9]; however, in some situations, rumoring supports the transmission of critical information through informal channels – a mechanism that may be used to inform or to correct information that is already on the web. Our results suggest that official news information prompts additional information sharing within the mass public, which could in turn be leveraged to produce faster and more thorough dissemination of emergency information than can be achieved via direct methods. Emergency managers and public officials would thus do well to join the conversation and understand how their participation may increase information sharing. The specific content of information that will be shared is also important during a disaster event. Those ideas that are most relevant or time-sensitive may well be the first to be passed along. Future work may explore how certain keywords are associated with faster response times to serial transmission. We have much to learn from online informal communication and rumoring in the case of extreme events.

ACKNOWLEDGMENTS

This research has been supported by the National Science Foundation under awards CMMI-1031853, CMMI-1031779 and the Office of Naval Research under award N00014-08-1-1015.

ADDITIONAL AUTHORS

REFERENCES

1. Allport, G., and Postman, L. *The psychology of rumor.* Henry Holt, 1947.

2. Anthony, S. Anxiety and rumor. *The Journal of Social Psychology 89* (1973), 91–98.

3. Bordia, P., and DiFonzo, N. Problem solving in social interactions on the Internet: Rumor as social cognition. *Social Psychology Quarterly 67*, 1 (2004), 33–49.

4. Boyd, D., Golder, S., and Lotan, G. Tweet, tweet, retweet: Conversational aspects of retweeting on twitter. In *Proceedings of the 43rd Hawaii International Conference on System Sciences* (2010), 1–10.

5. Buckner, H. T. A theory of rumor transmission. *Public Opinion Quarterly 29*, 1 (1965), 54.

6. Butts, C. T., Sutton, J., and Spiro, E. S. *Hazards, Emergency Response, and Online Informal Communication Project Data*. University of California, Irvine and University of Colorado, Colorado Springs, 2011.

7. Caplow, T. Rumors in war. *Social Forces*, 3 (1947), 298–302.

8. Centola, D., and Macy, M. Complex contagions and the weakness of long ties. *American Journal of Sociology 113*, 3 (2007), 702–34.

9. Danzig, E., Galanter, L., Thayer, P., Relations, I. R. H., and Administration, U. S. F. C. D. *The effects of a threatening rumor on a disaster-stricken community*. National Research Council, 1958.

10. Drabek, T. E. *Disaster in Aisle 13: A Case Study of the Coliseum Explosion at the Indiana Fairgrounds, October 31, 1963*. The Ohio State University Press, Columbus, OH, 1968.

11. Dunbar, R. Groups, gossip, and the evolution of language. *New aspects of human ethology* (1997), 77–89.

12. Erickson, B., Nosanchuk, T., Mostacci, L., and Dalrymple, C. The flow of crisis information as a probe of work relations. *Canadian Journal of Sociology/Cahiers canadiens de sociologie* (1978), 71–87.

13. Festinger, L. Informal social communication. *Psychological review 57*, 5 (Sept. 1950), 271–82.

14. Fragale, A. R., and Heath, C. Evolving informational credentials: the (mis)attribution of believable facts to credible sources. *Personality and social psychology bulletin 30*, 2 (Feb. 2004), 225–36.

15. Gomez Rodriguez, M., Leskovec, J., and Krause, A. Inferring networks of diffusion and influence. In *Proceedings of the 16th ACM SIGKDD international conference on Knowledge discovery and data mining*, ACM (2010), 1019–1028.

16. Granovetter, M. *Getting a Job: A Study of Contacts and Careers*. Harvard University Press, Cambridge, MA, 1974.

17. Jaeger, M. E., Anthony, S., and Rosnow, R. L. Who Hears What from Whom and With What Effect: a Study of Rumor. *Personality and Social Psychology Bulletin 6* (1980), 473–478.

18. Kapferer, J. A mass poisoning rumor in europe. *Public Opinion Quarterly 53*, 4 (1989), 467.

19. Moody, J. The Importance of Relationship Timing for Diffusion. *Social Forces 81* (2002), 25.

20. Oh, O., Kwon, K., and Rao, H. An exploration of social media in extreme events: Rumor theory and twitter during the haiti earthquake 2010. In *Proceedings International Conference on Information Systems (ICIS)* (2010).

21. Onnela, J.-P., Saramäki, J., Hyvönen, J., Szabó, G., Lazer, D., Kaski, K., Kertész, J., and a L Barabási. Structure and tie strengths in mobile communication networks. *Proceedings of the National Academy of Sciences of the United States of America 104*, 18 (May 2007), 7332–6.

22. Perry, R., and Green, M. The role of ethnicity in the emergency decision-making process. *Sociological Inquiry 52*, 4 (1982), 306–334.

23. Rosnow, R., Esposito, J., and Gibney, L. Factors influencing rumor spreading: Replication and extension. *Language & Communication* (1988).

24. Scanlon, T. Post-disaster rumor chains: A case study. *Mass Emergencies 2*, 126 (1977), 22–27.

25. Schachter, S., and Burdick, H. A field experiment on rumor transmission and distortion. *The Journal of Abnormal and Social Psychology 50*, 3 (1955), 363.

26. Shibutani, T. *Improvised News: A Sociological Study of Rumor.* Bobbs-Merrill, 1966.

27. Starbird, K., and Palen, L. Pass It On?: Retweeting in Mass Emergency. In *Proceedings of the 7th International ISCRAM Conference, Seattle, WA*, no. December 2004 (2010), 1–10.

28. Sutton, J. Twittering tennessee: Distributed networks and collaboration following a technological disaster. In *Proceedings of the 7th International ISCRAM Conference–Seattle*, vol. 1 (2010).

29. Sutton, J., Palen, L., and Shklovski, I. Backchannels on the front lines: Emergent uses of social media in the 2007 southern california wildfires. *Crisis*, May (2008).

30. Turner, R. H. Rumor as Intensified Information Seeking: Earthquake Rumors in China and the United States. In *Disasters, Collective Behavior, and Social Organization*, R. R. Dynes and K. J. Tierney, Eds. University of Delaware Press, Newark, NJ, 1994, 244–256.

31. Victor, J. S. *Satanic Panic: The Creation of a Contemporary Legend*. Addison-Wesley, New York, 1993.

32. Walker, C., and Beckerle, C. Effect of anxiety on rumor transmission. *Journal of Social Behavior and Personality 2* (1987), 353–360.

33. Walker, C. J., and Blaine, B. The Virulence of Dread Rumors : A Field Experiment. *Language & Communication 11*, 4 (1991), 291–297.

34. Weenig, M. W., Groenenboom, a. C., and Wilke, H. a. Bad news transmission as a function of the definitiveness of consequences and the relationship between communicator and recipient. *Journal of personality and social psychology 80*, 3 (Mar. 2001), 449–61.

35. Wu, S., Hofman, J., Mason, W., and Watts, D. Who says what to whom on twitter. In *Proceedings of the 20th international conference on World wide web*, ACM (2011), 705–714.