UNIVERSITY OF CALIFORNIA,
IRVINE

A Big Social Media Data Approach to Studying the Psychological Impact of Collective Trauma and Mishaps in Emergency Communications

DISSERTATION

submitted in partial satisfaction of the requirements
for the degree of

DOCTOR OF PHILOSOPHY

in Psychology and Social Behavior

by

Nickolas M. Jones

Dissertation Committee:
Professor Roxane Cohen Silver, Chair
Professor Karen Rook
Associate Professor E. Alison Holman

2018

DEDICATION

To

Mom, Γιαγιά, and Grammpa

TABLE OF CONTENTS

ACKNOWLEDGMENTS

It has been a long road for me (10+ years of education!). I never expected that I would have ended up on this journey—especially such a long one. Yet here I am after a rocky start at a community college after an unfulfilling career in real estate. Wow. My achievements thus far belong to other people in my life: People who have formed the powerful foundation in which I have taken root and grown; people who have encouraged me to be more than I thought I could ever be; people who have helped me to become the best version of myself both personally and intellectually.

The Powerful Womxn
I first want to thank Dr. Roxane Cohen Silver, my wonderful advisor, my second mom, and my champion. Over the past five years, she has guided me, shaped my thinking, encouraged me, challenged me (read: kicked my butt when I needed it), given me support when things seemed bleak, and took an unexpected interest in my life. She is one of the most dedicated mentors I have ever seen, and I am very lucky to have had her at the helm of my graduate school experience. Thank you, Roxy, for believing in me and helping me to see how my interest in technology was perfectly compatible with your area of expertise, an area I believe is one of the most important fields of our time. Thank you for the meetings, the writing feedback, the nominations, and of course, the Persian food with glorious ocean vistas! I will forever be grateful for this time with you, and thankful that you went along with Karen's plan (see below). I am so very proud of the *important* work we have done together and I hope that we will be life-long collaborators and friends.

I would also like to thank Dr. Karen Rook, my co-advisor, who graciously (read: luckily!) leaned over to Roxy (5.5 years ago) during a grad recruitment faculty meeting and convinced her they should both take me on as a student. Karen and I have many academic connections that precede my start at UC Irvine and these connections put me at ease during our interview together in 2013. She was a critical part of my first year as I struggled to navigate graduate school and she was the voice of reason and understanding that a lost boy needed. Throughout my time at UCI, she has made time to meet with me, encouraged my efforts, and assured me that I was doing good work. Karen, I am so grateful for you and I cannot imagine what would have happened to me if you weren't around. Thank you for taking a chance on me. Our time to publish together is on the horizon—I can't wait to talk to about what I have planned when all this is finally done!

I also want to thank Dr. Alison Holman, my badass collaborator and committee member. I am so impressed by you at every turn. I have always loved our time nerding out over statistics during long discussions about potential papers. I have always appreciated your critical eye for methods and theoretical perspectives. Alison, you are such an amazing scholar, a fire, a brilliance. I sure am lucky know you.

I also want to thank Dr. Dana R. Garfin who guided me during my first years in the program and, honestly, still does(!). Dana speaks my language. She has been an inspiration to me in many ways. She has embodied what it means to be a life-long learner and I continue to learn from her.

I also want to thank Dr. Joanne Zinger who pulled my application from the pile of applicants and championed my admission since the very beginning. Getting to know you and working with you over these past years has been such a joy for me. I am so lucky to know such a passionate instructor, dedicated mentor, and seriously good cook. Thank you for taking a chance on me, too. I am so damn grateful. I look forward to many more years of friendship and good times with you!

My friends (collaborators, past and future!)
This journey was not one I took alone. I want to thank Jason Kyler-Yano (my bromantic partner)—I am so grateful for you (I miss thee) and I frequently listen to the recording of our acapella in the PV stairwell. My thanks, my love, my best wishes, and my warmest regards go to Amanda Acevedo, Marie Cross, and Dr. Rebecca Robin Thompson for being incredible friends, dedicated pub mates, and gorgeous brunchettes during this nutty time of life. I love you girls so much and rely on you for being loving and understanding humans. You are beautiful, strong, and brilliant and I am lucky to be counted among your friends.

Special thanks, love, and admiration go to so many others who have helped me on this road: Cyle Metzger (The Grind didn't know what we could do!), Dr. Jeffrey Hunger & Burrel Vann Jr. (my CSUF forever-brothers, since the beginning, I love you), Rupa Jose (queen!), Dr. Adam Fine (a type of tree that shall never be forgotten!), Dr. Sean Wojcik (my idol), Dr. Emily Urban (Urb-E, a nickname that never took), Johnny Hunter (my brother in tech), Sean Goldy (go for the Goldy!), Emma Grisham (Watson), Brittany Page & Jesse Dollemore (constant with love and support for me! Brittany is the best part, though), k3vin.jc, The Silver Stress and Coping Lab (Sarah Redmond, Pauline Lubens, Josiah Sweeting et al.), and my Greek Festival Family (Bill, Linda [RIP], Derbeh, Liz, Seta, Dreena, Jeena, Pat, et al. +50 over the past 20 years).

The fam
Through it all, my family has always been there for me. They continue to encourage me, they keep me going, they believe that I can do all things I set my mind to, and they trust me to become the best I can be. All that I do is for them, and this will never change. I am such a lucky human to have been granted a home full of people who love me as they do. There is no way for me to express how grateful I am for my mother (Alexandra Efremidou, the MOTHERSHIP), my Γιαγιά (Catherine Jones), my Gramps (Michael L. Jones, whose name I have carried since I was 14), and my step-dad (Terry Miller). They have enriched my life in very important ways that have shaped me into the person I am today. Education has always been number one for them and fortunately, so have I. I love you all.

Finally, I want to thank Michael D. Salzman, Cheesemaker Extraordinaire. He is my companion, my support, my love, my mouman. He has been with me since I started this program and has luckily (!) stuck around through the madness. Mou, growing with you these past years has been a wonderful, fulfilling, and warming experience. I am so grateful for all that you give me. We believe in each other and support each other, we laugh like nobody's business, and I look forward to many more years of pastrami reubens, pickles, nightly forbidden bites of cheese (I'm lactose intolerant; yes, we know, it's a waste), endless Greek dancing, and relaxing evenings involving grape/mint hookah with you.

CURRICULUM VITAE

Nickolas M. Jones
Department of Psychology and Social Behavior
University of California, Irvine
nmjones@uci.edu
http://nickolasjones.com

EDUCATION

Psychology and Social Behavior, University of California, Irvine — Ph.D., 2018
Mentor: Roxane Cohen Silver
Specialization: Social Psychology
Minor: Quantitative Methods

Psychology, California State University, Fullerton — B.A. 2009, M.A. 2011

RESEARCH POSITIONS

Research Fellow — 3/2017 – 9/2018
Program of Research in Translational Technology Enabling High Quality Care, Department of Medicine, University of California, Irvine
Dana Mukamel & Dara Sorkin, Directors

Graduate Student Researcher — 1/2017 – 9/2018
National Center for Child Traumatic Stress, UCLA
"Examining the Short- and Long-term Impact of Mass Violence in Communities." Melissa Brymer & Robert Pynoos, PIs
National Institute of Justice, 1/1/17 - 12/31/19

Graduate Student Researcher — 7/2013 – 9/2018
Silver Stress and Coping Lab
Department of Psychological Science, University of California, Irvine
Roxane Cohen Silver, Advisor

QUANTITATIVE & TECHNICAL SKILLS

Data management, visualization, and analysis:

- Hierarchical Linear Modeling
- Structural Equation Modeling
- Logistic Regression
- Generalized Estimating Equations
- Non-linear Models / Spline-regression
- Data Imputation Techniques

Statistical packages:
R, Stata 14, SPSS

Large-scale project design, implementation, and management:

- Big data analytics
- Web-scraping Twitter data
- Smartphone data collection
- Experience sampling
- Natural language processing techniques

Coding experience:
R, Python, HTML, CSS, JavaScript, PHP

HONORS AND AWARDS

Martha Newkirk Award for Excellence in Research by a Graduate Student, School of Social Ecology, UCI ($1,000)	2018
Graduate Student Mentoring Award, UCI	2018,15,14
SPSP travel award to attend the Society of Australasian Social Psychologists' Summer Conference at the Australian National University, New South Wales, Australia ($750)	2018
SPSP Media Psychology Pre-Conference Registration Award	2017
BA awarded *cum laude*, California State University Fullerton	2009

INTRAMURAL RESEARCH GRANT

Chair Research Award, University of California Irvine
"Patient decision making in an era of unlimited information and misinformation on the Internet"
Emilie Chow, Dana B. Mukamel, Nickolas M. Jones, PIs
Department of Medicine, 7/2017 - 12/2018. Funding: $25,000

PUBLICATIONS

Jones, N. M., Thompson, R. R., Dunkel Schetter, C., & Silver, R. C. (2017). Distress and rumor exposure on social media during a campus lockdown. *Proceedings of the National Academy of Sciences of the USA*, *144*, 11663-11668 (published ahead of print October 17, 2017). doi: 10.1073/pnas.1708518114

Jones, N. M., Wojcik, S. P., Sweeting, J., & Silver, R. C. (2016). Tweeting negative emotion: An investigation of Twitter data in the aftermath of violence on college campuses. *Psychological Methods*, *21*, 526-541. doi: 10.1037/met0000099

Jones, N. M., Garfin, D. R., Holman, E. A., & Silver, R. C. (2016). Media use and exposure to graphic content in the week following the Boston Marathon bombings. *American Journal of Community Psychology*, *58*, 47-59. doi:10.1002/ajcp.12073

Guadagno, R. E., **Jones, N. M.**, Kimbrough, A. K., & Mattu, A. (2016). Translating social media psychological research. *Translational Issues in Psychological Science*, *2*, 213-215. doi: 10.1037/tps0000087

JOURNAL SPECIAL ISSUE

Jones, N. M. (Guest Associate Editor) (2016). *Translational Issues in Psychological Science: Special Issue on Social Media*, Rosanna Guadagno, Special Issue Editor.

MANUSCRIPTS UNDER REVIEW OR IN PREPARATION

Jones, N. M., & Silver, R. C. This is not a drill: Anxiety on Twitter following the Hawaii false missile alert.

Jones, N. M., Brymer, M., & Silver, R. C. Using big data to study the impact of mass violence: Opportunities for the traumatic stress field.

Redmond, S. A., **Jones, N. M.**, Holman, E. A., & Silver, R. C. Who watches an ISIS beheading? A longitudinal mixed-methods study of a national U.S. sample.

Relihan, D., **Jones, N. M.**, Holman, E. A., & Silver, R. C. Social identity, media, and the stress of terrorism.

Thompson, R. R., **Jones, N. M.**, Holman, E. A., & Silver, R. C. Media exposure to mass violence events can fuel a cycle of distress.

INVITED TALKS / CONFERENCE PRESENTATIONS

Jones, N. M., & Silver, R. C. (2017, December). *Distress and rumor exposure on social media during a campus lockdown*. Recommendations presented to the Chancellor's Cabinet of the University of California, Irvine, Irvine, CA.

Jones, N. M. (2017, January). *Uncertainty during a campus lockdown: An analysis of rumor communication and distress*. Data blitz presented at the annual meeting of the Society for Personality and Social Psychology, San Antonio, TX.

Jones, N. M. (2016, March). *Using social media to answer psychological questions.* Presented at Saddleback Community College, Mission Viejo, CA

Beals, K. P., Ghavami, N., **Jones, N. M.**, & Conlon, C. (2013, April). *Diversity in Research Workshop*. Panelist presentation at the annual meeting of the Western Psychological Association, Reno, NV.

CONFERENCE POSTERS

**Student mentee*

*Philbrick, G. M., *Wilder, R. A., *Le, K., **Jones, N. M.** (2018, May). *Social support, life satisfaction, and patterns of social media use*. Poster presented at the annual Undergraduate Research Opportunity Program conference, Irvine, CA.

*Wilder, R. A., *Philbrick, G. M., *Le, K., **Jones, N. M.** (2018, May). *Social media application use and belongingness*. Poster presented at the annual Undergraduate Research Opportunity Program conference, Irvine, CA.

Redmond, S. A, **Jones, N. M.**, Holman E. A., & Silver, R. C. (2017, November). *Who watches an ISIS beheading? A longitudinal mixed methods study of a national U.S. sample*. Poster presented at the annual conference of the International Society for Traumatic Stress Studies, Chicago, IL.

Jones, N. M., Thompson, R. R., Dunkel Schetter, C., & Silver, R. C. (2017, November). *Distress and rumor exposure on social media during a campus lockdown*. Poster presented at the Psychology of Technology Conference, Berkeley, CA.

Jones N. M., Wojcik, S. P., Sweeting J., & Silver, R. C. (2015, May). Tweeting negative emotion: Affective signal in community samples of Twitter users in the aftermath of the 2014 Isla Vista killings. Poster presented at the UC Irvine Data Science Initiative's Afternoon of Data Science, Irvine, CA.

Jones, N. M., Garfin, D. R., Holman, E. A., & Silver, R. C. (2015, February). *Media use and exposure to graphic content in the week following the Boston Marathon bombings.* Poster presented at the annual meeting of the Society for Personality and Social Psychology, Long Beach, CA.

*Page, B., & **Jones, N. M.** (2014, February). *Can moral foundations predict perceptions of cheating behaviors? The role of purity in perceiving deception.* Poster presented at the annual meeting of the Society for Personality and Social Psychology, Austin, TX.

Jones, N. M., Warstadt, M., Schultz, R., McKinnis, D. L., & Beals, K. P. (2010, January). Do LGB individuals distance themselves from the majority community following the passage of anti-gay referenda? A quantitative and qualitative look. Poster presented at the annual meeting of the Society for Personality and Social Psychology, Las Vegas, NV.

McKinnis, D. L., Fadakar, G., Ortiz, C., **Jones, N. M.**, & Beals, K. P. (2010, January). Non-recognition of same-sex unions and relationship commitment: A look at the effects of the California Supreme Court's decision to uphold Proposition 8. Poster presented at the annual meeting of the Society for Personality and Social Psychology, Las Vegas, NV.

McKinnis, D. L., **Jones, N. M.**, Schultz, R., Fadakar, G., Ortiz, C., Martinez, R., & Beals, K. P. (2009, April). *Effects of Proposition 8 on gay, lesbian, and bisexual individuals*. Poster presented at the annual meeting of the Western Psychological Association, Portland, OR.

MEDIA COVERAGE

Huffington Post, Mashable.com, Southern California Public Radio (KPCC) 2017

Arizona Public Radio (KNAU) 2016

OTHER PROFESSIONAL SERVICE

Student Member, Faculty Search Committee, Department of Psychological Science, University of California, Irvine 11/2016 - 3/2017

Ad Hoc Reviewer: Emotion, Translational Issues in Psychological Science

MEMBERSHIP IN PROFESSIONAL ASSOCIATIONS

Society for Personality and Social Psychology

TEACHING EXPERIENCE

Instructor, School of Social Ecology, University of California, Irvine — Summer 2017
Stats Bootcamp for incoming graduate students

Graduate Teaching Assistant, Department of Psychological Science, University of California, Irvine. *Classes*: Quantitative methods in psychology (graduate-level statistics series), Applied statistics for psychological research, Research methods, Social relationships, Psychology fundamentals — 10/2013 - 12/2016

Part-time Faculty Member, Department of Psychology, California State University, Fullerton. *Classes*: Computer applications in psychology, Human sexual behavior, Research methods — 8/2011 - 8/2013

REFERENCES

Roxane Cohen Silver, Ph.D.
Professor of Psychology, Public Health, & Medicine
Department of Psychological Science
University of California, Irvine
rsilver@uci.edu

Karen Rook, Ph.D.
Professor
Department of Psychological Science
University of California, Irvine
ksrook@uci.edu

Dana Mukamel, Ph.D.
Professor of Medicine, Public Health, & Nursing
Director, iTEQC Research Program
Department of Medicine
University of California, Irvine
dmukamel@uci.edu

Melissa Brymer, Ph.D.
Director of Terrorism and Disaster Programs
UCLA–Duke National Center for Child Traumatic Stress
MBrymer@mednet.ucla.edu

ABSTRACT OF THE DISSERTATION

A Big Social Media Data Approach to Studying the Psychological Impact of Collective Trauma and Mishaps in Emergency Communications

By

Nickolas M. Jones

Doctor of Philosophy in Psychology and Social Behavior

University of California, Irvine, 2018

Professor Roxane Cohen Silver, Chair

This dissertation examines how “Big Data” collected from Twitter can be used to understand how collective traumas (e.g., mass violence, campus shootings, and large-scale life-threatening warnings) impact individuals and communities. Prior research suggests that when collective traumas occur, people use social media channels (e.g., Twitter) to express psychological discomfort and seek critical updates, especially when updates from emergency management officials are lacking. Study 1 demonstrates the methodological value of mining locally-generated Twitter data after the 2015 San Bernardino terrorist attack. Results show that negative emotion increased 6.2% on the day of the attack and remained elevated for several days after. Studies 2a and 2b combine the methodological strengths of survey and Twitter data to understand distress responses, rumor exposure, generation, and transmission during a protracted lockdown event at a major university. Among about 3,400 students trapped in the lockdown, results indicate that a) rumor exposure and social media channel use during the lockdown were each associated with distress, b) students who trusted social media for critical updates reported increased acute stress, and c) exposure to rumors was associated with consulting Twitter for critical updates. In Study 2b, rumor generation and transmission occurred in the 90-minute gap in which no updates were transmitted to the campus during the lockdown. Study 3 exclusively uses Twitter data to explore the community- and individual-level impact of exposure to a life-threatening, but false, ballistic missile alert sent to smartphones of all residents and visitors in Hawaii in January 2015. Analyses at three time-scales reveal that a) anxiety expressed on Twitter increased 4.6% on the day of the false alert and remained elevated for at least two days, b) anxiety increased 3.4% every 15-minutes during the 38-minute alert period (incubation of threat), and c) users who expressed low pre-alert anxiety on Twitter exhibited the highest increase in anxiety expression at the onset of the alert (9.5%) and their anxiety remained elevated for 41 hours, compared to medium (5.6%; 23 hours) and high pre-alert anxiety users (-10.5%; 0 hours). The benefits of this big data approach for advancing knowledge and theory across all three studies are discussed.

Chapter 1: Introduction

Introduction

The increasing use of technologies that provide unfettered access to new media platforms for expressing emotions and seeking information after a collective trauma (e.g., terrorist attack, school shooting) offers researchers new opportunities to study how disasters impact people and the communities in which they live. For several decades, the collective experience of natural (e.g., hurricanes, earthquakes) and anthropogenic (e.g., mass shootings, terrorist attacks) traumas has been the subject of investigation by researchers interested in the clinical implications of exposure to such events. Across dozens of studies of collective traumas, researchers find that some people who experience these events exhibit negative outcomes like psychiatric problems (e.g., anxiety, depression), criterion-based conditions (e.g., posttraumatic stress disorder [PTSD]), health problems (e.g., medical conditions), and functional impairment (see Norris et al., 2002, for a review).

However, it is only recently that social scientists have sought to understand these events employing multi-method and ecological research approaches, due primarily to the ubiquity and availability of "big data" on social media. These data enable researchers to study collective traumas with sample sizes and on time scales not economically feasible in the past and with unprecedented granularity. As discussed in this dissertation, big social media data derived from new media platforms (e.g., Twitter) offer methodological advantages that are critical for exploring the impact of large-scale traumatic events that may be inaccessible to the traditional post-disaster researcher.

Using big social media data to study collective traumas

The rigorous study of the impact of collective traumas is difficult at best (Silver, 2004). Standard experimental methods that use random assignment to understand causal relationships in the context of collective traumas are impossible due to the unpredictable nature of large-scale traumatic events and ethical issues that preclude manipulating exposure. Beyond method, there are other prohibitive challenges that come into play (Silver, 2004). For example, when a collective trauma occurs, it may be difficult to quickly obtain ethics board approval, funding, and access to the affected population. Such barriers are implicated in the loss of insight into an event's impact (North & Pfefferbaum, 2002). Moreover, if a researcher waits too long to enter the field, individuals under study may forget symptoms experienced closer to the event, thereby biasing results (North, Smith, & Spitznagel, 1997).

Ideally, researchers should use longitudinal quasi-experimental designs (e.g., interrupted time-series) with representative samples to make inferences about an event's immediate and long-term psychological effects. In practice, however, such an endeavor can be very costly at best and impossible at worst. For example, most studies of collective traumas begin after the fact and lack any pre-event assessments, thereby precluding a researcher's ability to make unambiguous claims about an event's impact. Infrequently, however, preferred designs have been used by researchers who were already in the field conducting research on some other phenomenon and, after the onset of a collective trauma, altered their focus to study the event's impact (e.g., Bravo, Rubio-Stipec, Canino, Woodbury, & Ribera, 1990). Although this offers the methodological benefit of having pre-event data, researchers in this circumstance rarely have a control group against which to compare post-disaster trends in the affected community (Norris et al., 2002).

Recently, some researchers have avoided these prohibitive methodological challenges by turning to social media data on Twitter to study the impact of unexpected or ongoing community traumas (De Choudhury, Monroy-Hernandez, & Mark, 2014; Doré, Ort, Braverman, & Ochsner, 2015; Glasgow, Fink, & Boyd-Graber, 2014; Jones, Wojcik, Sweeting, & Silver, 2016). Twitter is a widely used social media platform that allows a user to post status updates (tweets), of 280 characters or less, to his or her timeline. These posts are instantly shared across a personal network of followers (individuals who subscribe to a user's account), who can then reply to the tweets or post them to their own timelines (retweeting), thus sharing tweets across many networks. Fortunately, researchers can freely connect to Twitter, download public tweets from any users who have public accounts, and analyze the content of those tweets.

Although the 280-character limit of tweet content may seem paltry, tweets contain a wealth of information about what users think and feel and provide insights into their online behavior in a variety of contexts. This knowledge-generating bounty is accessible using text-analysis techniques developed outside of psychology. Specifically, the systematic study of words has been a staple of the field of linguistics for decades (Biber, Conrad, & Reppen, 1998). It is only recently that psychologists have entered this arena, riding on the coat-tails of computer scientists who have been steadily developing the software to run computational analyses of text as computing power has increased dramatically over the past several years. Most researchers exploring tweets to study communities following trauma have relied on the text analysis platform developed by Pennebaker and his colleagues (Linguistic Inquiry and Word Counter [LIWC]; see Pennebaker, Mehl, & Niederhoffer, 2003; Tausczik & Pennebaker, 2010), which provides word dictionaries that have been empirically shown to map onto psychological constructs (e.g., negative emotion). Words used in tweets are checked against validated dictionaries, and researchers infer psychological meaning from a tweet based on whether dictionary words are present in the tweet's content.

A methodological advantage of analyzing Twitter data is that tweets are archival and time-stamped, allowing for time-series analyses of their content. The flexibility of using these data, in aggregate, to understand the time-course of a trauma's psychological impact on communities has been demonstrated across several studies. For example, this form of data has proven useful for examining changes in emotion over time after the Sandy Hook Elementary School massacre in Newtown, CT (Doré et al., 2015), as well as in an analysis of how individuals used death-related words after this event (Glasgow et al., 2014). These kinds of data have also been used to demonstrate long-term negative emotion expression in Mexican communities exposed to constant violence during the Mexican Drug War (De Choudhury et al., 2014) and to demonstrate negative emotion expression before, during, and after several incidents of college campus violence (Jones et al., 2016).

Traditional and new media after a disaster

When a large-scale traumatic event unfolds, whether local or national, the news media are often the first point of contact for people interested in finding out what happened. Indeed, researchers have long recognized the importance of information-seeking in the aftermath of collective traumas. For example, using Uncertainty Reduction Theory (Berger & Calabrese, 1975) as a lens, researchers who studied media use after a disaster (e.g., 9/11 terrorist attacks) argued that information-seeking behavior is a common response among individuals and serves as a means by which individuals can alleviate anxiety induced by the resulting uncertainty of the disaster context (Heath & Gay, 1997; Lachlan, Spence, & Seeger, 2009; Seeger, Sellnow, & Ulmer, 2003). Recently, researchers have identified distinct patterns of post-disaster media use across traditional (e.g., newspaper, radio, print) and new (e.g., social media) media channels. For example, researchers found that age was negatively associated with new media use, and positively associated with traditional media use (Jones, Garfin, Holman, & Silver, 2016).

Moreover, individuals reported using traditional and new media complementarily in the aftermath of the Boston Marathon bombings and the propensity to do so was stronger for younger individuals. Although these findings suggest that older people are less likely to report using new media after a collective trauma, a recent Pew Research Center (2017) study reports that 85% of U.S. adults use their smartphones to obtain news, and adults over 65 are driving this increasing trend. In addition, roughly two-thirds of adults over 65 use mobile news platforms.

Although seeking information is theorized to serve a palliative function after a large-scale traumatic event, research has also found that being indirectly exposed to a collective trauma via news coverage has the potential to transmit distress to those well outside the directly exposed community. Large studies of collective traumas using representative U.S. samples have demonstrated that repeated indirect, media-based exposure to collective traumas is associated with trauma-related symptoms. For example, after controlling for pre-event media habits and preexisting mental health conditions, Holman and colleagues (2014) found that six or more hours of media use after the Boston Marathon bombings was associated with higher acute stress symptomatology (DSM-IV; American Psychiatric Association, 2000) than was direct exposure to the bombings. Similar media-based distress response patterns have been found after several collective traumas in the U.S., including the Oklahoma City Bombing (Pfefferbaum et al., 2001), the 1990 Gulf War (Cantor, Mares, & Oliver, 1993), and the 9/11 terrorist attacks (Ahern et al., 2002; Schlenger et al., 2002; Schuster et al., 2001; Silver et al., 2013; Silver, Holman, McIntosh, Poulin, & Gil-Rivas, 2002).

Ambiguous threats and their alerts

In some cases, large-scale traumas (e.g., active shooter and adverse weather events) are accompanied by warnings and alerts transmitted by local or national emergency management agencies before and/or during an event. Such transmissions have the principal aim of informing members of the public of imminent threats so that individuals may act to secure personal safety and property. Emergency management agencies and their ability to effectively communicate with populations at risk of threatening events can play an important role in saving lives (Rodríguez, Díaz, Santos, & Aguirre, 2007).

However, when emergency alerts occur without follow up information or are difficult to verify, their transmission may result in increasing the ambiguity of disaster situations and anxiety felt by those caught in their path (Breznitz, 1984). Moreover, when both emergency management agencies and the media lack critical information to help those affected clarify the timing and extent of a threat, people increasingly turn to unofficial channels (e.g., social media) to piece together what is happening (Hughes & Palen, 2009; Mazer et al., 2015; Palen & Anderson, 2016; Palen & Liu, 2007; Spiro et al., 2012). Researchers acknowledge that social media platforms have been instrumental in providing critical emergency updates during some disasters (Palen & Hughes, 2018). However, this increasing reliance on social media by members of the public is concerning because during some disasters, these platforms have been hotbeds for the proliferation of misinformation and rumors (Starbird, Maddock, Orand, Achterman, & Mason, 2014).

Most studies of disaster communication are concerned with human behavior after receiving disaster alerts (e.g., evacuation behavior; Ripberger et al., 2015), death tolls resulting from mishaps in communications (e.g., false alarms; Simmons & Sutter, 2009), or the organizational implications (e.g., credibility loss; Dow & Cutter, 1998) of transmitting alerts to threatened communities. To date, however, no studies have assessed the psychological impact of disaster communications gone wrong.

Plan for the dissertation

Until recently, relatively few studies have attempted to use Twitter data to study the impact of collective trauma, despite the promise these data hold. Additionally, none to date have used these data to study the psychological reactions to emergency alerts nor the lasting psychological implications of their receipt. This dissertation describes the results of three studies that use big social media data before, during, and after collective traumas to answer theoretical questions of interest. The first chapter (Chapter 2) outlines the process and utility of collecting Twitter data to compare the impact of a collective trauma (terrorist attack) against a control community. Chapter 3 highlights the value of using these data in conjunction with traditionally-collected survey research data to explore the association between rumor exposure and distress after a shooting and lockdown at a major U.S. university. Chapter 4 explores how Twitter data can be used to group individual users in psychologically meaningful ways to analyze anxiety responses on Twitter during and following a false ballistic missile alert transmitted to an entire U.S. state. In each case, Twitter data are coded for psychological information embedded in tweets and analyzed at varying time-scales to showcase the granularity of analysis possible with these data. Implications of this method for studying collective traumas and the findings from these studies are discussed in Chapter 5.

References

Ahern, J., Galea, S., Resnick, H., Kilpatrick, D., Bucuvalas, M., Gold, J., & Vlahov, D. (2002). Television images and psychological symptoms after the September 11 terrorist attacks. *Psychiatry: Interpersonal and Biological Processes*, *65*, 289–300. https://dx.doi.org/10.1521/psyc.65.4.289.20240

American Psychiatric Association. (2000). *Diagnostic and statistical manual of mental disorders: DSM-IV-TR*. Washington, DC: American Psychiatric Association.

Berger, C. R., & Calabrese, R. J. (1975). Some explorations in initial interaction and beyond: Toward a developmental theory of interpersonal communication. *Human Communication Research*, 99–112. https://dx.doi.org/10.1111/j.1468-2958.1975.tb00258.x

Biber, D., Conrad, S., & Reppen, R. (1998). *Corpus linguistics: Investigating language structure and use*. Cambridge, UK: Cambridge University Press. https://doi.org/10.1017/CBO9780511804489

Bravo, M., Rubio-Stipec, M., Canino, G. J., Woodbury, M. A., & Ribera, J. C. (1990). The psychological sequelae of disaster stress prospectively and retrospectively evaluated. *American Journal of Community Psychology*, *18*, 661–680.

Breznitz, S. (1984). *Cry wolf: The psychology of false alarms*. Hillsdale, NJ: Lawrence Erlbaum Associates, Inc.

Cantor, J., Mares, M.-L., & Oliver, M. B. (1993). Parents' and children's emotional reactions to TV coverage of the Gulf War. In B. S. Greenberg & W. Gantz (Eds.), *Desert Storm and the mass media* (pp. 325–340). Cresskill, NJ: Hampton Press.

De Choudhury, M., Monroy-Hernandez, A., & Mark, G. (2014). Narco emotions: Affect and desensitization in social media during the Mexican drug war. In *Proceedings of the*

SIGCHI conference on Human Factors in Computing Systems (pp. 3563–3572). New York, NY: ACM. https://dx.doi.org/10.1145/2556288.2557197

Doré, B., Ort, L., Braverman, O., & Ochsner, K. N. (2015). Sadness shifts to anxiety over time and distance from the national tragedy in Newtown, Connecticut. *Psychological Science*, *26*, 363–373. https://dx.doi.org/10.1177/0956797614562218

Dow, K., & Cutter, S. L. (1998). Crying wolf: Repeat responses to hurricane evacuation orders. *Coastal Management*, *26*, 237–252. https://dx.doi.org/10.1080/08920759809362356

Glasgow, K., Fink, C., & Boyd-Graber, J. L. (2014). “Our grief is unspeakable”: Automatically measuring the community impact of a tragedy. In *Proceedings of the 8th International AAAI Conference on Weblogs and Social Media*. Palo Alto, CA: AAAI Press. Retrieved from http://www.aaai.org/ocs/index.php/ICWSM/ICWSM14/paper/view/8086

Heath, R. L., & Gay, C. D. (1997). Risk communication involvement, uncertainty, and control’s effect on information scanning and monitoring by expert stakeholders. *Management Communication Quarterly*, *10*, 342–372. https://dx.doi.org/10.1177/0893318997010003004

Holman, E. A., Garfin, D. R., & Silver, R. C. (2014). Media’s role in broadcasting acute stress following the Boston Marathon bombings. *Proceedings of the National Academy of Sciences*, *111*, 93–98. https://dx.doi.org/10.1073/pnas.1316265110

Hughes, A. L., & Palen, L. (2009). Twitter adoption and use in mass convergence and emergency events. *International Journal of Emergency Management*, *6*, 248–260.

Jones, N. M., Garfin, D. R., Holman, E. A., & Silver, R. C. (2016). Media use and exposure to graphic content in the week following the Boston Marathon bombings. *American Journal of Community Psychology*, *58*, 47–59. https://dx.doi.org/10.1002/ajcp.12073

Jones, N. M., Wojcik, S. P., Sweeting, J., & Silver, R. C. (2016). Tweeting negative emotion: An investigation of Twitter data in the aftermath of violence on college campuses. *Psychological Methods*, *21*, 526–541. https://dx.doi.org/10.1037/met0000099

Lachlan, K. A., Spence, P. R., & Seeger, M. (2009). Terrorist attacks and uncertainty reduction: Media use after September 11. *Behavioral Sciences of Terrorism and Political Aggression*, *1*, 101–110. https://dx.doi.org/10.1080/19434470902771683

Mazer, J. P., Thompson, B., Cherry, J., Russell, M., Payne, H. J., Kirby, E. G., & Pfohl, W. (2015). Communication in the face of a school crisis: Examining the volume and content of social media mentions during active shooter incidents. *Computers in Human Behavior*, *53*, 238–248. https://dx.doi.org/10.1016/j.chb.2015.06.040

Norris, F. H., Friedman, M. J., Watson, P. J., Byrne, C. M., Diaz, E., & Kaniasty, K. (2002). 60,000 disaster victims speak: Part I. An empirical review of the empirical literature, 1981–2001. *Psychiatry: Interpersonal and Biological Processes*, *65*, 207–239. https://dx.doi.org/10.1521/psyc.65.3.207.20173

North, C. S., & Pfefferbaum, B. (2002). Research on the mental health effects of terrorism. *Journal of the American Medical Association*, *288*, 633–636. https://dx.doi.org/10.1001/jama.288.5.633

North, C. S., Smith, E. M., & Spitznagel, E. L. (1997). One-year follow-up of survivors of a mass shooting. *American Journal of Psychiatry*, *154*, 1696–1702.

Palen, L., & Anderson, K. M. (2016). Crisis informatics—New data for extraordinary times. *Science, 353*, 224–225. https://dx.doi.org/10.1126/science.aag2579

Palen, L., & Hughes, A. L. (2018). Social media in disaster communication. In H. Rodríguez, W. Donner, & J. E. Trainor (Eds.), *Handbook of disaster research* (pp. 497–518). Cham, Switzerland: Springer International Publishing. https://doi.org/10.1007/978-3-319-63254-4_24

Palen, L., & Liu, S. B. (2007). Citizen communications in crisis: Anticipating a future of ICT-supported public participation. In *Proceedings of the SIGCHI conference on Human factors in computing systems* (pp. 727–736). New York: ACM Press. https://dx.doi.org/10.1145/1240624.1240736

Pennebaker, J. W., Mehl, M. R., & Niederhoffer, K. G. (2003). Psychological aspects of natural language use: Our words, our selves. *Annual Review of Psychology*, *54*, 547–577. https://dx.doi.org/10.1146/annurev.psych.54.101601.145041

Pew Research Center. (2017). *Growth in mobile news use driven by older adults*. Retrieved from http://www.pewresearch.org/fact-tank/2017/06/12/growth-in-mobile-news-use-driven-by-older-adults/

Pfefferbaum, B., Nixon, S. J., Tivis, R. D., Doughty, D. E., Pynoos, R. S., Gurwitch, R. H., & Foy, D. W. (2001). Television exposure in children after a terrorist incident. *Psychiatry: Interpersonal and Biological Processes*, *64*, 202–211. https://dx.doi.org/10.1521/psyc.64.3.202.18462

Ripberger, J. T., Silva, C. L., Jenkins-Smith, H. C., Carlson, D. E., James, M., & Herron, K. G. (2015). False alarms and missed events: The impact and origins of perceived inaccuracy in tornado warning systems. *Risk Analysis*, *35*, 44–56. https://dx.doi.org/10.1111/risa.12262

Rodríguez, H., Díaz, W., Santos, J. M., & Aguirre, B. E. (2007). Communicating risk and uncertainty: Science, technology, and disasters at the crossroads. In H. Rodríguez, E. L. Quarantelli, & R. R. Dynes (Eds.), *Handbook of disaster research* (pp. 476–488). New

York, NY: Springer. https://doi.org/10.1007/978-0-387-32353-4_29

Schlenger, W. E., Caddell, J. M., Ebert, L., Jordan, B. K., Rourke, K. M., Wilson, D., … Kulka, R. A. (2002). Psychological reactions to terrorist attacks: Findings from the National Study of Americans' Reactions to September 11. *Journal of the American Medical Association*, *288*, 581–588. https://dx.doi.org/10.1001/jama.288.5.581

Schuster, M. A., Stein, B. D., Jaycox, L. H., Collins, R. L., Marshall, G. N., Elliott, M. N., … Berry, S. H. (2001). A national survey of stress reactions after the September 11, 2001, terrorist attacks. *New England Journal of Medicine*, *345*, 1507–1512. https://dx.doi.org/10.1056/NEJM200111153452024

Seeger, M. W., Sellnow, T. L., & Ulmer, R. R. (2003). *Communication and organizational crisis*. Westport, CT: Greenwood Publishing Group.

Silver, R. C. (2004). Conducting research after the 9/11 terrorist attacks: Challenges and results. *Families, Systems, & Health*, *22*, 47–51. https://dx.doi.org/10.1037/1091-7527.22.1.47

Silver, R. C., Holman, E. A., Andersen, J. P., Poulin, M., McIntosh, D. N., & Gil-Rivas, V. (2013). Mental- and physical-health effects of acute exposure to media images of the September 11, 2001, attacks and the Iraq War. *Psychological Science*, *24*, 1623–1634. https://dx.doi.org/10.1177/0956797612460406

Silver, R. C., Holman, E. A., McIntosh, D. N., Poulin, M., & Gil-Rivas, V. (2002). Nationwide longitudinal study of psychological responses to September 11. *Journal of the American Medical Association*, *288*, 1235–1244. https://dx.doi.org/10.1001/jama.288.10.1235

Simmons, K. M., & Sutter, D. (2009). False alarms, tornado warnings, and tornado casualties. *Weather, Climate, and Society*, *1*, 38–53. https://dx.doi.org/10.1175/2009WCAS1005.1

Spiro, E. S., Fitzhugh, S., Sutton, J., Pierski, N., Greczek, M., & Butts, C. T. (2012). Rumoring during extreme events: A case study of Deepwater Horizon 2010. In *Proceedings of the 4th Annual ACM Web Science Conference* (pp. 275–283). New York, NY: ACM. https://dx.doi.org/10.1145/2380718.2380754

Starbird, K., Maddock, J., Orand, M., Achterman, P., & Mason, R. M. (2014). Rumors, false flags, and digital vigilantes: Misinformation on Twitter after the 2013 Boston Marathon bombing. In *iConference 2014 Proceedings* (pp. 654–662). Grandville, MI: iSchools. https://dx.oi.org/10.9776/14308

Tausczik, Y. R., & Pennebaker, J. W. (2010). The psychological meaning of words: LIWC and computerized text analysis methods. *Journal of Language and Social Psychology*, *29*, 24–54. https://dx.doi.org/10.1177/0261927X09351676

Chapter 2: Using big data to study the impact of mass violence: Opportunities for the traumatic stress field

Abstract

Studying the community impact of mass violence using a Big Data approach from social media data (e.g., Twitter) offers traumatic stress researchers an unprecedented opportunity to study and clarify theoretical assumptions using large-scale, observational ecologically valid data. We describe challenges and benefits of working with Twitter data and briefly review studies that used Twitter data to explore community responses to mass violence. We then demonstrate the use of Twitter data to examine community responses to a specific event: The 2015 San Bernardino terrorist attack in which 14 people were killed and 22 were wounded. In a 6-week window around this attack, we evaluated the time-course of community-level negative emotion. We downloaded 1.16 million tweets representing 26,030 users from San Bernardino, CA and a matched control community, Stockton, CA. All tweets were coded in R using the Linguistic Inquiry and Word Count (LIWC) negative emotion dictionary (Pennebaker et al., 2015). A piecewise regression technique with a discontinuity analysis was used to evaluate pre- and post-event trajectories of negative emotion across the study window. Controlling for within-user variability, negative emotion increased by 6.2% (standardized $b = 0.182$, $SE = .014$, $p < .001$) in San Bernardino on the day of the attack and remained elevated for five days; no elevation was observed in Stockton. We discuss how data-driven text analytic techniques are useful for exploring Twitter content generated after collective traumas and describe challenges and opportunities accompanying analyses of social media data to understand the impact of mass violence on affected populations.

Introduction

The rising interest in "big data" analytics among social scientists (National Research Council, 2013) has led to the development and refinement of computational tools for analyzing complex cross-sectional and longitudinal data. This has been particularly the case for the analysis of large scale corpora of social media text. Easy-to-use packages in open-source statistical programming languages (e.g., R, Python) now allow social scientists to engage with their data in ways that were once relegated to individuals with extensive training in programming languages (i.e., computer scientists). Indeed, a quick internet search yields many text- and video-based tutorials for extracting information embedded within social media data. Big, archival data (especially those generated on social media platforms like Twitter) and the techniques used for analyzing them come in many shapes, sizes, and flavors. This variety translates to a robust toolkit with which trauma researchers can answer important questions relevant to the field. In gaining expertise with these tools, there is an opportunity for new and seasoned traumatic stress researchers to generate new knowledge about how people within affected communities respond to and cope with mass violence in its immediate aftermath and over time.

In this paper, we provide a brief introduction to methodological considerations of working with Twitter data. We highlight studies at the intersection of technology and traumatic stress across several disciplines to demonstrate the value of this approach for studying mass violence. We then employ this approach, using a strong quasi-experimental design, to study an incident of mass violence (i.e., the December 2015 terrorist attack in San Bernardino, CA). Finally, we end with a discussion of how big social media data can be used to develop a rich understanding of the impact of collective traumas and advance knowledge and theory in the field.

Working with big data

As many traumatic stress researchers know, studying the impact of collective trauma using rigorous study designs with traditional research methods is challenging (Silver, 2004). Mass violence events are typically unexpected and occur without warning. This precludes a researcher's ability to collect pre-event data against which to compare post-event trends. Moreover, getting into the field quickly after an event can be difficult because of delays in ethics board approval or acquiring funding and staff to conduct research (Silver, 2004; Steinberg, Brymer, Steinberg, & PFefferbaum, 2006). Finally, studies of large-scale traumas do not typically allow for the formation of control groups, a hallmark of true experiments and an important feature of any strong quasi-experimental design (Shadish, Cook, & Campbell, 2002).

Recently researchers in informatics, sociology, communications, and psychology have circumvented these challenges, tapping Twitter data across the globe to study the impact of collective traumas. The data analyzed in these studies are comprised of "tweets," which are short messages written by Twitter users that containing opinions, random thoughts, and emotions, all of which can be shared publicly if the user desires. The primary analytic units of tweets are the words people use to write them, and many social scientists use an automated means of capturing the emotion, mood, and cognition expressed in these words. The most popular method used to accomplish this goal is the Linguistic Inquiry and Word Count (LIWC; Pennebaker, Booth, Boyd, & Frances, 2015) program, which is comprised of empirically-validated dictionaries that contain words related to several psychological constructs (e.g., emotion, cognition) and linguistic features related to psychological states (e.g., pronouns). Words in these dictionaries are then automatically compared to the words in individual tweets and the software produces a count of word-matches that can be analyzed in various ways depending on the research question.

There are important features of Twitter data that make them ideally suited for studying collective trauma. First, tweets are free and relatively easy to retrieve with some technical understanding of how to work with Twitter's Application Programming Interface (API). The text data they contain allow for the extraction of psychological information generated naturally, without interference from a researcher. Researchers often achieve this extraction using the LIWC program, but they may also use other computational text-analytic techniques (e.g., word frequency, inverse frequency). This relatively straight-forward method for capturing emotion in tweets is powerful, especially when coupled with the fact that tweets are archival by design. That is, each tweet has a timestamp that allows a researcher to compute pre-event trends of emotion expression that can serve as a baseline against which to compare post-event trends. Moreover, because tweets are rather easy to collect, it is also much easier to include tweets from a control group, a necessary but rare feature of collective-trauma studies.

Despite the virtues of Twitter data, however, there are several methodological challenges researchers should keep in mind. In any study of the impact of a collective trauma, identifying users in a target location is key, and there is marked variation in how researchers have sought to achieve this goal. Unsurprisingly, the majority of studies of collective traumas that use Twitter data rely on analyses of geo-coded tweets, that is, tweets tied to a specific geographic location via geocoordinates (Lin, Margolin, & Wen, 2017). This method offers total certainty about where a tweet was generated and, in using this approach, a researcher can capture tweets from residents and visitors in a target location affected by a collective trauma. However, as some Twitter researchers have noted (see, e.g., Jones, Wojcik, Sweeting, & Silver, 2016), very few users opt in to having their tweets geo-coded. In fact, estimates of the percentage of Twitter users who opt into having their tweet location made public range between ~6% and ~8% (Jones et al., 2016; Lin, Margolin, & Wen, 2017), although the true number is only known by Twitter.

Other Twitter studies rely on hashtags and keywords related to a collective trauma to source Twitter data (De Choudhury, Monroy-Hernandez, and Mark, 2014). This approach is reasonable, given that when mass violence occurs, communities tend to respond on social media with hashtags like #BostonStrong or #PrayForParis. Although this may result in downloading tweets relevant to a specific event, however, there is no guarantee that these tweets will represent users in the affected location, especially when the event under study garners national (or international) media attention. Importantly, using this strategy to source Twitter data a) only allows a researcher to access tweets generated seven days before they initiate their search because of restrictions set by Twitter in the availability of data from their API, and b) limits the search to a small, but undetermined subset of all available Twitter data (https://developer.twitter.com/en/docs/ tweets/search/overview/standard).

Because of these limitations, some researchers have used other methods to approximate a user's location. For example, some researchers take advantage of the fact that downloaded Twitter data come with a variable called "location," a field that users can optionally complete as part of their Twitter profile (Doré, Ort, Braverman, and Ochsner, 2015). This field can contain useful information like "Purdy, MO" or unusable information like "New Berlin, Luna." The advantage here is that users who provide a real location likely live there. The downside to this method is that it can be labor intensive for researchers to sift through each user's location to identify individuals to retain in the study. If automated means of identifying locations of interest are employed, misspelled locations may be excluded unless a human identifies them or the automated system is built to account for variations in spelling. It is also unclear whether users differ systematically in their propensity to provide this information.

Rather than relying solely on geo-coded tweets, keywords, or the location field to source Twitter data for analysis, researchers can identify users in a community of interest in two other ways. The first involves scraping tweets generated in a geographic region by interfacing with the Twitter API and downloading tweets using a bounded geographic box, defined by a set of geocoordinates that demarcate the box geospatially. This produces a dataset of tweets generated in a geographic area and provides information about the users to whom those tweets belong. Researchers can then construct a computational focus group (Lin, Keegan, Margolin, & Lazer, 2014), comprised of users with public accounts in the dataset, and then scrape the most recent 3,200 tweets, the maximum number of tweets accessible for free through the Twitter API, from each user's Twitter timeline. Although this method provides researchers with a timeline of data for each user before and after the event in question, the data are still derived from the very small subset of users who opt in to having their tweets geocoded. Moreover, this method has the potential to capture tweets from residents *and* visitors of the geographic location under study, in which case community-based claims of an event's impact may be overstated.

The second alternative method is community specific. To target likely residents in a community impacted by mass violence, researchers can first identify Twitter accounts of local organizations or agencies in that community and then download the list of users who follow those accounts (cf. Jones et al., 2016). The rationale is that users who follow the Twitter account of the Riverside, CA City Hall are likely to be residents and it is unlikely that residents of New York City would have any reason or interest in doing so.

After compiling lists of users who follow local accounts, researchers can download the most recent 3,200 tweets from each user who possesses a public account (9 – 13% of users have private accounts; see Jones et al., 2016). An additional advantage of this method is that it lends itself to the inclusion of data from a closely matched control community, an important feature of any observational study attempting to make a strong inferential claim about the impact of a mass violence event. This method is not without error. It is possible that accounts belonging to users outside of the community may be included when scraping the list of users who follow local accounts. For example, residents of an area who have moved away may still follow local accounts of their former community.

Big data in a trauma context

Keeping the pros and cons of each approach in mind, we present findings from several studies that have leveraged the methodological flexibility of working with Twitter data to understand how communities are impacted by a variety of collective traumas. As an early example, De Choudhury and colleagues (2014) examined Twitter data, coupled with county-level homicide data, to study the impact of the Mexican drug war across several Mexican communities facing ongoing, protracted exposure to violence. These researchers used a keyword search to harvest approximately 3.1 million tweets over the course of 2 years. They found that negative affect expressed in tweets declined over time, despite rising homicide rates in the communities under study. These authors suggested that this trend is perhaps indicative of a desensitization effect in which exposure to increasing violence numbs psychological responses to it. Importantly, this study demonstrated the value of triangulating on a community response by coupling Twitter data with related objective data (i.e., local homicide rates) to make a strong circumstantial case about the importance of the Twitter results obtained in the context under study.

Researchers have also used Twitter data to explore temporal and geospatial psychological responses to school shootings. For example, Doré, Ort, Braverman, and Ochsner (2015) explored how emotion expression changed after the Sandy Hook Elementary School shooting as a function of time and distance from this tragedy. The authors used Twitter's Streaming API, a real-time data collection method using keywords, to source their data and relied on the location field to approximate a user's location. They used the data to test two opposing, theoretically-based hypotheses about the pattern of sadness and anxiety expected. Specifically, they tested whether sadness and anxiety would decay at the same rate over time and distance from the school shooting (nonspecific distance hypothesis) or whether sadness and anxiety would change differentially across time and distance, based on appraisals of the event (construal-level hypothesis). Analyses revealed that expressions of sadness were most prominent around the Newtown, CT area immediately after the shooting and that emotion expression shifted toward anxiety with time and geographic distance from the school. Consistent with the construal-level hypothesis, authors found that causal thinking mediated these relationships.

Twitter data have also proven useful for understanding responses to international terrorist attacks. Lin, Margolin, and Wen (2017) examined geo-located Twitter data from 6,514 users identified as part of their computational focus group in Paris during the attacks in November 2015. They found marked increases in sadness, anxiety, and anger among individuals in Paris, and found that the negative affect lingered for days following the attacks there. Comparable patterns of negative emotion expression appeared among a similarly-obtained sample of tweets in the aftermath of the Brussels terrorist attack 4 months later, although the psychological response was not as strong.

The utility of Twitter data has also been demonstrated when examining community responses to college campus shootings. For example, Jones and colleagues (2016) sourced Twitter data by identifying Twitter accounts local to communities that experienced incidents of college campus violence (e.g., Isla Vista killings) and downloading tweets from users who followed those accounts. They also sourced data from control communities to allow for a comparison of patterns between the target and control communities. Across three instances of violence on college campuses, they found that event-related negative emotion expression on Twitter increased sharply following each event and was sustained for several days before returning to baseline.

Extracting meaning from Twitter content

Using LIWC dictionaries to count words in tweets to capture psychological constructs is a top-down approach that imposes an analytic frame onto the data. This is the most commonly used method in Twitter studies and it is suitable for understanding emotion expressed in Twitter data over time. However, Twitter data are also rich with meaning that can be extracted to understand context-specific aspects of an event that may be important. From a natural language processing (NLP) perspective, the exercise of extracting meaning from a text corpus (i.e., a collection of text documents) is an automated, data-driven (or a bottom-up) approach. Several computational techniques have been developed for analyzing text data, and they vary considerably in technical sophistication depending on the problem at hand. However, simple NLP techniques used by computational linguists may be sufficient for traumatic stress researchers to quickly get a sense of tweet content over time.

One such technique is called an *n*-gram analysis (Silge & Robinson, 2017) in which a corpus is tokenized into words that comprise all documents in the text. The words are automatically counted to produce each word’s frequency and commonly used words (e.g., the, and) are stripped out. In a unigram analysis, the researcher can request a list of the most common words in the corpus. Unigrams can be presented in a variety of ways, including via “word clouds," which have gained popularity recently because of the ease with which they can be created and the rough overview of content they provide. Nonetheless, the flexibility of the n-gram approach is that the "n" of n-gram refers to the number of co-occurring words a researcher chooses to identify in the corpus. If the interest is in identifying common word-pairs, this would be a bigram analysis. Thus, for example, if a document in a corpus contains the phrase "I went to the store," all possible adjacent word-pairings are extracted (e.g., I went, went to, to the, the store). These pairings are created and counted for all documents in the corpus to provide a frequency of the most commonly used pairs. This technique has not yet been used in Twitter studies of collective traumas previously but may be particularly useful for the trauma field because a) it is a relatively easily-to-employ automated process, and b) it may reveal important themes generated on Twitter after a collective trauma that may be relevant for understanding context-specific aspects of an event that can help explain psychological responses.

The current study

This study employed the method described by Jones and colleagues (2016) to locate Twitter users in a community (which had not yet been tested in the context of a mass violence event), along with users in a matched control community. Using tweets from these samples, we sought to model negative emotion expression over time and explored the utility of using word clouds and n-grams to examine the community impact of the December 2, 2015, mass violence in San Bernardino, California. On this date, 14 people were killed and 22 were wounded during a training event and holiday party for the San Bernardino Public Health Department. The perpetrators, who were husband and wife, fled the scene and were later killed in a shootout with police. At the time, this incident was considered the deadliest mass shooting in the U.S. since 2012.

Research questions

RQ1: What is the time course of negative emotion expression in San Bernardino after a mass violence event compared to a matched control community?

RQ2: What is the informational efficacy of a word cloud and n-gram analysis of tweet content generated in the immediate aftermath of the attack?

Method

Participants and procedure

Comparison community selection. Using procedures outlined in Wicke and Silver (2009), a list of potential comparison communities that matched San Bernardino on demographic and economic metrics available from the U.S census was created. Although several candidate communities were identified, Stockton, California, was selected primarily because its ethnic distribution was similar to San Bernardino more so than the other candidate communities.

Twitter data collection and cleaning procedure. Procedures outlined by Jones et al. (2016) were employed to source Twitter data generated by likely residents from each community. Twitter accounts owned by local government and commercial organizations (e.g., city hall, local radio stations) and likely followed by residents of San Bernardino and Stockton, respectively, were identified. This list included 43 local accounts in San Bernardino and 18 accounts in Stockton. Next, the twitteR package (Gentry, 2015) for R (R Core Team, 2017) was used to interface with Twitter's API and scrape the most recent 4,000 followers (the average number of followers across San Bernardino accounts) of each San Bernardino account and 10,000 followers of each Stockton account. Given that we identified far fewer Stockton community accounts, we opted to download many more Stockton followers in an attempt to source a similar number of users across communities. After filtering out "verified" follower accounts (via a variable provided by Twitter) that likely belonged to businesses, celebrities, and other public figures – rather than local community members – a list of 47,962 user accounts in San Bernardino and 22,409 in Stockton was retained. Despite our efforts to equalize the user samples across communities, the number of Stockton users was nearly half that of San Bernardino. Nonetheless, we identified tens of thousands of users in each community.

This list was then read by an R script that interfaced with the Twitter API and downloaded the most recent 1,000 tweets generated by each user. Selecting the number of tweets to download from each user depends on how far back in time one wishes (or needs) to go to source tweets – the more tweets requested at this step will ensure that the data, in aggregate, will have been generated around the event in question. Moreover, the number of tweets one requests also depends on when data collection begins relative to the event. We began data collection roughly a year and a half after the attack and felt that moderate usage (two tweets per day for ~547 days, or 1,000 tweets), coupled with the large sample of users was sufficient for the purposes of this study. We also considered computing limitations; requesting the maximum number of tweets for over 47,000 users may have increased users included in all our analyses, but this would have been computationally prohibitive for a basic desktop computer.

In all, 1.16 million tweets representing 26,030 users who tweeted in a 12-week window around the attack (cf. Jones et al., 2016) were downloaded (see Table 1 for details by community). Data collection procedures for this study were reviewed by the Institutional Review Board of the University of California, Irvine and assessed not to constitute human subjects research.

Table 1. Number of users and tweets

	12-week window		2-week window	
Community	**# tweets (%)**	**# users (%)**	**# tweets (%)**	**# users (%)**
San Bernardino	784,479 (67.5)	17,980 (69.0)	150,481 (68.9)	12,611 (69.2)
Stockton	376,903 (32.5)	8,050 (31.0)	68,052 (31.1)	5,618 (30.8)
Total	1,161,382	26,030	218,533	18,229

Time. A primary goal was to explore the time course of negative emotion expressed in the days and weeks before and after the event (RQ1). To do so, the timestamp associated with each tweet was used to group tweets by hour. It is important to note that when downloading tweets from the Twitter API, the timestamp for each tweet is expressed in Universal Time, Coordinated (UTC) and should be converted to the local time of the communities under study for ease of analysis. After being converted to Pacific Standard Time (PST), tweets were automatically coded for their temporal distance from the event (i.e., December 2, 2015, around 11 am) by subtracting the date and time the tweet was generated from the event date and time. For example, if a tweet was generated on November 29, 2015 at 11:00 am, it was coded as -96 (4 days or 96 hours before the event). Likewise, a tweet generated on December 3 at 7:00 pm was coded as 20 (20 hours after the event). Converting timestamps (typically downloaded as a string variable) to a date and time format easily readable by R and other packages is important so that calculations can be conducted on these dates.

Measures

Negative Emotion. A custom R script was used to compare the words in each tweet to a list of negative emotion words (e.g., sad, hate, tragedy) available in the LIWC program. Each tweet was then coded dichotomously such that if it contained at least one negative emotion word it was assigned a 1 (all others were coded 0). This allowed for a calculation of a proportion of tweets containing negative emotion in a given time frame (e.g., day, hour).

Analytic strategy

Visualizations and statistical analyses. Data were manipulated using tidytext (Silge & Robinson, 2016) and descriptive visualizations were created in R using ggplot2 (Whickam, 2009). We also employed the generalized additive model (GAM) smoothing function when creating plots with ggplot2 to depict a non-linear line-of-best-fit across time. To estimate linear trajectories over time and evaluate the significance of changes in negative emotion expression before and after the attack, a piecewise regression with a discontinuity analysis was conducted in Stata 14.2 (College Station, TX), using procedures outlined by others (Jones et al., 2016; Mitchell, 2012). This approach is well-suited for statistically evaluating non-linear changes in time-series data for which points of change are expected. First, a "knot" is placed at the moment of the event to signify a break in the time-series data. Doing so allows for an estimation of trends before and after the knot. A strength of this approach is that it provides a clear way to describe non-linear trends in time-series data for which standard mathematical functions do not readily apply. Additionally, in this application of a piecewise regression, the use of a knot placed at the time of the event allows for a regression discontinuity analysis (Thistlethwaite & Campbell, 1960) in which a predicted value of negative emotion is calculated at the knot, as if the event had not occurred. This value can then be statistically compared to a value representative of negative emotion that occurred at that knot-point.

Many users tweeted multiple times each day throughout the 12-week time frame. If a user was more prolific than others and expressed negative emotion across many of his or her tweets, that user could have undue influence on the results due to the correlated nature of the tweets. These dependencies violate the assumption of independence of residuals for an ordinary least squares regression analysis. To compensate, tweets were clustered within users in all regression analyses (Primo, Jacobsmeier, & Milyo, 2007; UCLA Statistical Consulting Group, 2015). In all contrasts, Stockton served as the reference group.

Tweet content analysis. To understand the content of tweets generated immediately after the attack, two text mining procedures in R were employed using tidytext (see Silge & Robinson, 2018, for a description of these techniques and the code to execute them). First, tweets were tokenized and common English-language words (e.g., the, and) were stripped out. Counts of remaining words generated the day of the attack were produced. A word cloud comprised of the most commonly used words after the attack was then generated using the wordcloud (Fellows, 2013) package in R. A second word cloud – with positive and negative words differentiated – was produced using the Bing sentiment lexicon function in tidytext, which labels words based on their valence.

Next, a bigram analysis was conducted in each of 12 hours after the attack to further extract tweet content as it unfolded in the immediate aftermath of the attack. Bigrams in each hour were assigned an importance weight using the term frequency-inverse document frequency (tf-idf) function in tidytext, a common technique that weights less commonly used words more heavily than commonly used words.

Results

Negative emotion trends

Raw proportions of negative emotion expression calculated each hour across the 12-week window are presented in Figure 1. There was marked variation in negative emotion expression in San Bernardino before the attack due to an unrelated event, the anniversary of the shooting of a young girl in the community. Also, relative to Stockton, this figure depicts a marked increase in negative emotion in San Bernardino on the day of the mass violence event.

Figure 1. Hourly negative emotion expression among likely residents of San Bernardino and Stockton 6 weeks before and after the San Bernardino mass violence event

(n_{tweets} = 1.16 million; n_{users} = 26,030)

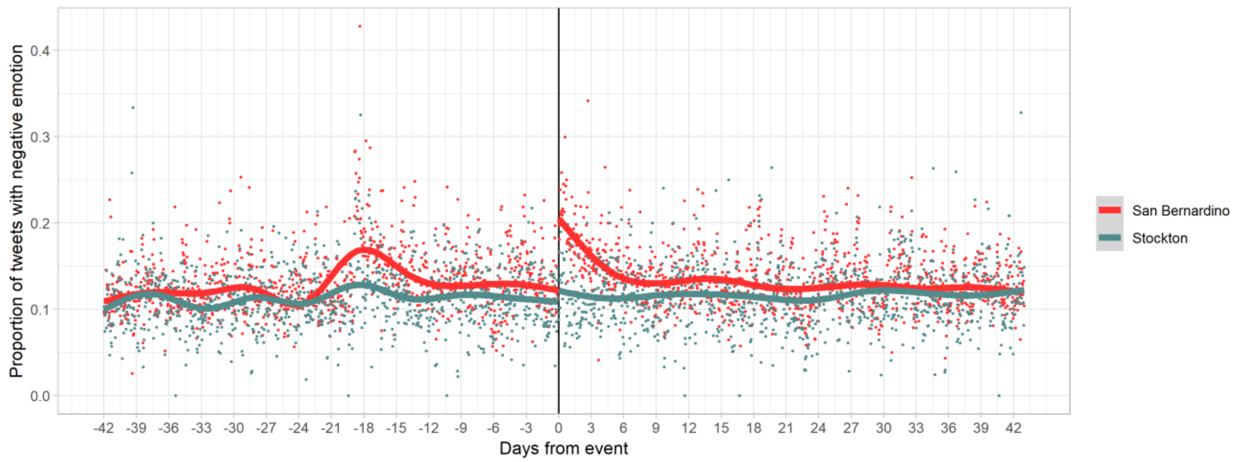

Given the spike in negative emotion roughly 18 days before the attack, constructing a piecewise regression to evaluate the significance of emotion trends was problematic because the spike was not anticipated and would require an additional knot. However, the negative emotion trend exhibited by tweets in San Bernardino after the unexpected spike stabilized by the week before the attack, and negative emotion expression during this week appeared similar to the baseline negative emotion expressed before the anniversary of the young girl's death. Thus, we restricted the piecewise regression to the week before and after the attack (14-day window), aggregating tweets by day, to accurately model the trends in negative emotion expression in this time frame (see Figure 2).

Figure 2. Daily negative emotion expression among likely residents of San Bernardino and Stockton 7 days before and after the mass violence event

(n_{tweets} = 218,533; n_{users} = 18,229)

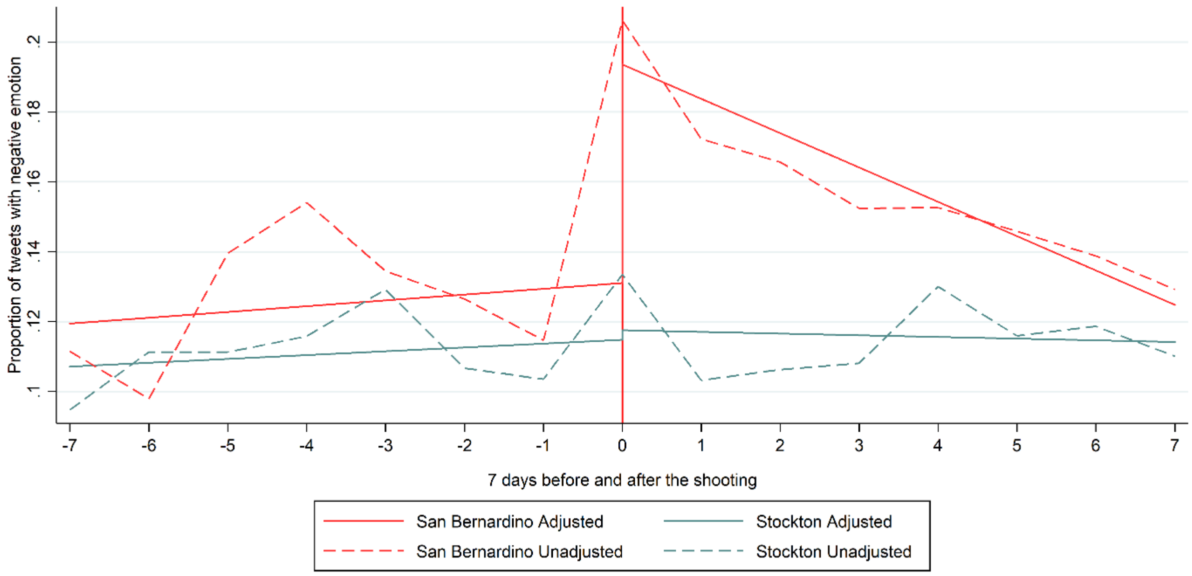

Before the attack, no difference in negative emotion trajectories was observed between San Bernardino and Stockton (standardized *b* = -.001, *SE* = .001, *t* = -.35, *p* = .72). In other words, both communities expressed negative emotion at the same rate before the mass violence event. Relative to baseline, negative emotion increased 6.2% (standardized *b* = .18, *SE* = .01, *t* = 12.33, *p* < .001) in San Bernardino on the day of the attack. No such increase was observed in Stockton (.02%; standardized *b* = .01, *SE* = .01, *t* = 0.46, *p* = .64). In San Bernardino, negative emotion decreased 1.0% each day (standardized *b* = -.03, *SE* = .002, *t* = -13.96, *p* < .001) until returning to baseline roughly six days after the attack. In contrast, the post event trend of negative emotion expression in Stockton remained stable after the attack (standardized *b* = -.001, *SE* = .002, *t* = -0.54, *p* = .59).

Tweet content surrounding the attack

Results from the word clouds from tweets collected by San Bernardino users are presented in Figure 3. The general word cloud reveals that "shooting" was the most commonly used word on the day of the attack. Ancillary words related to the attack also appear, although less prominently (e.g., mass, victims, police, safe). When differentiating the words by negative and positive valence, other dimensions of this event became clearer: this was a terror event, a massacre, and a bomb was involved. There was also much discussion about a suspect and the word "dead" was very prominent, likely reflective of the shootout in which both suspects were killed. The word clouds appeared to converge on descriptive aspects of the attack.

Figure 3. General and sentiment word clouds from tweets generated the day of the San Bernardino mass violence event

(n_{tweets} = 38,708; n_{users} = 6,788)

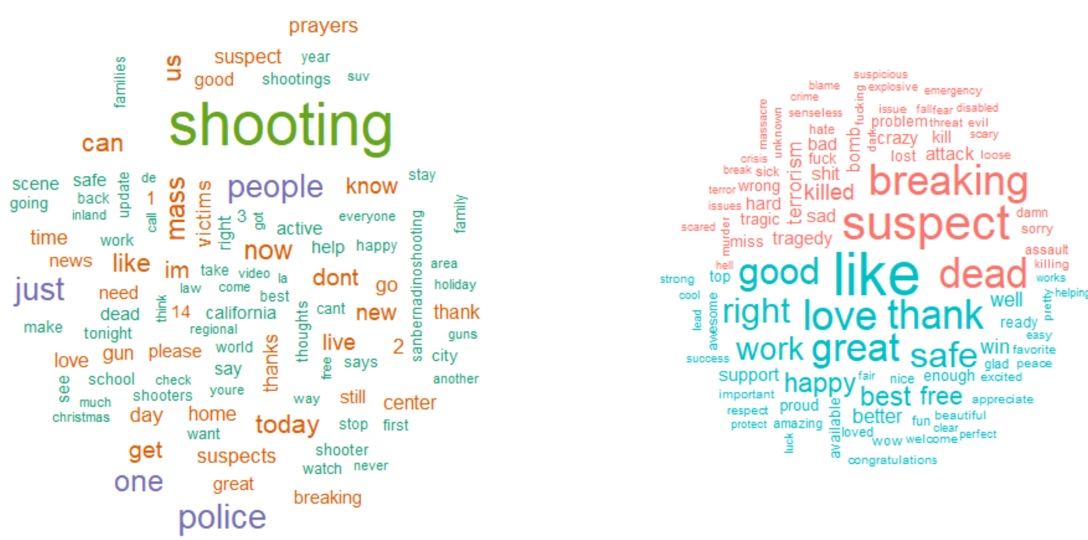

General word cloud **Sentiment word cloud**

Results from the bigram analysis of San Bernardino tweets are presented in Figure 4.

Important word pairs identified in this analysis provide clearer detail about the timing of event-related content tweeted in the 12 hours following the attack. Less than an hour after the attack, users began tweeting about a "shooter situation." As time unfolded, users tweeted about a "terrorist incident" (hour 3), a male and female perpetrator (hour 6) and eventually used each shooter's name (hour 7). In the 11th hour, a flurry of tweets was generated about the event, coinciding with a press conference held by the chief of police.

Figure 4. Top five bigrams in each of the 12 hours following the San Bernardino mass violence event

Discussion

Our aim was to use a robust method for sourcing Twitter data in the context of a mass violence event using a method previously validated as efficacious in the context of school shootings on university campuses. Importantly, this method circumvented many of the drawbacks of other approaches (e.g., geotagged tweets and keyword searches), and the ease with which the data could be sourced allowed for the inclusion of a control community against which to compare trends in San Bernardino. Overall, the results of the linguistic word count analyses revealed that following a mass violence event, negative emotion increased sharply and remained elevated for several days after the event; no such pattern was observed in Stockton, the matched control community. This pattern is consistent with past studies that have used Twitter data to explicate the impact of mass violence on communities (Doré, Ort, Braverman, and Ochsner, 2015; Jones et al., 2016; Lin, Margolin, & Wen, 2017).

We also used freely available standard text mining tools to explore their efficacy in revealing thematic content of tweets not captured by coding for emotion using a top-down approach (i.e., LIWC dictionaries). The results of the word clouds and bigram analyses demonstrate that bigrams provide more contextual detail than word clouds, although word clouds do provide a basic idea of how Twitter data can be mined to identify themes surrounding a collective trauma. Interestingly, the bigram analysis revealed a pattern of event-related information flow generated on Twitter in which information moved from ambiguous (e.g., shooter reported, male female) to concrete (e.g., suspects dead, suspect names) as time unfolded. This trend is consistent with news reporting about the incident over time and likely reflects exposure to the news and the updates released to the public during and after the ordeal.

Limitations

Past work shows that Twitter users are not representative of the general population, as they are typically younger (18–29) and from urban communities (Pew Research Center, 2018). Although this does limit the extent to which the findings may generalize to a full community under study, the demographics of San Bernardino (e.g., lower median age relative to the nation) suggests that many residents are likely to be users. Concerns about representativeness will always be woven into studies of individuals using technologies that are not yet widely adopted in the general population.

User location. Despite its utility, this method for sourcing data from likely residents of San Bernardino and Stockton is not without limitations. The first is that there is no guarantee that all users identified currently lived in these target communities. Moreover, lists of users who follow large commercial or governmental accounts contain more opportunities for error than those of hyper-local community accounts (e.g., local high school). For example, it could be the case that users that make it into a follower list belong to businesses or other governmental agencies or representatives, rather than individual residents. Potential avenues for mitigating both issues would involve filtering out users who indicated living somewhere other than the target community (via their profile location field) and filtering out users with an excessive number of tweets, relative to others in the sample, because organizations tend to tweet more than individual users. It is important that the benefit of such filtering procedures be weighed against the risk of excessive data loss when employing them, thereby prohibiting a researcher's ability to analyze the data in a meaningful way.

Limited content. As basic tools for extracting thematic content from a corpus, techniques such as word clouds and n-grams adequately revealed some aspects of the San Bernardino attack as they unfolded. However, it did not reveal any information about the neighborhood lockdown during the shootout between the perpetrators and police, a particularly tense time for some residents, coupled with live local news coverage via a helicopter over the area. Moreover, the bigrams did not identify any additional emotional content to shed more light on the community's response in the immediate aftermath. The utility of this method could be limited to assisting researchers in generating an event-related word list with which to tag event-related tweets for further qualitative examination. Future work should address this possibility to bring a mixed-methods frame to the analysis of large-scale Twitter data.

Future directions

The strategy for coding emotion in tweets, as we have outlined it, is just one way to model negative emotion in a community after a mass violence event. As others have noted (Jones et al., 2016), when dealing with social media data, how the data are coded depends on the research question. Twitter data coded for emotion are flexible in that they can be manipulated and modeled several ways. For example, rather than calculating the proportion of tweets with negative emotion over time, as we have done, it is possible instead to analyze the proportion of users who expressed negative emotion across time. It is also possible to examine only the raw counts of negative emotion words used each hour or day across the study period, although patterns of raw counts may be difficult to interpret.

An example of an analysis afforded by this coding flexibility is to examine differential emotion expression in a community based on pre-event tweet patterns of the users in a sample. For example, perhaps individuals who tweet less negative emotion than others do before a collective trauma exhibit a stronger reaction to the event both immediately and over time. Such an analysis would go a long way in clarifying theoretical assumptions of reactions to collective trauma using large-scale, ecologically valid data from many thousands of users. Social media user behavior might also be important to consider. For example, users who share more news links in the aftermath of an event may display differential negative emotion relative to those who share fewer news stories. These types of analyses are possible and worthy of exploration to fully capture psychological and behavioral underpinnings of the overall trends observed in longitudinal Twitter studies.

Extracting content from tweets, as we have outlined it in this paper, is relatively simple to achieve. There are, however, more statistically sophisticated methods for analyzing corpora. For example, Latent Dirichlet allocation (LDA) is a topic modeling technique that algorithmically accounts for the fact that a corpus of text is made up of many documents, each containing many topics, and that each topic is made up of many words. Importantly, words can overlap in topics and topics can overlap in documents. This technique has proven useful for delineating topics in book chapters, but it is not clear whether it can surmount the "noise" in Twitter data. To mitigate this noise, it may be prudent to use the techniques discussed earlier to identify and filter down to event-related tweets and then use LDA to extract their thematic content. Insofar as topic modeling, in all its forms, can be used to extract thematic content from Twitter data, it would also be useful to model this content over time. Indeed, there is an R package that does this (i.e., "tm" package). Future work might explore how Twitter content generated after a mass violence event can reveal information about how communities mobilize material and social resources when a trauma occurs and elucidate long-term impact and recovery efforts.

As demonstrated by De Choudhury et al. (2015), there is immense value in pairing Twitter data with other data sources to triangulate on explanatory factors that may elucidate longitudinal trends in Twitter studies. These authors found, for example, that over the course of two years, negative affect in tweets generated by users in communities impacted by the Mexican drug war decreased over time. To understand this trend, they folded in community-level rates of death due to violence to provide evidence of affective desensitization. Although linking Twitter data with other relevant data sources may not always be possible, it is important that researchers interested in using Twitter data attempt to do so to strengthen their explanations of patterns they find.

Conclusion

In this paper we described the challenges and benefits of working with Twitter data to highlight the important contribution data like these can make to the traumatic stress field. The considerations and analytic techniques we have highlighted can also be applied to big social media data sourced from other platforms with text-based data (e.g., Facebook, Reddit). Despite Facebook's recent policy change limiting access to its users' data through the platform's API, it has created opportunities for researchers to apply for access to its data to conduct social science research. Until other data sources become available, Twitter is the only widely-used social media platform offering researchers free access to public data generated by its users. Although Twitter studies have been ongoing since the platform was launched over a decade ago, the potential these data have for elucidating psychological processes on a large-scale is far from exhausted. We hope that by introducing a selection of studies that highlight the different approaches with which researchers have sourced local Twitter data and demonstrating how these data can be analyzed, trauma researchers will be energized to answer their own research questions by incorporating a Big Data approach.

References

De Choudhury, M., Monroy-Hernandez, A., & Mark, G. (2014). Narco emotions: Affect and desensitization in social media during the Mexican drug war. *Proceedings of the SIGCHI conference on Human Factors in Computing Systems* (pp. 3563–3572). New York, NY: ACM. https://dx.doi.org/10.1145/2556288.2557197

Doré, B., Ort, L., Braverman, O., & Ochsner, K. N. (2015). Sadness shifts to anxiety over time and distance from the national tragedy in Newtown, Connecticut. *Psychological Science*, *26*, 363–373. https://dx.doi.org/10.1177/0956797614562218

Fellows, I. (2013). *wordcloud: Word clouds (Version 2.5) [Computer software]*. Retrieved from https://cran.r-project.org/web/packages/wordcloud/wordcloud.pdf

Gentry, J. (2016). *twitteR: R based Twitter client (Version 1.1.9) [Computer software]*. Retrieved from https://cran.r-project.org/web/packages/twitteR/twitteR.pdf

Jones, N. M., Wojcik, S. P., Sweeting, J., & Silver, R. C. (2016). Tweeting negative emotion: An investigation of Twitter data in the aftermath of violence on college campuses. *Psychological Methods*, *21*, 526–541. https://dx.doi.org/10.1037/met0000099

Lin, Y., Keegan, B., Margolin, D., & Lazer, D. (2014). Rising tides or rising stars? Dynamics of shared attention on Twitter during media events. *PloS One*, *9*, e94093. https://dx.doi.org/10.1371/journal.pone.0094093

Lin, Y., Margolin, D., & Wen, X. (2017). Tracking and analyzing individual distress following terrorist attacks using social media streams. *Risk Analysis*, *37*, 1580–1605. https://dx.doi.org/10.1111/risa.12829

Mitchell, M. N. (2012). *Interpreting and visualizing regression models using Stata*. College Station, TX: Stata Press.

National Research Council. (2013). *Frontiers in massive data analysis*. Washington, DC: The National Academies Press. https://doi.org/10.17226/18374.

Pennebaker, J. W., Booth, R. J., Boyd, R. L., & Frances, M. E. (2015). *Linguistic Inquiry and Word Count: LIWC2015*. Austin, TX: Pennebaker Conglomerates (www.LIWC.net).

Pew Research Center (2018). *Social media use in 2018*. Retrieved from http://www.pewinternet.org/2018/03/01/social-media-use-in-2018/

Primo, D. M., Jacobsmeier, M. L., & Milyo, J. (2007). Estimating the impact of state policies and institutions with mixed-level data. *State Politics & Policy Quarterly*, *7*, 446–459. https://dx.doi.org/10.1177/153244000700700405

R Core Team. (2018). *R: A Language and Environment for Statistical Computing*. Vienna, Austria: R Foundation for Statistical Computing. Retrieved from http://www.R-project.org/

Shadish, W. R., Cook, T. D., & Campbell, D. T. (2002). *Experimental and quasi-experimental designs for generalized causal inference*. Boston, MA: Houghton, Mifflin and Company.

Silge, J., & Robinson, D. (2016). tidytext: Text mining and analysis using tidy data principles in R. *The Journal of Open Source Software*, *1*. https://dx.doi.org/10.21105/joss.00037

Silge, J., & Robinson, D. (2018). *Text Mining with R: A tidy approach*. Retrieved from https://www.tidytextmining.com/

Silver, R. C. (2004). Conducting research after the 9/11 terrorist attacks: Challenges and results. *Families, Systems, & Health*, *22*, 47–51. https://dx.doi.org/10.1037/1091-7527.22.1.47

Steinberg, A. M., Brymer, M. J., Steinberg, J. R., & Pfefferbaum, B. (2006). Conducting research on children and adolescents after mass trauma. In F. H. Norris, S. Galea, M. J. Friedman, & P. J. Watson (Eds.), *Methods for disaster mental health research* (pp. 243-253). New York, NY: Guilford Press.

Thistlethwaite, D. L., & Campbell, D. T. (1960). Regression-discontinuity analysis: An alternative to the ex post facto experiment. *Journal of Educational Psychology*, *51*, 309–317. https://dx.doi.org/10.1037/h0044319

UCLA Statistical Consulting Group. (2015). *Regression with Stata: Beyond OLS*. Retrieved from http://www.ats.ucla.edu/stat/stata/webbooks/reg/chapter4/statareg4.htm

Wicke, T., & Silver, R. C. (2009). A community responds to collective trauma: An ecological analysis of the James Byrd murder in Jasper, Texas. *American Journal of Community Psychology*, *44*, 233–248. https://dx.doi.org/10.1007/s10464-009-9262-8

Wickham, H. (2009). *ggplot2: Elegant Graphics for Data Analysis*. Springer-Verlag New York. Retrieved from http://ggplot2.org

Acknowledgements

This Chapter has been submitted for publication in a peer-reviewed journal. I thank my coauthors for their invaluable contributions to this manuscript. I also thank Sarah Kastor for assisting me with identifying the control community in this study.

Jones, N. M., Brymer, M., & Silver, R. C. (submitted for publication). Using big data to study the impact of mass violence: Opportunities for the traumatic stress field.

Chapter 3: Distress and rumor exposure on social media during a campus lockdown

Abstract

During crisis events, people often seek out event-related information to stay informed of what is happening. However, when information from official channels is lacking or disseminated irregularly, people may be at risk of exposure to rumors that fill the information void. We studied information-seeking during a university lockdown following an active-shooter event. In Study 1a, students in the lockdown (N = 3,890) completed anonymous surveys one week later. Those who indicated receiving conflicting information about the lockdown reported greater acute stress [standardized regression coefficient (*b*) = .07; *SE* = .01; 95% confidence interval (CI), .04, .10]. Additionally, those who reported direct contact with close others via text message (or phone) and used Twitter for critical updates during the lockdown were exposed to more conflicting information. Higher acute stress was reported by heavy social media users who trusted social media for critical updates (*b* = .06, *SE* = .01; 95% CI, .03, .10). In Study 1b, we employed a big data approach to explore the time course of rumor transmission across 5 hours surrounding the lockdown within a subset of the university's Twitter followers. We also examined the patterning of distress in the hours during the lockdown as rumors about what was happening (e.g., presence of multiple shooters) spread among Twitter users. During periods without updates from official channels, rumors and distress increased. Results highlight the importance of releasing substantive updates at regular intervals during a crisis event and monitoring social media for rumors to mitigate rumor exposure and distress.

Study 1a: Introduction

Since 2000, active shooter events in the United States have been on the rise and often result in higher casualty counts when they occur in educational institutions compared to other settings (Blair & Schweit, 2014). Crisis events like these are often accompanied by deficits of credible information as authorities attempt to piece together the facts of the unfolding events. As a consequence, individuals caught in the wake of the crisis are left vulnerable to rumors and conflicting information from unofficial communication channels they trust. Despite a robust scientific literature on the circumstances in which rumors are generated (Allport & Postman, 1946; Festinger et al., 1948; Rosnow, 1991) and transmitted (Fitzhugh, Gibson, Spiro, & Butts, 2016; Starbird, Maddock, Orand, Achterman, & Mason, 2014), little is known about the psychological correlates of exposure to rumors during crisis events. When rumors proliferate during dangerous and uncertain situations, what impact do they have on people who receive and believe them?

Crisis events are often ambiguous in nature and when uncertainty is high, appraisals of threat among individuals caught in their path can be heightened (Taha, Matheson, & Anisman, 2014). This, in turn, may instigate information-seeking behavior as a way of reducing situational uncertainty (Ball-Rokeach & DeFleur, 1976; Seeger, Vennette, Ulmer, & Sellnow, 2002) and consequently, the psychological distress such uncertainty engenders (Chen & Hong, 2010). In the past, this information seeking led people to their radios and televisions to acquire critical updates from official channels. However, when a crisis unfolds, people now increasingly acquire critical updates from social media (e.g., Twitter; Hughes & Palen, 2009; Mazer et al., 2015; Palen & Anderson, 2016; Palen & Liu, 2007; Spiro et al., 2012), along with traditional media channels (Jones, Garfin, Holman, & Silver, 2016a).

Moreover, when information from official channels is irregular or lacks new information, uncertainty and information-seeking behavior are likely sustained. As a result, people may also turn to unofficial channels, such as social media, to mitigate their discomfort. The challenge with social media as a resource for updates, however, is the lack of mechanisms for vetting the accuracy of the information being shared among users. This is particularly important because, as the rumor literature suggests, trust in an information source moderates whether rumors are believed and transmitted (Rosnow, Yost, & Esposito, 1986). In addition, once rumors begin to spread on social media, they are very difficult to undermine with updates or corrections (Starbird et al., 2016).

What people see when they are exposed to media-based coverage of a crisis event has been studied previously in the context of collective traumas, such as terrorist attacks and natural disasters that affect many people and occur without warning (Jones et al., 2016a). Large studies of representative U.S. samples demonstrate that, among other variables, repeated indirect, media-based exposure to collective traumas is associated with event-related distress, even when controlling for pre-trauma media habits and preexisting mental health conditions (Holman, Garfin, & Silver, 2014; Silver et al., 2013). For example, in a study following the Boston Marathon bombings (BMB), researchers found that six or more hours of BMB-related media use was associated with higher acute stress than was direct exposure to the bombings (Holman et al., 2014). This relationship is thought to be at least partly driven by the transmission of graphic, event-related imagery via news coverage (Jones et al., 2016a; Schlenger et al., 2002). However, when images are not relevant or available, other content, like rumors or conflicting information, may also contribute to the distress individuals experience during a crisis.

Current study

In the summer of 2016, a former student of a major university walked into the office of a professor and fired at least three shots: the first shots killed the professor; the other was self-inflicted, killing the assailant. In the immediate aftermath, confusion ensued. Authorities were unclear about the extent of the violence, they assumed the worst (i.e., a mass shooting), and mounted a massive police response. The campus community was alerted of the shooting via email, text, and social media and placed under lockdown for approximately two hours. Those on campus struggled to shelter in place, citing an inability to lock office and classroom doors. Additionally, during this time, communications from campus emergency response officials were irregular. Specifically, after the first lockdown alert, communications were silent for roughly 90 minutes, leaving many students, faculty, and staff in a dangerous and ambiguous crisis situation. This lapse in communication likely provided fertile ground for the generation and transmission of rumors as those affected attempted to grasp what was happening. Study 1a investigated these issues in a survey collected among students caught in the lockdown, which was completed after the event had been resolved and the details of the shooting were released.

Hypotheses and research question

Given that exposure to misinformation likely heightened uncertainty among students trapped in the lockdown, the following hypothesis will be examined:

> **H1**: Exposure to misinformation will be associated with increased distress responses during the lockdown.

Given the evidence linking media use to event-related distress in the aftermath of a collective trauma, it is hypothesized that:

> **H2**: Information channel use (e.g., phone calls and texts from close others, traditional

media, and social media) will be associated with increased distress responses during the lockdown.

Because trust in an information source or channel is theoretically related to whether rumors they generate are believed and transmitted it is hypothesized:

H3: Trust in information channels will moderate the relationship between number of channels used and distress responses during the lockdown.

Finally, in the context of this shooting and lockdown, it is unclear which information channels exposed people to misinformation. Therefore, the following broad research question will be investigated:

RQ1: Which information channels are associated with exposure to misinformation during the lockdown?

Method

Sample and procedures

Beginning seven days after the campus shooting, undergraduate and graduate students were invited to complete an anonymous, internet-based survey via a system-wide email sent on the researchers' behalf by the university administration. The survey was fielded to all 40,339 students listed in the university system. A reminder email was sent out a week later to bolster student participation. Participants who clicked the link to the survey were presented with an initial screen indicating that the purpose of the survey was to study the impact of the campus shooting. Participants were asked to complete the survey without consulting others and instructed to answer items as honestly as possible. Participants were also presented with contact information for the lead researchers and the campus counseling center and encouraged to call or visit the center if they felt a need to speak with someone about their feelings regarding the incident. Data were collected up to 29 days following the shooting, with the majority of responses (92%) collected within 16 days post-event. The participation rate was approximately 18% (*n* = 6,540). Of these, 3,890 (~60%) students reported having been in the lockdown; between 2,696 and 3,393 of these students had complete information on variables across analyses in Study 1. Of the 3,051 students who provided ethnicity data, nearly 40% identified as European American, roughly 30% identified as Asian American, and 14% identified as Latino American; the remainder identified as multi-racial/ethnic (8.6%), African American (2.7%), or other (6%). All procedures for this study were approved by the Institutional Review Board of the University of California, Irvine.

Measures

Dependent variables

Acute stress. Symptoms of acute stress were assessed using the Acute Stress Disorder Scale-5 (Bryant, 2016), which is based on the Diagnostic and Statistical Manual 5th edition (DSM-5; American Psychiatric Association, 2013). Respondents used a 5-point Likert-type scale ranging from 0 (*not at all*) to 4 (*very much*) to describe the extent to which they experienced each of 14 possible reactions "since the shooting and lockdown" (e.g., "Do you have distressing dreams about the lockdown/shooting?"). Responses were summed (range: 0 - 56) to create a continuous score for acute stress symptoms and to capture maximum variability in potential responses (MacCallum, Zhang, Preacher, & Rucker, 2002).

Conflicting information. Respondents were asked to indicate the extent to which they agreed with the statement: "*I received conflicting information from different sources about the details of the shooting*." Responses were measured on a 5-point Likert-type scale ranging from 0 (*not at all*) to 4 (*a great deal*). Because no valid information about the event was available to any students during the lockdown, this variable was used as a proxy for rumor exposure.

To validate this assumption, we examined responses to an open-ended survey item that asked, "*what particular parts of the event were most upsetting to you?*" We used an R script to code each response for whether it mentioned the word "rumor." Those who mentioned rumors reported higher exposure to conflicting information (standardized b =.31, SE =.04, p<.001). To further clarify what participants wrote about rumors, we examined the word pairs (i.e., bigram analysis) that occurred most commonly in the corpus of responses using a text analysis program called Meaning Extraction Helper (Boyd, 2017). The four word-pairs occurring most commonly were: *multiple shooter*, *rumor spread*, *rumor multiple*, and *rumor shooter*. The words "rumors" and "false information" appeared in other responses with less frequency but were also present. Given that many respondents wrote at length, we also conducted a trigram analysis, which analyzes the most common occurrence of 3 words appearing together across responses. This analysis revealed that "rumor [of] multiple shooters" was the most common response.

Independent variables

Traditional media and online news. On a 3-point Likert-type scale (0 = *not at all*, 1 = *some of the time*, 2 = *most or all of the time*), respondents indicated receiving critical updates from radio, television, and online news sites (e.g., CNN, NY Times, TMZ). Responses for each of these channels were dichotomously coded (0 = *not at all*, 1 = at least *some of the time*). They were aggregated to form a count of traditional media/online news sites used (range = 0 to 3).

Direct contact from close others. On a 3-point Likert-type scale (0 = *not at all*, 1 = *some of the time*, 2 = *most or all of the time*), respondents indicated whether they received critical updates from a group text message from a campus student organization, a text message from a friend, a text message from a family member, a phone call from a friend, and a phone call from a family member. Responses for each of these channels were dichotomously coded (0 = *not at all*, 1 = at least *some of the time*). They were aggregated to form a count of direct contact from close others (range = 0 to 5).

Social media. On a 3-point Likert-type scale (0 = *not at all*, 1 = *some of the time*, 2 = *most or all of the time*), respondents indicated whether they received critical updates from Twitter, Facebook, Snapchat, Reddit, and Instagram or some other platform not listed. Responses were dichotomously coded (0 = *not at all*, 1 = at least *some of the time*). They were aggregated to form a count of social media channels used (range = 0 to 5)

Channel trust. For each communication channel students reported using, they were asked to rate how much they trusted it for information about the shooting and lockdown. These ratings were reported on a 5-point Likert-type scale ranging from 1 (*strongly distrust*) to 5 (*strongly trust*). That is, if students reported using Twitter, they were asked to rate how much they trusted it for information and critical updates. Trust ratings were aggregated across each communication channel category to create a composite of trust for traditional media and online news (α = .90), direct contact from close others (α = .89), and social media (α = .90), respectively.

Relevant covariates

Prior school shooting experience. Respondents indicated whether or not they, or someone close to them, ever experienced a school shooting. A total of 16.11% (n = 611) reported previously having such an experience.

Prior trauma exposure. Respondents were asked whether they personally experienced a natural disaster (e.g., tornado, earthquake), community violence (e.g., shooting, civil unrest), combat during war, or any other form of violence prior to the shooting. Responses across these four items were summed and ranged from 0 – 4.

Affiliation with affected department. To capture psychological proximity to the department where the shooting took place, respondents reported their affiliation with the department. In all, 26.2% (n = 830) indicated being affiliated with the department either by being a department major, minor, having taken classes in the department, or some other reason.

Lockdown event exposure. Respondents were asked to indicate whether they experienced each of 9 exposures to the lockdown (e.g., “I was in [building name omitted] when the shooting occurred”). Affirmative responses across these items were summed to create an index of event exposure. Responses ranged from 0 to 9 exposures.

Alone. Respondents indicated whether they were alone or with others during the lockdown. Responses were dummy coded such that if a student indicated being alone, he or she was coded with a 1.

Analytic strategy

Statistical analyses were conducted in Stata 14 (College Station, TX). A series of ordinary least-squares regression analyses were conducted to examine correlates of acute stress and exposure to conflicting information, respectively. Because a number of collective traumas were prominent in the media during data collection, we opted to control for survey completion week to account for the potential influence of these events on participant responses. We also included statistical controls for age and gender. Descriptive statistics for all model variables are reported in Table 1.

Table 1. Descriptive statistics for all variables of interest (n = 3,162)

Variables	***M***	***SD***	**Min**	**Max**
Acute stress	10.57	10.11	0	56
Conflicting information exposure	3.37	.94	0	4
Traditional media count	1.53	.79	0	3
Direct contact from friends/family count	3.22	1.41	0	5
Social media count[a]	1.90	1.25	0	5
Trust in traditional media[b]	4.03	.84	1	5
Trust in direct contact from friends/family[b]	3.73	.85	1	5
Trust in social media[a]	3.29	.91	1	5
Covariates and sample characteristics				
Age	22.43	5.05	18	68
Survey completion week	1.88	.71	1	4
Prior trauma exposure	1.06	.88	0	4
Lockdown event exposure	2.60	1.43	0	9
	%			
Prior school shooting exposure	16			
Alone during lockdown	11			
Department affiliation	26			
Gender				
Female	67			
Male	32			
Other	1			

[a] In Model 2 where trust in social media was interacted with social media use, complete data were available for 2,696 participants. In Model 2, mean social media use was 2.22 (*SD* = 1.07); mean social media trust was 3.29 (*SD* = .91).
[b] Sample size varies due to missing data

Results

Results indicated that exposure to conflicting information was associated with acute stress related to the lockdown, after controlling for several relevant covariates (see Table 2). Traditional media use was not associated with acute stress, but direct contact with close others and social media use were each associated with greater acute stress.

Table 2. Correlates of acute stress among students in the lockdown

	Model 1 (n = 3,162)[a]			Model 2 (n = 2,696)[a]		
Variables	***b*(95% CI)**	**SE$_b$**	***t***	***b*(95% CI)**	**SE$_b$**	***t***
Completion week	-.05(-.08, -.01)*	.01	-3.10	-.08(-.13, -.03)*	.02	-3.23
Gender						
Male, female = 0	-.37(-.44, -.29)**	.03	-10.27	-.36(-.44, -.29)**	.03	-9.30
Other	.48(.16, .80)*	.16	2.98	.44(.08, .80)*	.18	2.44
Age	-.01(-.05, .01)	.01	-.93	-.001(-.01, .01)	.004	-.32
Prior shooting exposure	.17(.08, .26)**	.04	3.85	.21(.12, .31)**	.04	4.48
Prior trauma (violence, war, other)	.13(.10, .17)**	.01	8.07	.15(.10, .19)**	.02	7.09
Dept. affiliation, not affiliated = 0	.10(.03, .18)*	.03	2.86	.10(.02, .18)*	.04	2.53
Lockdown event exposure	.09(.05, .12)**	.01	5.00	.05(.02, .07)**	.01	4.00
Alone, with others = 0	-.11(.-21, -.01)*	.05	-2.25	-.11(-.22, -.01)*	.05	-2.14
Exposure to conflicting information	.07(.04, .10)**	.01	4.28	.08(.04, .12)**	.02	4.14
Count of traditional media use	.01(-.02, .05)	.01	.85	.01(-.03, .05)	.02	.40
Count of contact with friends/family	.13(.10, .17)**	.01	7.21	.09(.06, .12)**	.01	6.70
Count of social media use	.07(.03, .10)**	.01	3.99	-.15(-.28, -.03)*	.06	-2.42
Social media trust	--	--	--	-.06(-.15, .01)	.04	-1.55
Social media use X social media trust	--	--	--	.06(.03, .10)**	.01	3.66
Model Statistics	F(13, 3148) = 43.02; p< .001; R^2= .15			F(15, 2680) = 33.31; p< .001; R^2= .15		

*p < .05; **p < .001; all regression coefficients are standardized

[a] Sample sizes vary across models due to missing data.

We then examined whether trust in these communication channels moderated their respective relations with distress. No moderating effect was found for trust in direct contact with close others. However, greater acute stress was reported by heavy social media users who trusted social media channels for critical updates (see Table 1, Model 2; Figure 1).

Figure 1. Association between count of social media channels used during the lockdown and acute stress, moderated by level of trust in social media

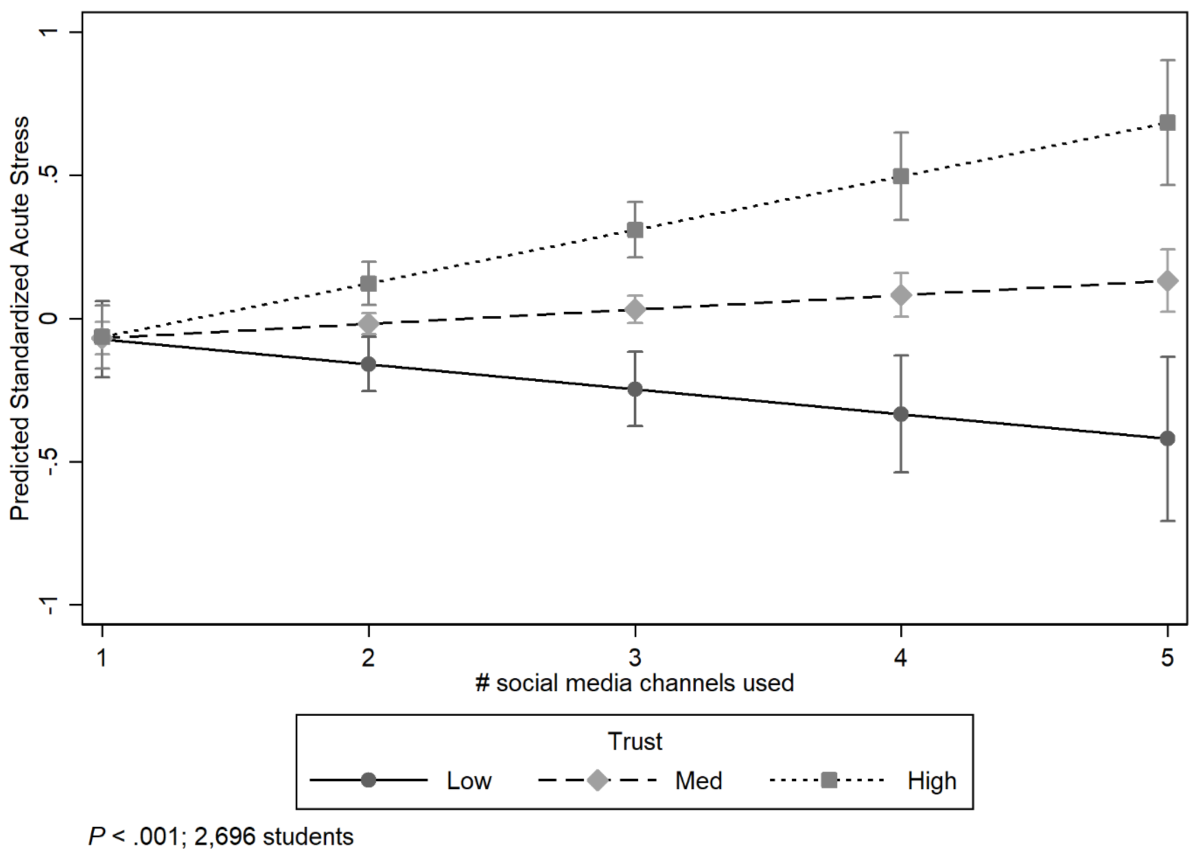

In addition, students who acquired critical updates via text messages from close others or via Twitter reported increased exposure to conflicting information compared to those who did not rely on these channels (Table 3).

Table 3. Communication channels students consulted during the lockdown and their independent associations with exposure to conflicting information (n = 3,393)

Variables	***b*(95% CI)**	**SE$_b$**	***t***
Direct contact			
Text message from a campus group	.41(.34, .48)**	.03	11.49
Text message from friend	.29(.16, .42)**	.06	4.58
Text message from family	.09(.01, .18)*	.04	2.16
Phone call from a friend	-.09(-.17, -.01)*	.04	-2.15
Phone call from family	-.03(-.12, .05)	.04	-.81
Social media			
Twitter	.08(.01, .14)*	.03	2.49
Facebook	.07(-.002, .14)†	.03	1.90
Snapchat	.03(-.04, .11)	.03	.90
Instagram	.01(-.06, .07)	.03	.15
Reddit	.07(-.02, .16)	.04	1.52
Model Statistics	F(10, 3382) = 25.28; p< .001; R^2= .07		

*p < .05; **p < .001; †p = .057; all regression coefficients are standardized.

Discussion

A prominent feature of this campus shooting and lockdown event was the spread of misinformation and rumors during a 90-minute gap in official updates. Study 1a was the first of its kind to examine whether exposure to misinformation during a dangerous and uncertain crisis situation is associated with event-related distress. Moreover, it aimed to understand the extent to which trust in information channels from which information was received moderated the relationship between channel use and distress. When updates from officials are irregular or lacking, it is important to understand how people attempt to stay abreast of the situation in an information vacuum filled with speculation and rumors. Thus, these findings may inform emergency management officials to efficiently combat the generation and spread of misinformation and rumors, thereby reducing uncertainty and distress among those affected.

Study 1b: Introduction

Because students in this first study who used Twitter reported increased exposure to conflicting information, in Study 1b we examined the time-course of community-level rumor generation and virality (i.e., degree to which rumors were circulated) among a subset of Twitter users who followed two official university Twitter accounts. Prior research on rumor posits that situational ambiguity, high importance, and anxiety are necessary conditions for rumor generation (Allport & Postman, 1946; Festinger et al., 1948; Rosnow, 1991). Therefore, rumor generation and transmission during the lockdown were likely highest during the ~90-minute gap during the lockdown in which no updates were transmitted to students.

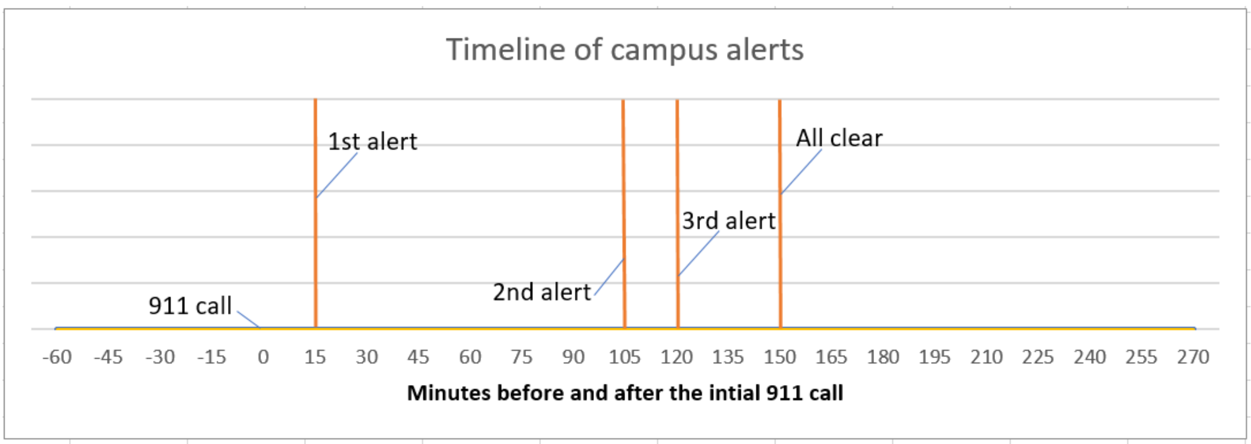

The following research questions were examined:

Research questions

RQ1: When updates are lacking, are rumors generated and transmitted on social media as the rumor literature suggests?

RQ2: Does event-related negative emotion follow the trajectory of rumor generation or rumor transmission?

Method

Twitter user selection

There are several challenges associated with searching for tweets in a geographic area using the Twitter API. Although Twitter does allow for searching based on geographic coordinates (geotags), only 1-5% of tweets are geotagged. Moreover, it is not currently possible to perform searches for tweets generated more than 3 days prior to the date of the search. To circumvent these challenges in the context of a campus shooting, researchers have relied on downloading tweets directly from followers of a university's Twitter account; the efficacy of this technique for approximating users who are likely to be students affiliated with the university has been demonstrated across several incidents of campus violence (Jones et al., 2016b). Thus, on the day of the shooting and lockdown, we used the twitteR package (Gentry, 2015) for R (R Core Team, 2016) to connect to the Twitter API and download the list of the most recent 20,000 followers of (or subscribers to) the university's main Twitter account, and 6,000 followers of the university's emergency management Twitter account.

We then removed users from this list based on the following criteria: non-English language account, private account (in which case tweets would not be publicly available), "verified" account (usually indicative of high-profile Twitter users or businesses), and accounts with more than 1,000 total tweets (to omit super users). After employing these exclusion criteria, 13,000 user accounts were available from which to pull tweets. Approximately two weeks after the shooting and lockdown, we interfaced with the API and requested the most recent 200 tweets from each user in our trimmed list.

Tweet collection and processing

Using R, we connected to Twitter via its Application Programming Interface (API) on the day of the lockdown and downloaded a list of the most recent 13,000 public followers of the university's primary and emergency response Twitter accounts. Two weeks later, we downloaded the most recent 200 tweets from each follower. Because tweets were time-stamped, we constrained our analysis to tweets generated in the hour leading up to the initial 911 call until the second hour after the lockdown was lifted, segmenting time into 15-minute blocks. Each tweet in this frame was tagged if it contained a rumor, defined as a statement verified to be blatantly false in the days following the incident. To capture community-level distress about the lockdown, we devised an R script to automatically tag tweets referencing the lockdown incident (for a similar method, see Jones et al., 2016b) and those containing negative emotion words using the Linguistic Inquiry and Word Counter (LIWC) negative emotion dictionary (Tausczik & Pennebaker, 2010). For example, tweets about the lockdown that also contained words like "distress" or "afraid" were tagged to reflect event-related negative emotions. Tagging tweets in this way allowed us to calculate a count of rumor tweets, as well as calculate the proportion of tweets with event-related negative emotions in each 15-minute segment over time. We then overlaid the official campus alerts sent to all university students during the same timeframe.

We downloaded nearly 2.3 million tweets. We then constrained our analysis to the time-frame immediately around the lockdown: approximately one hour before the 911 call up until the end of the second hour after the all-clear. Within this 5-hour window of time, we captured 11,617 tweets from 2,863 users. After removing duplicate tweets, we were left with 7,824 tweets from 2,515 users.

Measures

Rumor tweets. All tweets generated in the time-frame around the lockdown were manually coded for rumors. A coder was instructed to tag tweets in the sample that contained information that was not verified at the time of the lockdown. Given that virtually no information was available during the lockdown, aside from official reminders that the university was on lockdown, the task of identifying rumor tweets was relatively straightforward. In all, 38 tweets with rumors were identified.

Rumor virality. Every tweet downloaded via the Twitter API comes with a measure of how many times it was retweeted. This measure captures the virality of the tweet over its lifetime and is not tied to its virality at a given time point (e.g., during the lockdown). However, given the targeted nature of this event, virality was likely isolated to the lockdown as there would be no need to retweet any rumors about the lockdown after the "all clear".

Event-related negative emotion. We analyzed the linguistic content of each tweet using a custom R script that tallied the frequency with which words used in each tweet match words from the LIWC software's negative emotion dictionary. Similar to prior research (Jones et al., 2016*b*), we also employed an R script that used a 17-item custom word list to automatically identify and tag tweets about the shooting and lockdown. This list included context-specific words (e.g., lockdown, #[university name]strong, #prayfor[university name]) to bolster the script's efficacy in identifying lockdown-related tweets. Tweets containing at least one negative emotion word and one word referencing the event were coded with a 1 (all others coded with 0).

Analytic strategy

Data were imported into Stata 14 (College Station, TX) from R and tweets were combined into 15-minute blocks across time. The proportion of tweets with event-related negative emotion expression in each 15-minute block was calculated. We also calculated the quantity and virality (via retweet counts) of rumor tweets in each block, respectively. Event-related negative emotion expression was plotted across time and rumor generation count and rumor virality were overlaid, respectively, in two graphs.

Results

Within the corpus of tweets generated around the lockdown timeframe, 38 rumors were identified ($M_{\text{retweet count}}$= 179, SD = 427; min = 0, max = 2,299). Viral rumors (i.e., those retweeted most frequently) involved descriptions of a non-existent white male suspect and his movements. Other rumors involved claims of multiple deaths and warnings of multiple shooters at several locations on campus (Table 4).

Table 4. Example rumors and their virality

Text of retweeted rumors	***# retweets***
Description of perpetrator(s)	
• [user omitted] per scanner [university name omitted] suspects are male and female white male approx 6 ft tall	103
• [user omitted] [university omitted] shooting 2 victims per [university name omitted] newsroom campus on lockdown shooter described as 6ft white male wearing black	859
• [user omitted] police search for 6foot white male dressed in all black after two people shot dead at [university name omitted]	1,096
Warnings of multiple shooters and victims	
• [user omitted] multiple shooters on campus right now make sure to get into a safe place	9
• [user omitted] wtf multiple shooters people on stretchers 5 helicopters in the air im literally so scared right now	9
• [user omitted] 2 confirmed victims down multiple shooters on the loose been almost 30 min still not caught	15

As depicted in Figure 2, the bulk of rumors were generated during the 90-minute gap in communication from campus officials after the first lockdown alert and continued consistently until a second campus alert was disseminated to remind students about the lockdown.

Figure 2. Time-course of rumor generation, event-related negative emotions, and campus alerts in the hours before, during, and after the lockdown

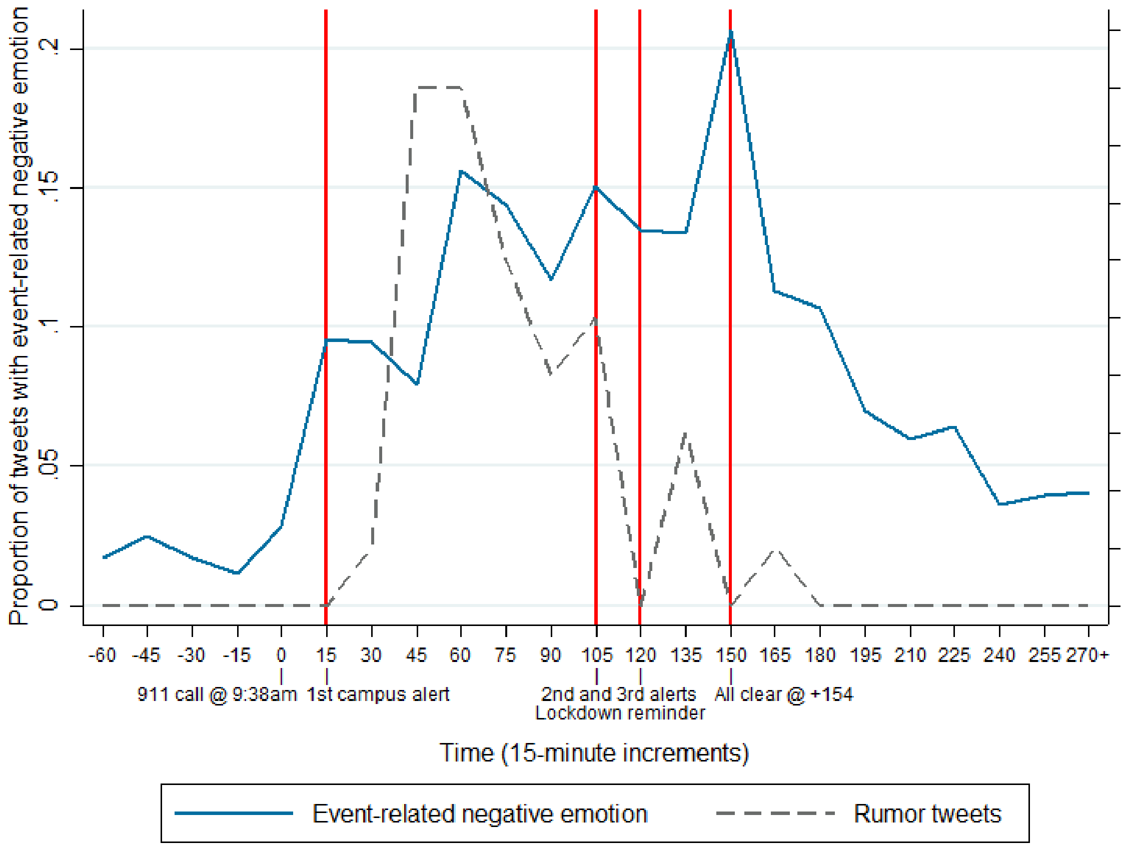

Although the number of rumors decreased after the second alert, the virality of the few rumors that were generated in the time-block after a third alert (again regarding the lockdown) far exceeded any rumor tweet in the preceding blocks (Figure 3). It is not clear why the rumors in this block went viral, but it might have been due to more people becoming aware of the situation and wanting to pass information on to others. Although we cannot directly link them together, event-related negative emotions tracked almost identically with rumor virality consistently across time (Figure 3). These findings suggest that while the amount of rumor generation over time may be informative for understanding the efficacy of rumor theory using social media data, the virality of rumors may also be implicated in the transmission of distress during a crisis.

Figure 3. Time-course of rumor virality, event-related negative emotions, and campus alerts in the hours before, during, and after the lockdown

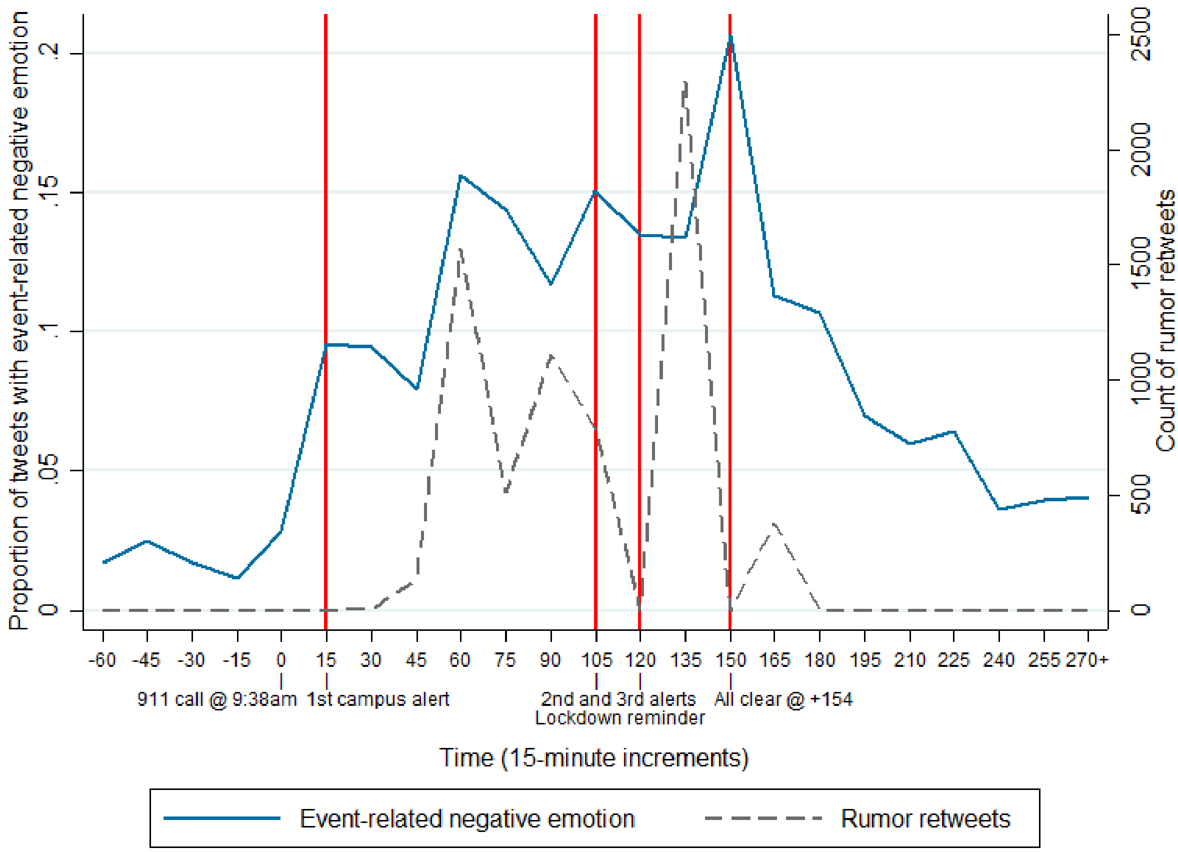

Discussion

Prior research on rumors demonstrates that situational ambiguity, high importance and anxiety are necessary conditions for rumor generation (Allport & Postman, 1946; Festinger et al., 1948; Rosnow, 1991). Consistent with this work, our findings indicate that during crisis events, when critical updates from official channels are irregular, rumors proliferate. Individuals who are caught in the path of a crisis event are often left feeling helpless and without situational control (Janoff-Bulman, 1992), which can lead people to see patterns in the information obtained that are not present (e.g., via illusory pattern perception; Whitson & Galinsky, 2008). In addition, situational stress may interfere with information processing via inhibited executive functioning (e.g., working memory, self-regulation processes; Henderson, Snyder, Gupta, & Banich, 2012). Taken together, these effects may diminish the myriad cognitive processes necessary for scrutinizing the veracity of unique and repeated information (Lewandowsky, Ecker, Seifert, Schwarz, & Cook, 2012), such as content propagated on social media platforms during a crisis. This causal chain likely plays a role in increasing the potency of rumor exposure during potentially threatening and ambiguous crisis situations.

To mitigate this problem, we offer several recommendations. First, emergency officials should disseminate frequent updates to the affected population, in real time. In the context of a school shooting, repeated alerts have been found to increase the perception of urgency among participants who received them (Stephens, Barrett, & Mahometa, 2013), a factor necessary for eliciting swift and appropriate action. Although we cannot explicitly test whether more frequent updates from official channels would have mitigated rumors and distress using the data we collected, crisis communication scholars posit that regular communications from emergency management officials are essential for mitigating uncertainty and rumors after a crisis (Seeger, Sellnow, & Ulmer, 2003). For example, as part of their response to a mass shooting at a shopping mall in Munich, Germany, local police urged the public via press conferences and social media to resist speculation about the attack and directly addressed rumors on social media as they became aware of them (Schuster, 2016). Despite false reports of additional shootings in the city and the overall lack of clarity about what was happening during the citywide lockdown, police chose to maintain transparency and constant contact with the public throughout the ordeal, a strategy likely appreciated by the public (Fischhoff, 2011). Had the Munich police remained silent, however, the budding rumors about shootings in other city locations would have likely filled the information void. As communities learn to manage active shooter crises and other emergencies, crisis communications like those employed by the Munich police department will be prudent to put in place.

Second, critical updates disseminated to the public should include new information, when possible. However, when new information is not available, updates should be tailored to reduce situational uncertainty (Reynolds & Seeger, 2005), thereby mitigating distress and rumors (Seeger, Sellnow, & Ulmer, 2003). Additionally, emergency management officials should attempt to counter the impact of rumors that arise during crisis situations by monitoring social media channels and encouraging individuals to keep a healthy skepticism about information coming from unofficial channels.

Furthermore, we believe the news media, which play a critical role in informing the public during crisis events (Seeger, Sellnow, & Ulmer, 2003), must share the responsibility for disseminating accurate information. The importance of this point is illustrated by the examples of conspiracy theories propagated on social media (and other channels) that resonate with individuals psychologically attuned to alternative narratives (Del Vicario et al., 2016). Although seemingly benign, conspiracy theories can lead people to deny that acts of horror, like the 9/11 terrorist attacks and the Sandy Hook Elementary School massacre, occurred at all. Consequently, denial narratives born from inconsistencies in news reporting can directly and negatively impact the individuals in communities devastated by these events (Wendling, 2017).

Although we examined the correlates of unofficial communication-channel use in our analyses in Study 1a, we acknowledge the important role of official channel use during a crisis. Unfortunately, because 96% of respondents in our sample indicated consulting official channels, and roughly 92% indicated trusting these channels somewhat or strongly, the lack of variation precluded our ability to include these variables in our statistical models. Also, we are unable to determine whether participants actively sought – or were passively exposed to – information from different channels. During the lockdown, participants may have sought information (especially during the 90-minute gap in communication from campus officials) by accessing news sites and social media, or they might have sent text messages to friends and family to see if they knew any details. However, students could have simultaneously received unsolicited messages via social media or text during this event, in which case their receipt of information could be considered passive. An additional limitation of Study 1a was that data collection occurred retrospectively (albeit soon after the event), and we did not employ real-time data collection methods during the lockdown (e.g., ecological momentary assessment), which would have been valuable for assessing exposure to conflicting information and distress responses. To compensate for this, we collected archival Twitter data – before, during, and after the lockdown – from thousands of users in Study 1b, providing supplemental data that occurred in real-time. This supplemental analysis of Twitter data fostered additional depth to our understanding of the crisis event we studied.

Conclusion

Exposure to rumors and conflicting information that arise out of the ambiguity of a crisis may have negative consequences for the people who receive and believe them. Moreover, the extent to which people trust the channels through which unofficial and conflicting information flow may exacerbate distress. Rumor generation during ambiguous crisis events is certain to continue. Therefore, social scientists should study the psychological impact of rumor exposure using methodological triangulation to understand the dynamic contextual features of and community responses to these events. Doing so will help to better elucidate the function and impact of crisis-related communications, or the lack thereof, on distress responses. Science on crisis communications and the media can be an ally in this challenging set of tasks.

References

Allport, G. W., & Postman, L. (1946). An analysis of rumor. *Public Opinion Quarterly*, *10*, 501–517.

American Psychiatric Association. (2013). *Diagnostic and statistical manual of mental disorders (5th ed.)*. Retrieved from http://dsm.psychiatryonline.org/book.aspx?bookid=556

Ball-Rokeach, S. J., & DeFleur, M. L. (1976). A dependency model of mass-media effects. *Communication Research*, *3*, 3–21. https://dx.doi.org/10.1177/009365027600300101

Blair, J. Pete, & Schweit, K. W. (2014). *A study of active shooter incidents in the United States between 2000 and 2013*. U.S. Department of Justice, Federal Bureau of Investigation, Washington D.C., Report.

Boyd, R. L. (2017). *MEH: Meaning Extraction Helper* (Version 1.4.15) [Computer software]. Retrieved from http://meh.ryanb.cc Accessed August 16, 2017

Bryant, R. A. (2016). *Acute Stress Disorder: What it is and how to treat it*. New York, NY: Guilford Publications.

Chen, C. Y., & Hong, R. Y. (2010). Intolerance of uncertainty moderates the relation between negative life events and anxiety. *Personality and Individual Differences*, *49*, 49–53. https://dx.doi.org/10.1016/j.paid.2010.03.006

Del Vicario, M., Bessi, A., Zollo, F., Petroni, F., Scala, A., Caldarelli, G., … Quattrociocchi, W. (2016). The spreading of misinformation online. *Proceedings of the National Academy of Sciences*, *113*, 554–559. https://dx.doi.org/10.1073/pnas.1517441113

Festinger, L., Cartwright, D., Barber, K., Fleischl, J., Gottsdanker, J., Keysen, A., & Leavitt, G. (1948). A study of a rumor: Its origin and spread. *Human Relations*, *1*, 464–486.

Fischhoff, B. (2011). Communicating about the risks of terrorism (or anything else). *American*

Psychologist, *66*, 520–531. https://dx.doi.org/10.1037/a0024570

Fitzhugh, S. M., Gibson, C. B., Spiro, E. S., & Butts, C. T. (2016). Spatio-temporal filtering techniques for the detection of disaster-related communication. *Social Science Research*, *59,* 137–154. https://dx.doi.org/10.1016/j.ssresearch.2016.04.023

Gentry, J. (2016). *twitteR: R based Twitter client* (Version 1.1.9) [Computer software]. Retrieved from https://cran.r-project.org/package=twittR

Henderson, R. K., Snyder, H. R., Gupta, T., & Banich, M. T. (2012). When does stress help or harm? The effects of stress controllability and subjective stress response on stroop performance. *Frontiers in Psychology*, *3,* 1–15. https://dx.doi.org/10.3389/fpsyg.2012.00179

Holman, E. A., Garfin, D. R., & Silver, R. C. (2014). Media's role in broadcasting acute stress following the Boston Marathon bombings. *Proceedings of the National Academy of Sciences*, *111*, 93–98. https://dx.doi.org/10.1073/pnas.1316265110

Hughes, A. L., & Palen, L. (2009). Twitter adoption and use in mass convergence and emergency events. *International Journal of Emergency Management*, *6,* 248–260.

Janoff-Bulman, R. (1992). *Shattered Assumptions: Toward a New Psychology Trauma*. New York, NY: The New Press.

Jones, N. M., Garfin, D. R., Holman, E. A., & Silver, R. C. (2016a). Media use and exposure to graphic content in the week following the Boston Marathon bombings. *American Journal of Community Psychology*, *58,* 47–59. https://dx.doi.org/10.1002/ajcp.12073

Jones, N. M., Wojcik, S. P., Sweeting, J., & Silver, R. C. (2016b). Tweeting negative emotion: An investigation of Twitter data in the aftermath of violence on college campuses. *Psychological Methods*, *21*, 526–541. https://dx.doi.org/10.1037/met0000099

Lewandowsky, S., Ecker, U. K., Seifert, C. M., Schwarz, N., & Cook, J. (2012). Misinformation and its correction: Continued influence and successful debiasing. *Psychological Science in the Public Interest*, *13*, 106–131. https://dx.doi.org/10.1177/1529100612451018

MacCallum, R. C., Zhang, S., Preacher, K. J., & Rucker, D. D. (2002). On the practice of dichotomization of quantitative variables. *Psychological Methods*, *7*, 19. https://dx.doi.org/10.1037/1082-989X.7.1.19

Mazer, J. P., Thompson, B., Cherry, J., Russell, M., Payne, H. J., Kirby, E. G., & Pfohl, W. (2015). Communication in the face of a school crisis: Examining the volume and content of social media mentions during active shooter incidents. *Computers in Human Behavior*, *53*, 238–248. https://dx.doi.org/10.1016/j.chb.2015.06.040

Palen, L., & Anderson, K. M. (2016). Crisis informatics—New data for extraordinary times. *Science*, *353*, 224–225. https://dx.doi.org/10.1126/science.aag2579

Palen, L., & Liu, S. B. (2007). Citizen communications in crisis: Anticipating a future of ICT-supported public participation. In *Proceedings of the SIGCHI conference on Human factors in computing systems* (pp. 727–736). New York: ACM Press. https://dx.doi.org/10.1145/1240624.1240736

R Core Team. (2018). *R: A Language and Environment for Statistical Computing*. Vienna, Austria: R Foundation for Statistical Computing. Retrieved from http://www.R-project.org/

Reynolds, B., & Seeger, M. W. (2005). Crisis and emergency risk communication as an integrative model. *Journal of Health Communication*, *10*, 43–55. https://dx.doi.org/10.1080/10810730590904571

Rosnow, R. L. (1991). Inside rumor: A personal journey. *American Psychologist*, *46*, 484–496. https://dx.doi.org/10.1037/0003-066X.46.5.484

Rosnow, R. L., Yost, J. H., & Esposito, J. L. (1986). Belief in rumor and likelihood of rumor transmission. *Language & Communication*, *6*, 189–194. https://dx.doi.org/10.1016/0271-5309(86)90022-4

Schlenger, W. E., Caddell, J. M., Ebert, L., Jordan, B. K., Rourke, K. M., Wilson, D., … Kulka, R. A. (2002). Psychological reactions to terrorist attacks: Findings from the National Study of Americans' Reactions to September 11. *Journal of the American Medical Association, 288*, 581–588. https://dx.doi.org/10.1001/jama.288.5.581

Seeger, M. W., Sellnow, T. L., & Ulmer, R. R. (2003). *Communication and organizational crisis*. Westport, CT: Greenwood Publishing Group.

Seeger, M. W., Vennette, S., Ulmer, R. R., & Sellnow, T. L. (2002). Media use, information seeking and reported needs in post crisis contexts. In B. S. Greenberg (Ed.), *Communication and Terrorism: Public and Media Responses to 9/11* (pp. 53–63). Cresskill, NJ: Hampton Press.

Shuster, S. (2016). *Munich police face questions over false tweets during shooting*. Retrieved from http://www.time.com/4421410/munich-shooting-germany-terrorism-tweets/

Silver, R. C., Holman, E. A., Andersen, J. P., Poulin, M., McIntosh, D. N., & Gil-Rivas, V. (2013). Mental- and physical-health effects of acute exposure to media images of the September 11, 2001, attacks and the Iraq War. *Psychological Science*, *24*, 1623–1634. https://dx.doi.org/10.1177/0956797612460406

Spiro, E. S., Fitzhugh, S., Sutton, J., Pierski, N., Greczek, M., & Butts, C. T. (2012). Rumoring during extreme events: A case study of Deepwater Horizon 2010. In *Proceedings of the 4th Annual ACM Web Science Conference* (pp. 275–283). New York, NY: ACM. https://dx.doi.org/10.1145/2380718.2380754

Starbird, K., Maddock, J., Orand, M., Achterman, P., & Mason, R. M. (2014). Rumors, false flags, and digital vigilantes: Misinformation on Twitter after the 2013 Boston Marathon bombing. In *iConference 2014 Proceedings* (pp. 654–662). Grandville, MI: iSchools. https://dx.doi.org/10.9776/14308

Stephens, K. K., Barrett, A. K., & Mahometa, M. J. (2013). Organizational communication in emergencies: Using multiple channels and sources to combat noise and capture attention. *Human Communication Research*, *39*, 230–251. https://dx.doi.org/10.1111/hcre.12002

Taha, S. A., Matheson, K., & Anisman, H. (2014). H1N1 was not all that scary: Uncertainty and stressor appraisals predict anxiety related to a coming viral threat. *Stress and Health*, *30*, 149–157. https://doi.org/10.1002/smi.2505

Tausczik, Y. R., & Pennebaker, J. W. (2010). The psychological meaning of words: LIWC and computerized text analysis methods. *Journal of Language and Social Psychology*, *29*, 24–54. https://dx.doi.org/10.1177/0261927X09351676

Wendling, M. (2017). *Sandy Hook to Trump: "Help us stop conspiracy theorists."* Retrieved from https://www.bbc.com/news/blogs-trending-39194035

Whitson, J. A., & Galinsky, A. D. (2008). Lacking control increases illusory pattern perception. *Science*, *322*, 115–117. https://dx.doi.org/10.1126/science.1159845

Acknowledgements

This chapter has been published in the peer-reviewed journal *Proceedings of the National Academy of Sciences*. I thank my coauthors for their invaluable contributions to this project. I also thank Maryann J. Gray, E. Alison Holman, and Dana Rose Garfin for their valuable input at the onset of Study 1a, and Carly Steinberger for her efforts coding tweets in Study 1b. I also thank Azim Shariff for his comments on an earlier version of this paper.

Jones, N. M., Thompson, R. R., Dunkel Schetter, C., & Silver, R. C. (2017). Distress and rumor exposure on social media during a campus lockdown. *Proceedings of the National Academy of Sciences*. *114*, 11663-11668 (published ahead of print October 17, 2017). https://dx.doi.org/10.1073/pnas.1708518114

Chapter 4: This is not a drill: Anxiety on Twitter following the 2018 Hawaii false missile alert

Abstract

The accuracy of emergency management agencies in transmitting alerts about dangerous threats to public safety is key for the protection of life and property. When alerts of imminent threats are believed to be real, uncontrollable, and impossible to escape, people who receive them often experience fear and anxiety as they await the threat's arrival (i.e., incubation of threat, Breznitz, 1968). However, what are the consequences when an alert turns out to be a false alarm? To evaluate this question, we explored anxiety on Twitter before, during, and after a false ballistic missile alert was transmitted across the State of Hawaii on the morning of January 13, 2018. In a 9-week window around the alert, we downloaded 1.2 million tweets representing 14,830 individuals to understand community- and individual-level anxiety expression among likely residents of Hawaii. Using piecewise regressions and change-point analyses at various timescales, we demonstrated that anxiety expressed on Twitter across the sample increased 4.6% on the day of the false alert and anxiety during the 38-minute alert period increased 3.4% every 15 minutes. In addition, users who expressed either low, medium, or high pre-alert anxiety exhibited differential anxiety responses post-alert, differential stabilization intervals (when anxiety stopped decreasing after the all-clear), and different post-alert baselines relative to their pre-alert levels. Taken together, findings suggest that false alarms of inescapable and dangerous threats are anxiety-provoking and that this anxiety can persist for many people. We offer several recommendations for how emergency management agencies should best communicate with the public after false alerts are transmitted.

Introduction

When danger is imminent, people often rely on local and state emergency management organizations for information to assess the severity of the threat and respond with appropriate measures to secure personal safety and protect property. Indeed, the ability of emergency management organizations to effectively communicate with populations at risk of threatening events plays an important role in saving lives (Rodríguez, Díaz, Santos, & Aguirre, 2007), and this is especially true for those who live under regular threat from natural disasters. However, when officials charged with disseminating information about impending threats falter, this might lead to a disaster (even if a threat becomes innocuous; Gilbert, 1998), and there may be a number of unintended consequences. For example, a lack of regular updates during a crisis can elicit negative psychological outcomes among individuals under threat (Jones, Thompson, Dunkel Schetter, & Silver, 2017).

False alarms are one example of how emergency management agencies falter. While relatively rare, false alarms occur when a transmitted warning of a large-scale impending threat (i.e., collective trauma) is no longer relevant, such as when a hurricane changes course and no longer threatens a geographic area. False alarms also occur when an emergency organization broadly transmits a warning about an impending threat that does not actually exist, such as false active shooter reports. The implications of false alarms of collective traumas have been studied by researchers across several disciplines. This body of work demonstrates how false alarms are related to loss of organizational credibility (Dow & Cutter, 1998; Ripberger et al., 2015), behavioral outcomes like diminished protective behavior among individuals under threat (Ripberger et al., 2015), and increased loss of life in tornado-prone areas with a high false-alarm ratio (Simmons & Sutter, 2009). However, little work has directly examined the psychological impact of exposure to potentially life-threatening events that turn out to be false. Can a false alarm of a collective trauma itself be a form of collective trauma? If so, how do people respond?

Research in psychology offers some clues about how individuals might respond to false alarms of collective traumas. Early experimental work on false alarms found that when the threat of a strong electric shock was perceived to be real, imminent, uncontrollable, and impossible to escape, individuals experienced a heightened physiological fear response (i.e., heightened heart rate) and subjective tension (Breznitz, 1984; Breznitz, 1985). During the moments in which participants anticipated the threat's arrival, subjective apprehension and worry increased over time (i.e., the incubation of threat; Breznitz, 1968, 1971, 1984). This work also demonstrated that when warnings of threat were cancelled, it took time for individuals to recover from their heightened state (Breznitz, 1984), suggesting that a cancellation of a threat warning does not immediately remedy the psychological consequences of being warned.

Attention to fear and worry in this work is reminiscent of studies that highlight anxiety as a key psychological reaction to collective traumas (e.g., terrorist attacks and natural disasters; Norris et al., 2002). For some individuals, lingering anxiety can cross clinical thresholds, developing into post-traumatic stress disorder (PTSD; Norris et al., 2002). Collective traumas often occur without warning and the threat they pose can be ambiguous. The inherent uncertainty in some disaster situations (Gilbert, 1998) can elicit psychological distress depending on how it is appraised by those who experience it (Folkman, Lazarus, Dunkel Schetter, DeLongis, & Gruen, 1986). Specifically, limited experience with a threat may increase situational ambiguity, thereby exacerbating anxiety (Lazarus, 1966). Thus, during dangerous or life-threatening situations in which information is lacking and ambiguity is high, uncertainty about the impending outcome may lead to anxiety (Taha, Matheson, & Anisman, 2014) and other negative psychological outcomes (Jones et al., 2017), especially among individuals for whom ambiguity is intolerable (Breznitz, 1984; Chen & Hong, 2010; Rosen, Knäuper, & Sammut, 2007).

When communities experience collective traumas, there is marked variation in how they respond, and these variations may depend on individual characteristics. For example, researchers have found that after 9/11, older adults had a steeper decline in their post-traumatic stress symptoms relative to younger individuals (Holman, Silver, Mogle, & Scott, 2016). Other work suggests that a number of individual-level psychological vulnerabilities present before a collective trauma may be related to post-event responses (cf. Seery, Holman, & Silver, 2010). Studies show that among both youth and adults, experiencing negative psychological states (e.g., anxiety) before a collective trauma is associated with an increased risk of PTSD in its aftermath (Asarnow et al., 1999; Nolan-Hoeksema & Morrow, 1991). Overall, these studies suggest that community-based studies of the impact of collective traumas should examine groups of individuals by their pre-event vulnerabilities when possible to assess whether post-event outcomes linger differentially over time.

Variation in psychological responses to collective trauma has been studied mostly using traditional research methodologies (e.g., surveys and interviews). However, rigorously studying the psychological impact of collective trauma is often difficult because of a lack of pre-event data and challenges entering the field in a timely manner (e.g., securing funding and ethics board approval), among other hurdles (Jones et al., 2016; Silver, 2004). Some social scientists have circumvented these challenges by using social media data to explicate psychological responses to collective trauma. Researchers have shown that Twitter data are particularly useful for evaluating community responses to school shootings (Doré, Ort, Braverman, & Ochsner, 2015; Jones et al., 2016; Jones et al., 2017), terrorist attacks (Gruebner et al., 2016; Lin, Margolin, & Wen, 2017), natural disasters (Gruebner et al., 2017; Murthy & Longwell, 2013) and other collective traumas (De Choudhury et al., 2014).

Twitter data offer unprecedented opportunities for theoretical insight because analyses of large-scale Twitter data possess several strengths compared to traditionally-collected data. First, access to Twitter data is free to anyone with some technical know-how. Second, these data are an ecologically valid observational data source, eliminating bias that may come with traditional survey methods (e.g., differential participation rates, demand characteristics). Third, the archival nature of Twitter data enables an examination of an emotion expression before and after collective traumas across an extended time frame that can be aggregated with granularity not possible with traditionally-collected data. Fourth, Twitter data can also be explicitly linked to a user's location via geo-coordinates if a user opts in to making this information public, or a user's location can be inferred based on the accounts they follow (see Jones et al., 2016).

The methodological strengths of working with Twitter data may also be useful for studying how individuals respond to false alarms of potentially life-threatening events. To date, the theoretical insights from experimental studies of false alarms have not been validated against large-scale, observational data. We harnessed Twitter data to explore anxiety responses to the 2018 Hawaii false ballistic missile alert. At 8:07am on January 13, 2018, Hawaii residents and visitors received an emergency alert from the Hawaii Emergency Management Agency on their smartphones stating that a ballistic missile was inbound to Hawaii, people should seek shelter, and that this alert was "NOT A DRILL" (caps in original). Media reports cited increased anxiety among residents during the ordeal (Silva, 2018) and some residents reached out to their loved ones to say goodbye (Pactol, 2018). Presumably many individuals thought that death or destruction was inevitable. However, a second message was transmitted 38 minutes later stating that there was no threat or danger, indicating that the original message had been a false alarm.

To study how individuals across the State of Hawaii responded to this event, we downloaded 1.2 million tweets generated by 14,830 Twitter users likely to be Hawaii residents over a period of time ranging from 6 weeks before (to establish a baseline) – to 18 days after (when data collection ended) – the alert. All tweets were coded for the presence of anxiety words from the Linguistic Inquiry and Word Count (LIWC; Pennebaker, Booth, Boyd, & Frances, 2015) anxiety dictionary. The data were examined at three different time-scales, outlined below, to understand the time course of anxiety expression across the entire sample of tweets, evaluate whether the incubation of threat could be observed using Twitter data, and explore whether users grouped into either low, medium, or high pre-alert anxiety expression displayed differential post-alert anxiety expression trends.

Hypotheses

H1: Consistent with other Twitter studies of collective trauma, anxiety will increase at the time of the alert and remain elevated for an undetermined period of time.

H2: Consistent with the incubation of threat phenomenon, anxiety will increase incrementally during the time from the transmission of the missile alert to the “all clear,” released 38 minutes later.

H3: Because prior research demonstrates that relative to others, low and high anxious individuals may be at greater risk for experiencing anxiety after a collective trauma (Asarnow et al., 1999; Nolan-Hoeksema & Morrow, 1991), we hypothesized that low and high anxiety users would exhibit the strongest increase in anxiety after the alert, stabilize later, and exhibit higher post-alert anxiety relative to their pre-alert baseline.

Method

Twitter data collection and measures

Procedures outlined in prior research (Jones et al., 2016) were used to source Twitter data generated by users likely to be Hawaii residents. First, Twitter accounts (n = 46) operated by local government and commercial organizations (e.g., city hall, local radio stations) that were likely to be followed by Hawaii residents were identified in the days following the alert. Next, the rtweet package (Kearney, 2017) for R Software (R Core Team, 2018) was used to interface with Twitter’s Application Programming Interface (API) and scrape the most recent 5,000 followers of each local account 18 days after the false missile alert. If an account had fewer than 5,000 followers, all followers were downloaded. After filtering out non-English language user accounts, user accounts created after December 2017 (because they would not have 6-weeks of pre-alert data), and both “private” and “verified” follower accounts (likely belonging to businesses, celebrities, and other public figures), a list of 32,239 user accounts was retained. This list was then read into an R script that interfaced with the Twitter API and downloaded the most recent 800 tweets generated by each user to scrape enough tweets dating back to six weeks before the alert. In all, 1.2 million tweets representing 14,830 users who tweeted within a 9-week window around the false alert (6 weeks before and 18 days after) were downloaded. Data collection procedures for this study were reviewed by the Institutional Review Board of the University of California, Irvine and assessed not to constitute human subjects research.

Measures

Anxiety expression. A custom R script was used to compare the words in each tweet to a list of 114 anxiety words available in the LIWC program. The words “threat” and “alarm” were removed from the dictionary because these words were specifically used to refer to the missile alert event. Each tweet was then coded dichotomously such that it was assigned a 1 if it contained at least one anxiety word (e.g., afraid, fear, scared; all others were coded 0). This protocol allowed for a proportion of tweets with anxiety to be calculated across analytic time frames in this study, and it provides a measure of anxiety expression that compensates for differential counts of tweets generated at each time-unit of analysis (e.g., minutes, hours, days).

Pre-alert anxiety. Users were grouped based on their level of pre-alert anxiety. Each user’s average proportion of pre-alert anxiety expression across the seven days before the alert was calculated. These proportions were then standardized across all users the z-scores were used to group users based on where they fell along this distribution. Users in the “low” group (*n* = 6,849) did not express any anxiety on Twitter before the alert. Users placed into the “medium” group (*n* = 1,394) expressed some anxiety, up to just below +1 *SD*; users in the “high” anxiety group (*n* = 503) expressed anxiety greater than +1 *SD*, relative to all other users.

Analytic strategy

Data were manipulated in R using tidytext (Silge & Robinson, 2016) and descriptive visualizations were created in R using ggplot2 (Whickam, 2009) and employed the generalized additive model (GAM) smoothing function to depict a non-linear line-of-best-fit across time.

Trajectories of anxiety before and after the alert were evaluated at several time-scales. To estimate linear trajectories of anxiety during these time scales and evaluate the significance of changes in anxiety during and after the alert period (relative to baseline), several piecewise regressions with discontinuity analyses were conducted in Stata 14.2 (College Station, TX) using procedures outlined by others (Jones et al., 2016; Mitchell, 2012). This approach is well-suited for statistically evaluating non-linear changes in time-series data for which points of change are expected. Knots are placed at specific time points to demarcate conceptual points of change in the time-series data and are included in a regression analysis to model trends before and after the knot. Additionally, in this application of a piecewise regression, the use of a knot placed at the time of the alert allows for a regression discontinuity analysis in which a predicted value of anxiety is calculated at the knot, as if the event had not occurred. This value can then be statistically compared to a value representative of anxiety that actually occurred at that knot-point. Importantly, these analyses were clustered around each user to account for within-user propensity for anxiety expression. The effect of the all-clear was evaluated in only a single time frame because the initial alert and the all-clear occurred within the same hour (38-minutes apart), and only one analysis was fine-grained enough to capture anxiety at this scale (e.g., 12-hour window).

9-week window. The proportion of daily anxiety was calculated across the 6 weeks preceding the alert and the 18 days that followed. Before attempting to model trajectories of anxiety across time with a piecewise regression, a change-point analysis was conducted in R using the *changepoint* package (Killick, Haynes, & Eckley, 2016) to determine the discrete time-point at which anxiety decreased to a stable level after the missile alert. This method employs the pruned exact linear time (PELT) algorithm (Killick & Eckley, 2014) to identify when the mean and variance of a variable deviate over time. In other words, the algorithm evaluates the mean and variance in a block at the start of time, and then evaluates whether they are significantly different from the mean and variance calculated in the next block of time. The package then outputs the time values where the algorithm identified significant changes in daily proportions of anxiety.

The change-point analysis indicated that anxiety increased on the day of the missile alert and stabilized two days later. We then conducted a piecewise regression placing a knot at the moment the alert was transmitted (January 13, 2018 at ~8:07a.m. HST) to mark the point at which anxiety was expected to increase and another was placed two days later.

12-hour window. Time across a period spanning six hours before and after the missile alert was parsed into 15-minute intervals (cf., Jones et al., 2017) and a proportion of tweets with anxiety was calculated for each interval. A piecewise regression analysis was conducted on all tweets in this time frame to evaluate whether anxiety increased the moment the alert was transmitted and the extent to which anxiety continued to increase during the 38-minute alert period. A knot was placed at the moment the alert was transmitted, and 45 minutes later (roughly 7 minutes after the all-clear was transmitted). In addition to evaluating immediate changes in anxiety resulting from the missile alert, this approach also allowed for an analysis of the slope of anxiety expression during the 38-minute alert period.

14-day window. Hourly proportions of tweets with anxiety generated in a 14-day window around the alert (7 days before and after) were calculated for each pre-alert anxiety group (i.e., low, medium, high) to provide a more fine-grained analysis of the immediate impact of the alert and its sustained effect across the following week for each group. To determine the discrete time-point at which anxiety leveled off for each group, a change-point analysis was conducted for each group individually. Post-alert change points were used to place a second knot in each piecewise regression to accurately model changes in anxiety over time for each group. In sum, a knot was placed at the moment the alert was transmitted and at the group-specific change point after the alert, when anxiety stabilized.

To evaluate whether the event had a lasting effect on each group in the days that followed, we conducted three ordinary least squares (OLS) regression analyses in which each group's pre-alert average proportion of anxiety expression (baseline anxiety) was compared to its post-alert average (new baseline), after stabilization in anxiety occurred.

Results

First, we examined daily proportions of anxiety across the full range of available data. We hypothesized that the day of the missile alert would be marked by a stark increase in anxiety among Twitter users across the state (H1). We found that across all users, anxiety increased 4.6% (standardized $b = .25$, $SE = .01$, $p < .001$, $t = 19.62$) on the day of the alert (see Figure 1).

Figure 1. Daily anxiety among Hawaii Twitter users 6 weeks before and 18 days after the false missile alert (n_{tweets} = 1.2 million; n_{users} = 14,830)

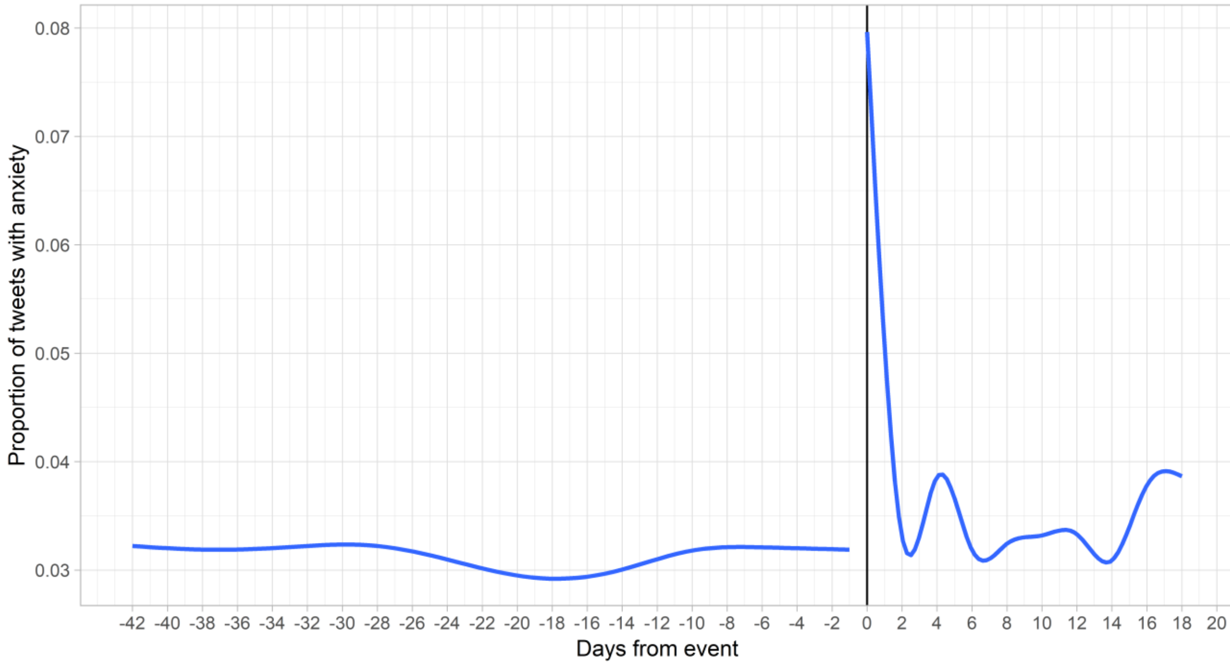

We further hypothesized that anxiety would increase incrementally during the period from the release of the alert to the "all clear" 38 minutes later, during which time users awaited the attack (H2; Breznitz, 1968, 1984). Consistent with the incubation of threat hypothesis, the results from 4,415 users (20,338 tweets) who tweeted in this time frame indicated that anxiety increased 3.4% in each 15-minute block during the alert period (standardized b = .12, SE = .02, t = 5.90, p < .001; see Figure 2), until the all-clear was transmitted.

Figure 2. Anxiety on Twitter, 6 hours before and after the ballistic missile alert (n_{tweets} = 20,338; n_{users} = 4,415)

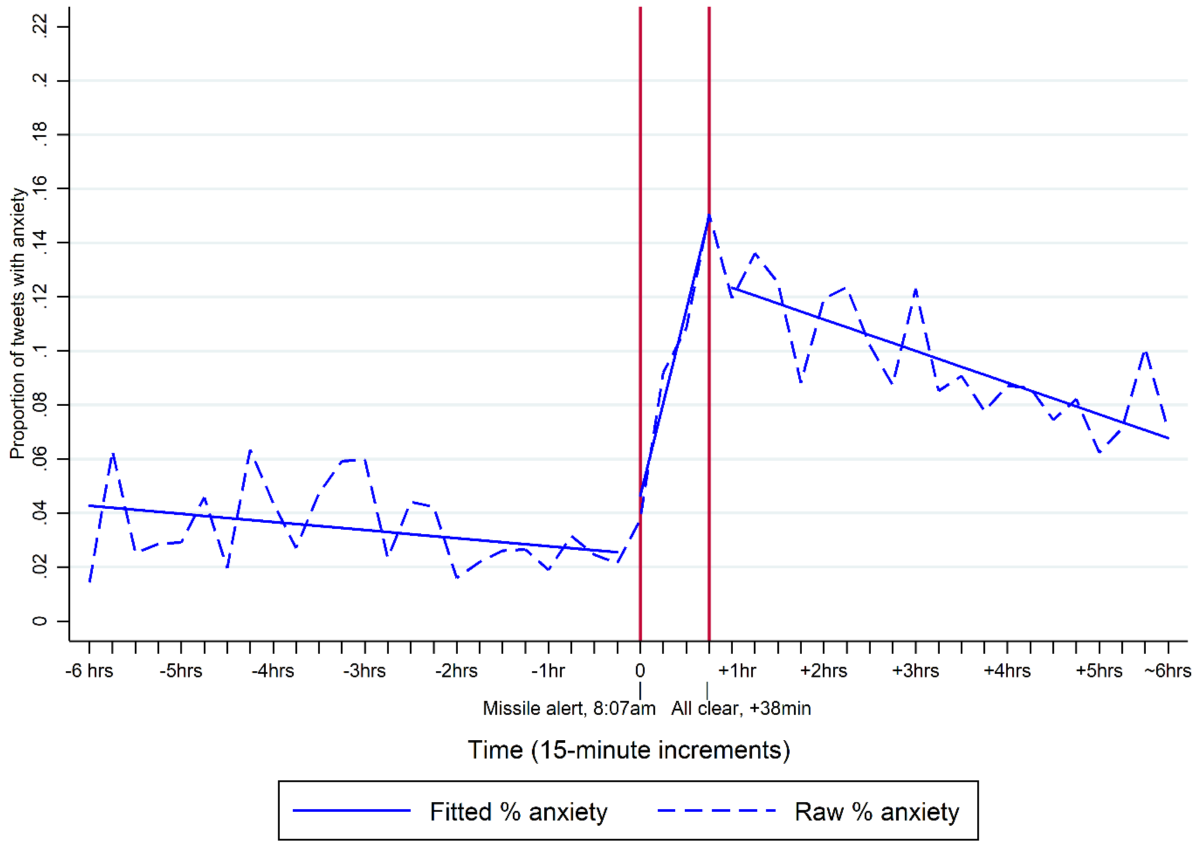

We also explored whether users placed into either a pre-alert low, medium, or high anxiety group displayed differential levels of anxiety at the onset of the alert, stabilization rates after it, and post-event mean-levels of anxiety expression relative to their pre-alert baselines. For each pre-alert anxiety group, we first used a change-point analysis to determine the discrete post-alert hour during which anxiety stabilized. After obtaining each group's stabilization point, we were then able to follow up with separate piecewise regressions in which a knot was placed at the onset of the alert and the point at which each group's anxiety expression stabilized.

In a window seven days before and after the missile alert, we found that for the low pre-alert anxiety group, anxiety increased 9.5% at the onset of the alert period (standardized *b* = .49, *SE* = .02, *t* = 22.87, *p* < .001). The medium pre-alert anxiety group also increased in anxiety expression (5.8%; standardized *b* = .30, *SE* = .03, *t* = 8.34, *p* < .001). Although we had hypothesized that the high pre-alert anxiety group would also increase, this group's anxiety expression decreased 8.8% (standardized *b* = -.46, *SE* = .07, *t* = -5.93, *p* < .001) at the onset of the alert.

Figure 3. Hourly anxiety by Twitter users with low, medium, and high pre-event anxiety, 7 days before and 7 days after the false missile alert (n_{tweets} = 324,010; n_{users} = 8,746)

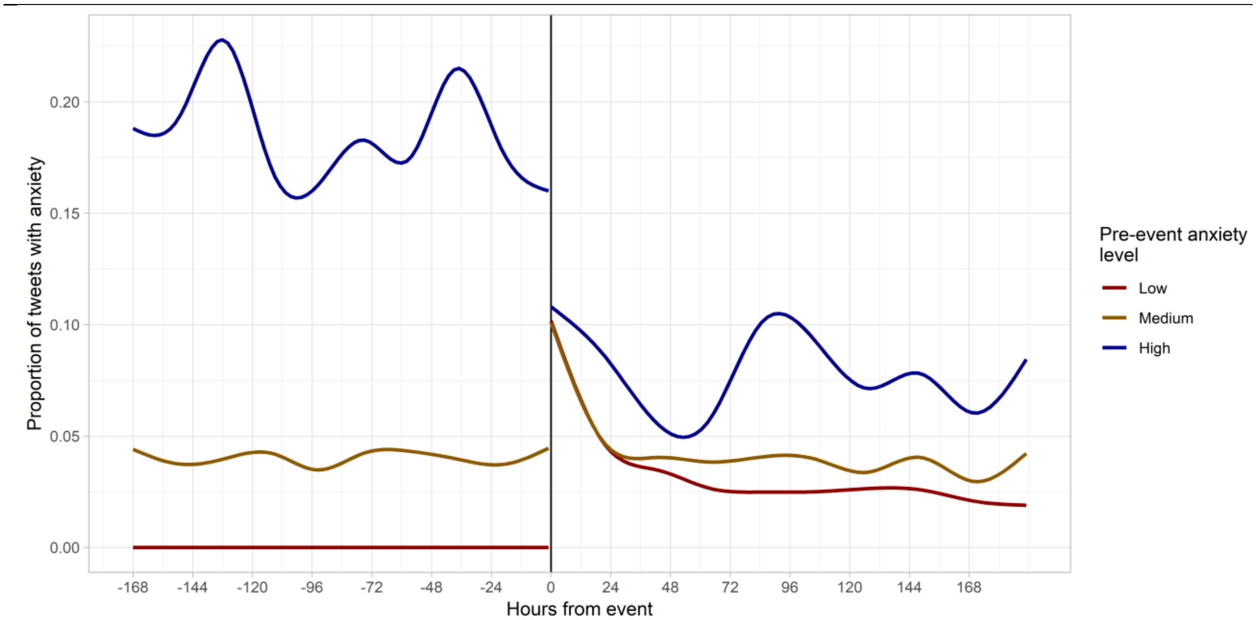

Post-event stabilization rates also differed (see Figure 4). As expected, the low pre-alert anxiety group took the longest to stabilize (~41 hours after the false missile alert). Contrary to our hypothesis, however, the high pre-alert anxiety group stabilized immediately after the alert was transmitted, the same point at which we observed a significant drop in anxiety for this group; the medium pre-alert anxiety group stabilized 23 hours post-event.

We also evaluated whether the post-alert baseline (after stabilization) differed from the pre-alert baseline for each group. The low pre-alert anxiety group, for whom anxiety expression before the alert was zero across a seven-day period, exhibited a new anxiety-baseline that was 2.5% higher than its pre-event baseline (standardized $b = .13$, $SE = .004$, $t = 28.15$, $p < .001$). The medium pre-alert anxiety group's new baseline was less than 1.0% lower than its pre-event baseline (standardized $b = -.01$, $SE = .005$, $t = -2.11$, $p = .035$). The high pre-alert anxiety group exhibited a new baseline that was 10.5% lower than its pre-event baseline (standardized $b = -.55$, $SE = .05$, $t = -10.76$, $p < .001$).

Figure 4. Change point analyses for each pre-alert anxiety group where time 169 (blue line) is the moment the alert was transmitted; a) low pre-alert anxiety b) medium pre-alert anxiety, c) high pre-alert anxiety

a)

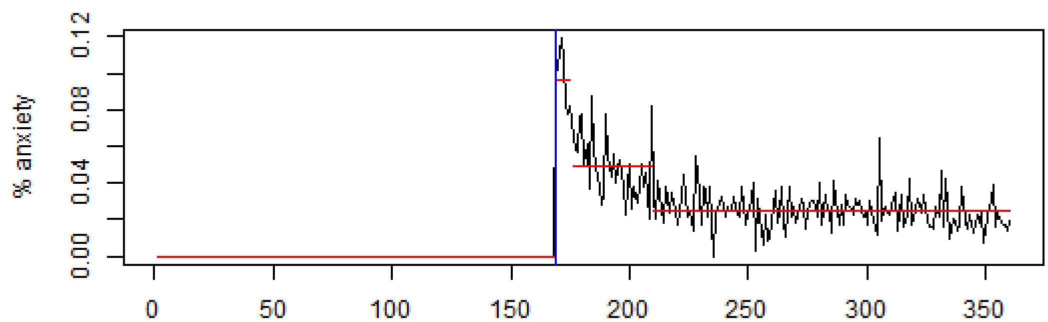

b)

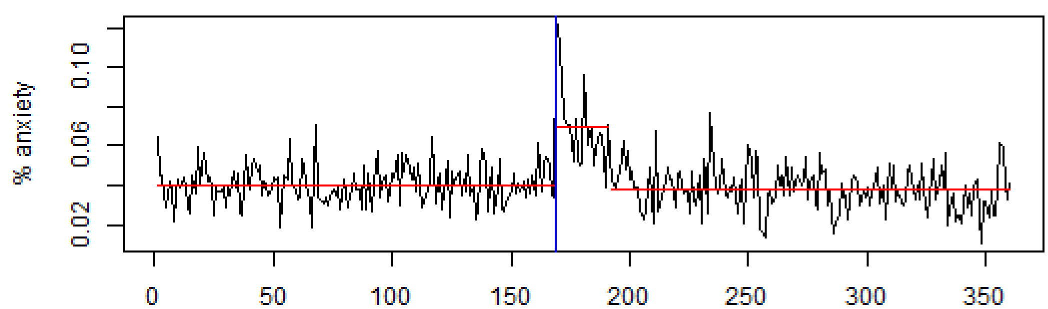

c)

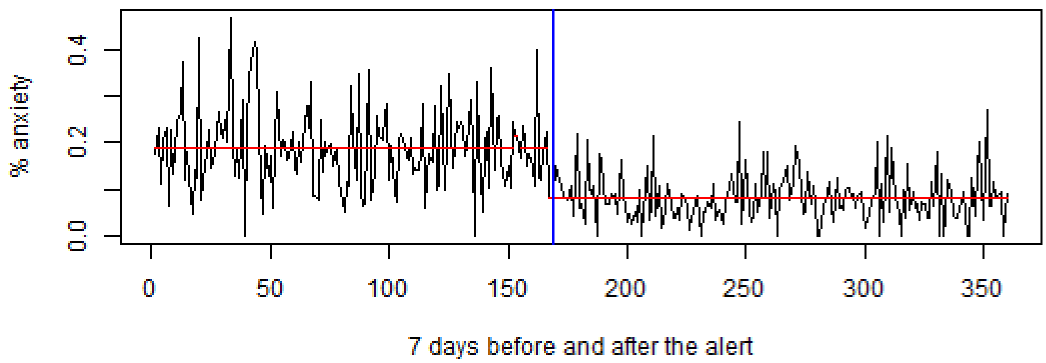

Discussion

Using a big data approach, our analyses provide evidence that the experience of a novel false alarm about an inescapable, life-threatening event is anxiety-provoking. Consistent with other studies using Twitter data analyzed in the context of a collective trauma, results indicate a marked increase in anxiety on the day of the event. The initial jump in anxiety observed is not novel in this regard, but it is the first time such an effect has been observed in the context of a false alarm using social media data.

In the most fine-grained level of analysis, results are consistent with the incubation of threat hypothesis (Breznitz, 1968, 1984), which states that anxiety experienced in anticipation of a threat will increase in the moments during which one waits for the threat to arrive. We found that anxiety increased 3.4% every 15 minutes until the transmission of an alert 38 minutes later reporting that the initial alert was a false alarm. Surprisingly, this trend was not thwarted by corrective tweets, explaining that the missile alert was a mistake, posted by the state's emergency management agency and by a local congressional representative shortly after the initial missile alert was transmitted. Despite these corrections being retweeted by at least 35,000 users, the incubation of threat phenomenon persisted. This is likely the case because users in our sample either a) did not see the messages dispelling the threat, or b) saw them but did not believe them, as evidenced by a user who tweeted: "*I'm still sheltering twenty minutes [later] who knows who's right?*"

Alternatively, this trend could simply reflect the fact that unique users became aware of the missile alert at different points in time during the alert period. For example, if a large number of users woke up to the alert 35 minutes into the ordeal and generated tweets expressing anxiety, this would explain why we observed a positive trend appeared during the alert period. This alternative explanation was evaluated among 556 users who tweeted in at least two of the three 15-minute intervals during the alert period. Consistent with the incubation of threat hypothesis, we found that anxiety increased incrementally across the alert period among these users as well. Overall, this suggests that evaluating tweets generated in the alert period, in aggregate, sufficiently captured the experience of the incubation of threat for the sample as a whole.

Consistent with early work suggesting that the cancellation of a threat does not immediately remedy reactions to a threatening situation (Breznitz, 1984), our results suggest that the experience of a false alarm may have a lingering impact on some individuals well after the threat is dispelled. For the group of users in this sample that did not express any anxiety before the event (i.e., the low pre-event anxiety group), anxiety increased the most and lingered the longest, relative to other groups, before stabilizing to a new baseline level 2.5% higher than what it was before the missile alert. Insofar as anxiety expression on Twitter can be assumed to be reflective of a user's life experience (Jones et al., 2016), this pattern is consistent with evidence demonstrating that people who are likely to have had lives devoid of psychologically impactful negative experiences are at increased risk of negative psychological outcomes following a traumatic event (Seery, Holman, & Silver, 2010). Moreover, the lingering presence of anxiety well after the threat was dispelled may be driven by some users in this group engaging in perseverative cognition, or “the repeated or chronic activation of the cognitive representation of one or more psychological stressors” (Brosschot, Gerin, & Thayer, 2006).

Research also suggests that individuals who have had more than an average number of negative life experiences, perhaps reflective of the high pre-alert anxiety group, should also be at increased risk for negative psychological outcomes after a traumatic event (Seery, Holman, & Silver, 2010). However, our results show a markedly different pattern. The high pre-alert anxiety group exhibited decrease in anxiety following alert period, declining to a new baseline level 10.5% lower than the group’s pre-event baseline. It is not readily clear why this group’s pattern was different than had been expected. This pattern might reflect a process in which the threat of death via a ballistic missile put disruptive daily stressors into perspective, thereby mitigating the anxiety individuals in this group typically express. Moreover, the realization that death was not imminent (after the all clear) may have been an emotional boon for users in this group compared to the others.

These suppositions are consistent with the notion of a leveling theory of human adaptation to life events. In this view, people react differently to positive and negative experiences and how one responds to these events depends on the totality of a person’s life. For example, hedonic leveling (Lucas, Clark, Georgellis, & Diener, 2003) posits that happy individuals have less to gain when they experience positive life events (e.g., marriage), whereas unhappy people have more to gain from positive life events. These findings are also reminiscent of adaptation level theory (Brickman, Coates, Janoff-Bulman, 1978), which suggests the importance of the contrast between what life after a significant event feels like compared to what life was like before. If these theoretical perspectives work similarly with respect to negative life events (e.g., the false missile alert), anxious individuals may have more to appreciate when they experience a near-miss and thus express less anxiety on social media after having “survived” what would have likely been construed as a deadly situation. This notion is also consistent with downward counterfactual thinking (Byrne, 2016) in which users in the high anxiety group may have recognized how much worse things could have been had the missile threat been real. This unexpected pattern may also be indicative of near-miss relief (Sweeney & Vohs, 2012), a phenomenon observed when an aversive event is avoided. The extent to which any of these explanations are correct is unknown and further work is needed to understand why this pattern emerged.

Although our results are similar to other studies of Twitter data that reveal increased negative emotion after a collective trauma, we acknowledge several limitations. Twitter data are not typically representative of the population, as they tend to be younger (aged 18-29) and from urban locations (Pew Research Center, 2018). The tweets we have collected, while informative, are not linked directly to any person-level psychological data (e.g., personality, history of negative life events). Although this precludes our ability to group users based on self-report characteristics, there is a tradition in psychology of relying on the words people use to provide a window into their psychological state (for a review, see Pennebaker, Mehl, & Niederhoffer, 2003; Tausczik & Pennebaker, 2010). Therefore, we consider a person’s pre-alert word usage pattern to be an acceptable proxy when other information is lacking. We also acknowledge that there is no guarantee that all users in our sample are residents of Hawaii; we did not rely on geolocation data because less than one percent (.07%) of tweets in our data were geocoded. However, the method we used to source locally-generated Twitter data has performed well in other studies of collective trauma (see Jones et al., 2016; Jones et al., 2017). In addition, the presence heightened anxiety, despite the error inherent this method, suggests that the effects we demonstrate would likely be even greater had we only captured tweets generated by residents.

Despite these limitations, the results lead to several recommendations for mitigating the psychological impact of impeding threats, false or otherwise. Early theorizing on false alarms posited the *false alarm* or "*cry wolf*" effect in which individuals may not believe the next threat warning because they lose faith in the credibility of systems or agencies responsible for disseminating them (Breznitz, 1984). Empirical work supports this. For example, residents of tornado-prone areas who perceive their local false alarm ratio to be high are less trusting of the National Weather Service (Ripberger, et al., 2015). Other work indirectly shows that such perceptions can be deadly (Simmons & Sutter, 2009). Such credibility loss is particularly important to combat in risk-prone locales like Hawaii, in which residents remain in targeting range of a ballistic missile from North Korea, live near active volcanic areas, and sometimes experience destructive weather events. In such a locale, the public's *continued* reliance on and trust in emergency management systems for information about impending danger are paramount to maintain. When emergency systems falter, research shows that credibility loss can be mitigated by a clear explanation of why the false alarm occurred in the first place (Breznitz, 1984). In the days and weeks following the Hawaii false missile alert, the media reported that the alert was sent in error by an employee with whom the agency had several past issues. This reporting likely raised questions about the agency's disciplinary procedures but may have assured the public that the entire affair was a fluke. Further reporting assured the public that new safeguards were in place to prevent one person from having the authority to transmit any message on the state-wide emergency system (Wamsley, 2018).

Although indications from this study suggest that social media channels were underutilized, we believe they can be an important channel through which critical updates are transmitted. For example, many users in our sample wondered why the all-clear alert took 38 minutes to transmit. As reporting about the missile alert unfolded, it came to light that the employee who transmitted the alert refused to cancel it (Wamsley, 2018). To circumvent this delay, the agency instead used its Twitter account to issue a statement that the alert was sent in error. These measures were enacted shortly after the alert was transmitted and undoubtedly prevented widespread panic. However, we know from users in our dataset that not everyone saw those corrections, although extensive media reporting highlighted the fact that corrections were sent over Twitter. We speculate that this experience likely led residents to follow social media accounts belonging to local emergency management agencies as people realized that in the future they may find critical updates on social media platforms. Indeed, the value of using social media accounts, in tandem with outlined procedures, to transmit critical updates, has been acknowledged by the Department of Homeland Security for almost a decade (Silver & Fischhoff, 2011). Agencies should also increase public outreach efforts to ensure that community members are connected to critical information channels during a crisis. Such measures would serve to bolster trust between the public and reporting agencies. We also believe it is good advice for the public to seek out and follow verified social media accounts belonging to their local emergency management agencies as soon as possible so that if another unfortunate event like this occurs (and it happens to be real), they will have access to the most up-to-date information about the potential threat.

In addition to revealing new information channels to follow, false alarms may also raise awareness of potential threats. For example, some individuals in our sample expressed the realization of not being prepared, for example: "D*efinitely scary. I live in a small studio a few blocks from the beach I was like take shelter where? I'm definitely not prepared*." In the months before the false alert, aggressive rhetoric between the American president and leader of North Korea heightened fears that a conflict might be imminent. As a result, in December 2017, Hawaii began performing emergency drills and siren testing to prepare for such an eventuality, although what to do during an attack was not clear to all residents (Kelkar, 2018). In our sample, many users expressed confusion about what actions to take and were unsure where to seek shelter when instructed to do so. It is prudent for members of the public to educate themselves about emergency preparations and for emergency management agencies to use all available channels to transmit concrete recommendations for protective action (Fischhoff, 2011).

Insights from our data also highlight the importance of the interface between emergency management agencies and the news media. Users in our sample searched for information on traditional media channels to no avail, evidenced by tweets like: "*missile alert in Hawaii but no news coverage*" and "*...nothing on the news. looked it up on twitter and people are as confused as we are...*" These tweets highlight a crucial point about the importance of information dissemination via the news media during a crisis. The public relies on the media for updates when a crisis unfolds. In the absence of critical updates from emergency management agencies and the news media, people will turn to alternative channels (e.g., Twitter) to understand the nature of a threat. This is consistent with other studies of ambiguous and threatening situations that have shown that people turn to social media channels for critical updates and are sometimes exposed to conflicting information (Jones et al., 2017). When a crisis like this unfolds, it is vitally important that emergency management officials reach out to the media to signal actions the members public should take or, in Hawaii's case, issue corrections that can quickly be disseminated broadly to a worried public.

Free and open access to public Twitter data, coupled with Hawaii's false missile alert, provided an opportunity to study, for the first time, how several thousand people responded psychologically to the threat of an inescapable, impending tragedy. While it is fortunate we were able to study this phenomenon without loss of life, we show that for many users, the anxiety elicited by this false alarm lingered well beyond the assurance that the threat was not real, which may have health consequences over time for some individuals (Holman et al., 2008). Thus, our results reveal how potently frightening a crisis period can be and highlight how intense experiences like this may have lasting effects which become even clearer after disaggregating users by their pre-event psychological state. This event serves as an example of how accidental crisis communication can become a disaster (Gilbert, 1998), and should inform emergency management agencies how they can better serve the communities they are charged with informing.

References

Asarnow, J., Glynn, S., Pynoos, R. S., Nahum, J., Guthrie, D., Cantwell, D. P., & Franklin, B. (1999). When the earth stops shaking: Earthquake sequelae among children diagnosed for pre-earthquake psychopathology. *Journal of the American Academy of Child & Adolescent Psychiatry*, *38*, 1016–1023. http://dx.doi.org/10.1097/00004583-199908000-00018

Breznitz, S. (1968). "Incubation of threat" in a situation of conflicting expectations. *Psychological Reports*, *22*, 755–756.

Breznitz, S. (1971). A study of worrying. *British Journal of Social and Clinical Psychology*, 10, 271-279. http://dx.doi.org/10.1111/j.2044-8260.1971.tb00747.x

Breznitz, S. (1984). *Cry wolf: The psychology of false alarms*. Hillsdale, NJ: Lawrence Erlbaum Associates, Inc.

Breznitz, S. (1985). False alarms: Their effects on fear and adjustment. *Issues in Mental Health Nursing*, *7*, 335–348. http://dx.doi.org/10.3109/01612848509009460

Brickman, P., Coates, D., & Janoff-Bulman, R. (1978). Lottery winners and accident victims: Is happiness relative? *Journal of Personality and Social Psychology*, *36*, 917–927.

Brosschot, J. F., Gerin, W., & Thayer, J. F. (2006). The perseverative cognition hypothesis: A review of worry, prolonged stress-related physiological activation, and health. *Journal of Psychosomatic Research*, *60*, 113–124. http://dx.doi.org/10.1016/j.jpsychores.2005.06.074

Byrne, R. M. (2016). Counterfactual thought. *Annual Review of Psychology*, *67*, 135-157.

Chen, C. Y., & Hong, R. Y. (2010). Intolerance of uncertainty moderates the relation between negative life events and anxiety. *Personality and Individual Differences*, *49*, 49–53. http://dx.doi.org/10.1016/j.paid.2010.03.006

De Choudhury, M., Monroy-Hernandez, A., & Mark, G. (2014). "Narco" emotions: Affect and desensitization in social media during the Mexican drug war. In *Proceedings of the 32nd Annual ACM Conference on Human Factors in Computing Systems* (pp. 3563–3572). New York, NY: ACM. http://dx.doi.org/10.1145/2556288.2557197

Doré, B., Ort, L., Braverman, O., & Ochsner, K. N. (2015). Sadness shifts to anxiety over time and distance from the national tragedy in Newtown, Connecticut. *Psychological Science*, *26*, 363–373. http://dx.doi.org/10.1177/0956797614562218

Dow, K., & Cutter, S. L. (1998). Crying wolf: Repeat responses to hurricane evacuation orders. *Coastal Management*, *26*, 237–252. http://dx.doi.org/10.1080/08920759809362356

Folkman, S., Lazarus, R. S., Dunkel-Schetter, C., DeLongis, A., & Gruen, R. J. (1986). Dynamics of a stressful encounter: Cognitive appraisal, coping, and encounter outcomes. *Journal of Personality and Social Psychology*, *50*, 992–1003. http://dx.doi.org/10.1037/0022-3514.50.5.992

Gilbert, C. (1998). Studying disaster. In E. L. Quarantelli (Ed.), *What is a disaster? Perspectives on the question* (pp. 11–18). New York, NY: Routledge.

Gruebner, O., Lowe, S. R., Sykora, M., Shankardass, K., Subramanian, S., & Galea, S. (2017). A novel surveillance approach for disaster mental health. *PLoS ONE*, *12*, e0181233. http://dx.doi.org/10.1371/journal.pone.0181233

Gruebner, O., Sykora, M., Lowe, S. R., Shankardass, K., Trinquart, L., Jackson, T., … Galea, S. (2016). Mental health surveillance after the terrorist attacks in Paris. *The Lancet*, *387*, 2195–2196. http://dx.doi.org/10.1016/S0140-6736(16)30602-X

Holman, E. A., Silver, R. C., Mogle, J. A., & Scott, S. B. (2016). Adversity, time, and well-being: A longitudinal analysis of time perspective in adulthood. *Psychology and Aging*, *31*, 640–651. http://dx.doi.org/10.1037/pag0000115

Holman, E. A., Silver, R. C., Poulin, M., Andersen, J., Gil-Rivas, V., & McIntosh, D. N. (2008). Terrorism, acute stress, and cardiovascular health: A 3-year national study following the September 11th attacks. *Archives of General Psychiatry*, *65*, 73–80. http://dx.doi.org/10.1001/archgenpsychiatry.2007.6

Jones, N. M., Thompson, R. R., Dunkel Schetter, C., & Silver, R. C. (2017). Distress and rumor exposure on social media during a campus lockdown. *Proceedings of the National Academy of Sciences*. *114*, 11663-11668 (published ahead of print October 17, 2017). http://dx.doi.org/10.1073/pnas.1708518114

Jones, N. M., Wojcik, S. P., Sweeting, J., & Silver, R. C. (2016). Tweeting negative emotion: An investigation of Twitter data in the aftermath of violence on college campuses. *Psychological Methods*, *21*, 526–541. http://dx.doi.org/10.1037/met0000099

Kearney, M. W. (2017). *rtweet: Collecting Twitter Data*. Retrieved from https://cran.r-project.org/package=rtweet

Kelkar, K. (2018). *Hawaii sounds first nuclear warning siren since Cold War*. Retrieved from https://www.pbs.org/newshour/nation/hawaii-sounds-first-nuclear-warning-siren-since-cold-war

Killick, R., & Eckley, I. A. (2014). Changepoint: An R package for changepoint analysis. *Journal of Statistical Software*, *58*, 1–19. http://dx.doi.org/10.18637/jss.v058.i03

Killick, R., Haynes, K., & Eckley, I. A. (2016). *Changepoint: An R package for changepoint analysis*. Retrieved from https://CRAN.R-project.org/package=changepoint

Lazarus, R. S. (1966). *Psychological stress and the coping process*. New York: McGraw-Hill.

Lin, Y.-R., Margolin, D., & Wen, X. (2017). Tracking and analyzing individual distress following terrorist attacks using social media streams. *Risk Analysis*, *37*, 1580–1605.

http://dx.doi.org/10.1111/risa.12829

Lucas, R. E., Clark, A. E., Georgellis, Y., & Diener, E. (2003). Reexamining adaptation and the set point model of happiness: Reactions to changes in marital status. *Journal of Personality and Social Psychology*, *84*, 527–539. http://dx.doi.org/10.1037/0022-3514.84.3.527

Mitchell, M. N. (2012). *Interpreting and visualizing regression models using Stata*. College Station, TX: Stata Press.

Murthy, D., & Longwell, S. A. (2013). Twitter and disasters: The uses of Twitter during the 2010 Pakistan floods. *Information, Communication & Society*, *16*, 837–855. http://dx.doi.org/10.1080/1369118X.2012.696123

Nolen-Hoeksema, S., & Morrow, J. (1991). A prospective study of depression and posttraumatic stress symptoms after a natural disaster: The 1989 Loma Prieta earthquake. *Journal of Personality and Social Psychology*, *61*, 115–121. http://dx.doi.org/10.1037/0022-3514.61.1.115

Norris, F. H., Friedman, M. J., Watson, P. J., Byrne, C. M., Diaz, E., & Kaniasty, K. (2002). 60,000 disaster victims speak: Part I. An empirical review of the empirical literature, 1981–2001. *Psychiatry: Interpersonal and Biological Processes*, *65*, 207–239. http://dx.doi.org/10.1521/psyc.65.3.207.20173

Pactol, C. C. (2018). False Missile Alert Causes Panic. Retrieved from https://themolokaidispatch.com/false-missile-alert-causes-panic/

Pennebaker, J. W., Booth, R. J., Boyd, R. L., & Frances, M. E. (2015). *Linguistic Inquiry and Word Count: LIWC2015*. Austin, TX: Pennebaker Conglomerates (www.LIWC.net).

Pennebaker, J. W., Mehl, M. R., & Niederhoffer, K. G. (2003). Psychological aspects of natural language use: Our words, our selves. *Annual Review of Psychology*, *54*, 547–577.

http://dx.doi.org/10.1146/annurev.psych.54.101601.145041

Pew Research Center (2018). *Social media use in 2018*. Retrieved from http://www.pewinternet.org/2018/03/01/social-media-use-in-2018/

R Core Team. (2018). *R: A Language and Environment for Statistical Computing*. Vienna, Austria: R Foundation for Statistical Computing. Retrieved from http://www.R-project.org/

Ripberger, J. T., Silva, C. L., Jenkins-Smith, H. C., Carlson, D. E., James, M., & Herron, K. G. (2015). False alarms and missed events: The impact and origins of perceived inaccuracy in tornado warning systems. *Risk Analysis*, *35*, 44–56. http://dx.doi.org/10.1111/risa.12262

Rodríguez, H., Díaz, W., Santos, J. M., & Aguirre, B. E. (2007). Communicating risk and uncertainty: Science, technology, and disasters at the crossroads. In H. Rodríguez, E. L. Quarantelli, & R. R. Dynes (Eds.), *Handbook of disaster research* (pp. 476–488). New York, NY: Springer. https://doi.org/10.1007/978-0-387-32353-4_29

Rosen, N. O., Knäuper, B., & Sammut, J. (2007). Do individual differences in intolerance of uncertainty affect health monitoring? *Psychology and Health*, *22*, 413–430. http://dx.doi.org/10.1080/14768320600941038

Seery, M. D., Holman, E. A., & Silver, R. C. (2010). Whatever does not kill us: Cumulative lifetime adversity, vulnerability, and resilience. *Journal of Personality and Social Psychology*, *99*, 1025–1041. http://dx.doi.org/10.1037/a0021344

Silge, J., & Robinson, D. (2016). tidytext: Text mining and analysis using tidy data principles in R. *JOSS*, *1*. https://doi.org/10.21105/joss.00037

Silva, D. (2018). Hawaii false alarm: Ensuing chaos is teachable moment, experts say. Retrieved from https://www.nbcnews.com/news/us-news/hawaii-false-alarm-ensuing-chaos-teachable-moment-experts-say-n838836

Silver, R. C. (2004). Conducting research after the 9/11 terrorist attacks: Challenges and results. *Families, Systems, & Health*, *22*, 47–51. http://dx.doi.org/10.1037/1091-7527.22.1.47

Silver, R. C., & Fischhoff, B. (2011). What should we expect after the next attack? *American Psychologist*, *66*, 567–572. http://dx.doi.org/10.1037/a0024893

Silver, R. C., Holman, E. A., Andersen, J. P., Poulin, M., McIntosh, D. N., & Gil-Rivas, V. (2013). Mental- and physical-health effects of acute exposure to media images of the September 11, 2001, attacks and the Iraq War. *Psychological Science*, *24*, 1623–1634. http://dx.doi.org/10.1177/0956797612460406

Simmons, K. M., & Sutter, D. (2009). False alarms, tornado warnings, and tornado casualties. *Weather, Climate, and Society*, *1*, 38–53. http://dx.doi.org/10.1175/2009WCAS1005.1

Sweeny, K., & Vohs, K. D. (2012). On near misses and completed tasks: The nature of relief. *Psychological science*, *23*, 464-468.

Taha, S. A., Matheson, K., & Anisman, H. (2014). H1N1 was not all that scary: Uncertainty and stressor appraisals predict anxiety related to a coming viral threat. *Stress and Health*, *30*, 149–157. http://dx.doi.org/10.1002/smi.2505

Tausczik, Y. R., & Pennebaker, J. W. (2010). The psychological meaning of words: LIWC and computerized text analysis methods. *Journal of Language and Social Psychology*, *29*, 24–54. http://dx.doi.org/10.1177/0261927X09351676

Wamsley, L. (2018). *Worker who sent Hawaii false alert thought missile attack was imminent*. Retrieved from https://www.npr.org/sections/thetwo-way/2018/01/30/581853255/hawaii-missile-drill-stated-this-is-not-a-drill-resulting-in-false-alert

Wickham, H. (2009). *ggplot2: Elegant Graphics for Data Analysis*. Springer-Verlag New York. Retrieved from http://ggplot2.org

Acknowledgements

This Chapter has been submitted for publication in a peer-reviewed journal. I thank my coauthor for her invaluable contributions to this manuscript.

Jones, N. M., & Silver, R. C. (submitted for publication). This is not a drill: Anxiety on Twitter following the 2018 Hawaii false missile alert.

Chapter 5: Epilogue

Epilogue

Twitter data are an invaluable means with which social scientists can quickly and effectively study the impact of collective traumas. Not only can these data shed light on the initial and early psychological reactions of a community, but they can also be used to separate users in psychologically meaningful ways to examine differential impact of collective traumas and trajectories of emotion expression over time. As a large-scale observational data source, Twitter can supplement traditional data collection methods in ways that allow a researcher to a) triangulate on the impact of a collective trauma and b) explore theoretical assumptions regarding how individuals respond to collective trauma. This dissertation contributes to the larger bodies of social psychological and trauma studies by explicating how using big data generated on social media can aid in understanding how people respond to collective traumas, how rumors can exacerbate distress when they proliferate on social media without corrections, and how false alarms of potentially deadly situations can have a lingering psychological effect on thousands of people, even after the threat is dispelled.

Chapter 2 examined the efficacy of a novel methodological approach for downloading Twitter data by sourcing users who follow Twitter accounts in the community under study. This method circumvents challenges associated with relying on geo-coded Twitter data and those challenges associated with conducting traditional studies of collective trauma (Silver, 2004). The study’s aim was to understand the community impact of the 2015 terrorist attack in San Bernardino, California, compared to a control community matched on important demographic variables. Results highlighted how these data can be examined in several longitudinal frames to understand the time course of negative emotion expression before, during, and after a collective trauma. Specifically, negative emotion on Twitter increased 6.2% in San Bernardino on the day of the attack, relative to the comparison community in which no increase was observed. Additionally, this study demonstrated how computational approaches to content analysis of tweets generated in the immediate aftermath of a collective trauma reflect a process of information dissemination as authorities and the news media pieced together what happened. Overall, the method used for sourcing Twitter data in the context of a terrorist attack was sufficient for studying its impact on the affected community.

Chapter 3 leveraged the strengths of both traditionally-collected survey data and Twitter data for evaluating a) the relationship between acute stress and rumor exposure among several thousand students affected by a shooting and protracted lockdown at a large U.S. university, and b) rumor generation, transmission, and negative emotion expression on Twitter among followers of the university's social media accounts. The results from this study revealed that students trapped in the lockdown relied on many social media sites for information about the severity of the situation. Students who used many social media sites for critical updates and trusted what they saw reported more acute stress than individuals who trusted social media less. Students who used Twitter for critical updates also reported more exposure to conflicting information during the lockdown. Consistent with the rumor literature (Allport & Postman, 1946; Festinger et al., 1948; Rosnow, 1991), an analysis of Twitter data from users who followed the university's emergency and primary Twitter accounts revealed that rumors were generated and retweeted several thousand times during the lockdown period when no critical updates were transmitted from the university authorities. Negative emotion expression on Twitter also increased during this time frame. Using a mixed methods approach, this study triangulated on this collective trauma to understand information seeking and psychological responses. This work resulted in several policy recommendations for increasing communication and easing psychological discomfort with uncertainty during a crisis.

Chapter 4 tested theoretical assumptions about how individuals respond to a false alarm of a potentially deadly, inescapable, collective trauma: The 2018 Hawaii ballistic missile alert. Using Twitter data generated by likely residents across the State of Hawaii (sourced using the technique described in Chapter 2), the results indicated that anxiety about this event increased on the day of the missile alert and lingered for at least two days among the entire sample. Consistent with the incubation of threat hypothesis (Breznitz, 1968, 1971, 1984), a fine-grained analysis of tweets generated six hours before and after the alert period showed that anxiety increased 3.4% every 15 minutes during the 38-minute alert period (until the all clear was officially transmitted).

Users were also grouped by their pre-alert tendency to express anxiety on Twitter (low, medium, and high pre-alert anxiety). Doing so enabled an analysis of differential anxiety responses at the onset of the alert, stabilization rates after, and whether groups settled to new baseline levels of anxiety expression after the alert compared their pre-event baselines. The results showed that relative to other groups, the low pre-alert anxiety group exhibited the largest increase in anxiety at the onset of the alert (9.5%), and this group's expression of anxiety remained elevated the longest (41 hours) before returning to a new baseline 2.5% higher than their pre-alert level. Surprisingly, this analysis also revealed that the high pre-alert anxiety group exhibited a pattern contrary to what was hypothesized. This group's anxiety decreased 8.8% at the onset of the alert and immediately stabilized at a new level 10.5% lower than their baseline for at least seven days after the alert.

This is likely one of the first studies to take such an approach to examine whether the emergent patterns corroborate theoretical assumptions of how life before a collective trauma influences reactions to, and recovery from, a collective trauma. Prevailing perspectives of reactions to collective trauma (Asarnow et al., 1999; Nolan-Hoeksema & Morrow, 1991) suggest that users high in anxiety, measured before the ballistic missile alert, may be at greater risk of increased anxiety resulting from the missile alert. However, this was not the case; the observed pattern was in the opposite direction of what was expected. High pre-alert anxiety users expressed less anxiety at the onset of the alert and thereafter, for up to seven days, compared to their pre-alert baseline.

Taken together, these studies comprise a novel approach to studying the impact of collective traumas and explore whether predictions derived from classic social psychological theories (e.g., rumor theory, psychology of false alarms) play out as expected when analyzing big data. As demonstrated in this dissertation, Twitter data make it possible to examine how rumors and false alarms impact people in large-scale, dangerous situations and allow for rigorous quasi-experimental designs to be employed for making strong claims about the impact of these events. The studies in this dissertation capitalized on the strengths of Twitter data for answering important theoretical questions by analyzing data at several time scales and manipulating the data so as to parse users into psychologically meaningful groups to understand differential impact of collective traumas. This flexibility in design and manipulation constitutes the most exciting and promising aspects of this approach.

Limitations and other considerations

Throughout this dissertation, the advantages and drawbacks of using tweets sourced from followers of local Twitter accounts have been discussed. For example, these data primarily come from young, urban residents (Pew Research, 2018), and the patterns found in each study may not generalize to the population as whole (although see a rebuttal to this limitation in Chapter 3). Furthermore, the method employed in this dissertation cannot guarantee that all individuals in these studies were residents of a geographical area of interest. Despite this potential for error, the method reliably captured patterns one would expect to observe after collective traumas. Indeed, these patterns would have likely been stronger had the error been absent. These limitations do highlight the importance of incorporating other data sources with which to triangulate on the impact these events have on individuals and the communities they comprise. Researchers should strive to obtain survey data or other community-level data (e.g., monthly emergency room encounters, county-level health metrics), when attempting to use Twitter data to understand the impact of a collective trauma.

Another important consideration regarding the use of big social media data is that access to them may be here today and gone tomorrow. Increasing concerns around user rights and privacy, especially resulting from the Facebook/Cambridge Analytica scandal (Granville, 2018), have tipped the hand in favor of national governmental bodies keen on regulating the flow of information from social media companies to advertisers and other third parties (including social scientists). The European Union recently enacted the General Data Protection Regulation (GDPR), which sets out binding standards that give users the right to see the data companies have collected on them and provide pathways to have those data removed from company servers (https://gdpr-info.eu/). This regulatory standard in Europe has also taken hold in the United States as Facebook has campaigned to regain the trust of its users. Other social media platforms, like Twitter, have followed suit. The result is that some inconsequential data points (e.g., user profile image links) downloaded as part of this dissertation in early 2018 are currently no longer available. Moving forward, increased regulations, or changing company policies as companies cope with compliance, could preclude budding social scientists from using social media data, especially if they lack the funding to purchase access to these rich data sources.

Future Directions

Although Twitter is just one platform in a sea of social media applications people use throughout the day, it is the leading platform for posting feelings, ideas, criticism, and sharing information in the public space. Other platforms (e.g., Facebook and Snapchat) are geared toward the exchange of information between smaller networks of individuals who know each other, and the norms associated with using these platforms revolve around maintaining privacy through user control over who has permission to access personal content. Thus, public data at the individual level is limited and rapidly conducting large-scale studies, like those presented in this dissertation, is currently not possible using these other platforms.

At this time, Twitter is the only social media company that provides social scientists free access to its data and this access enables researchers to examine much more than how users respond to collective traumas. Evidence from this dissertation suggests that information from the news media spills into the public sphere, and an analysis of what users tweet in the aftermath of a terrorist attack reflects the release of event-related information by news organizations. Because most organizational Twitter accounts are public, there is an opportunity to scrape data from emergency management organizations, news media, and governmental agencies to understand the time course of information flow into the public sphere when disaster strikes. In the past, studies of the impact of the news media on human behavior in certain disaster contexts required enough time for newspaper articles to be posted to university library databases. However, with instant access to Twitter streams available now, researchers can quickly download and evaluate information transmitted on Twitter accounts owned by news organizations and link that information with human behaviors that unfold, in real time, after a collective trauma. Future work could examine, for example, how information released to the public before a hurricane makes landfall is related to evacuation behavior.

Access to Twitter data also offers many opportunities for researchers to continue evaluating theoretical assumptions about social psychological phenomena (e.g., rumor transmission). Indeed, researchers interested in the spread of misinformation on social media have conducted several studies to explore how rumors (Arif et al., 2017; Vosoughi, Roy & Aral, 2018) and conspiracy theories (Starbird, 2017) spread on Twitter after political upheavals and collective traumas. In this era of "fake news" and statements like "truth isn't truth" being elevated into the collective conscious of the nation, this work is important for understanding what strategies are necessary for combating rumors and conspiracy theories that are easily peddled in digital spaces and sow distrust in local and national institutions as they transmit critical information during disaster periods.

One of the most intriguing findings from this dissertation reveals that users with high anxiety expression before the 2018 false ballistic missile alert in Hawaii, exhibited less anxiety expression in the seven days following the alert period. Although several potential explanations of this pattern were offered, it clearly presents a challenging new area in which to explore how social media users' emotions and psychological states before a collective trauma play a role in their responses after. Are high anxiety Twitter users simply categorically different than participants used in studies that demonstrate the role of pre-event anxiety on post-event responses to collective traumas (Asarnow et al., 1999 and Nolan-Hoeksema & Morrow, 1991)? The answer is unknown, but future work should aim to group users in meaningful ways, using pre-event data, to capture the full range of psychological responses that unfold on Twitter in the immediate aftermath of a collective trauma, and thereafter. This effort will likely reveal other unexpected patterns and potentially inform new theoretical perspectives on how human interaction with technology can inform responses to a collective trauma.

References

Allport, G. W., & Postman, L. (1946). An analysis of rumor. *Public Opinion Quarterly*, *10*, 501–517.

Arif, A., Robinson, J. J., Stanek, S. A., Fichet, E. S., Townsend, P., Worku, Z., & Starbird, K. (2017). A closer look at the self-correcting crowd: Examining corrections in online rumors. In *Proceedings of the 2017 ACM Conference on Computer Supported Cooperative Work and Social Computing* (pp. 155–168). New York, NY: ACM. chttps://dx.doi.org/10.1145/2998181.2998294

Asarnow, J., Glynn, S., Pynoos, R. S., Nahum, J., Guthrie, D., Cantwell, D. P., & Franklin, B. (1999). When the earth stops shaking: Earthquake sequelae among children diagnosed for pre-earthquake psychopathology. *Journal of the American Academy of Child & Adolescent Psychiatry*, *38*, 1016–1023. https://dx.doi.org/10.1097/00004583-199908000-00018

Breznitz, S. (1968). "Incubation of threat" in a situation of conflicting expectations. *Psychological Reports*, *22*, 755–756.

Breznitz, S. (1971). A study of worrying. *British Journal of Social and Clinical Psychology*, 10, 271-279. http://dx.doi.org/10.1111/j.2044-8260.1971.tb00747.x

Breznitz, S. (1984). *Cry wolf: The psychology of false alarms*. Hillsdale, NJ: Lawrence Erlbaum Associates, Inc.

Festinger, L., Cartwright, D., Barber, K., Fleischl, J., Gottsdanker, J., Keysen, A., & Leavitt, G. (1948). A study of a rumor: Its origin and spread. *Human Relations*, *1*, 464–486.

Granville, K. (2018, March 19). Facebook and Cambridge Analytica: What you need to know as fallout widens. *The New York Times*, Retrieved from https://www.nytimes.com/2018/03/19/technology/facebook-cambridge-analytica-explained.html

Nolen-Hoeksema, S., & Morrow, J. (1991). A prospective study of depression and posttraumatic stress symptoms after a natural disaster: The 1989 Loma Prieta earthquake. *Journal of Personality and Social Psychology*, *61*, 115–121. https://dx.doi.org/10.1037/0022-3514.61.1.115

Pew Research Center (2018). *Social media use in 2018*. Retrieved from http://www.pewinternet.org/2018/03/01/social-media-use-in-2018/

Rosnow, R. L. (1991). Inside rumor: A personal journey. *American Psychologist*, *46*, 484–496. https://dx.doi.org/10.1037/0003-066X.46.5.484

Silver, R. C. (2004). Conducting research after the 9/11 terrorist attacks: Challenges and results. *Families, Systems, & Health*, *22*, 47–51. https://dx.doi.org/10.1037/1091-7527.22.1.47

Starbird, K. (2017). Examining the alternative media ecosystem through the production of alternative narratives of mass shooting events on Twitter. Retrieved from http://faculty.washington.edu/kstarbi/Alt_Narratives_ICWSM17-CameraReady.pdf

Vosoughi, S., Roy, D., & Aral, S. (2018). The spread of true and false news online. *Science*, *359*, 1146–1151. https://dx.doi.org/10.1126/science.aap9559