

Gradients of Fear and Anger in the Social Media Response to Terrorism

Matt Baucum
University of Southern California
baucum@usc.edu

Richard John
University of Southern California
richardj@usc.edu

Abstract

Research suggests that public fear and anger in wake of a terror attack can each uniquely contribute to policy attitudes and risk-avoidance behaviors. Given the importance of these negative-valenced emotions, there is value in studying how terror events can incite fear and anger at various times and locations relative to an attack. We analyze 36,259 Twitter posts authored in response to the 2016 Orlando nightclub shooting and examined how fear- and anger-related language varied with time and distance from the attack. Fear-related words sharply decreased over time, though the trend was strongest at locations near the attack, while anger-related words slightly decreased over time and increased with distance from Orlando. Comparing these results to users' pre-attack emotional language suggested that distant users remained both angry and fearful after the shooting, while users close to the attack remained angry but quickly reduced expressions of fear to pre-attack levels.

1. Introduction

Understanding public reactions to highly salient terror attacks is crucial to appraising the overall risk that terrorism poses to society. Terror attacks have the potential to cause psychological and economic damage that can far outlast their immediate effects [1, 2, 3, 4], and it is important for researchers and public officials to anticipate the contours of such effects as best as possible.

Intuitively, any highly publicized act of terror should cause heightened levels of negative affect among the public. Yet one of the clearest and more nuanced psychological findings regarding terrorism's emotional impact is that *fear* and *anger* seem to serve different psychological functions in the aftermath of a disaster event. One study [5] demonstrated that an experimental induction of fear increased participants' terrorism risk perceptions, whereas inducing anger

decreased them. In the wake of the September 11th terror attacks, individuals' self-reported fear predicted preferences for "defensive" anti-terror policies (e.g., deporting suspected terrorists), while self-reported anger predicted support for "offensive" policies (e.g., aggressive military action in the Middle East) [6]. Fear and anger also have distinct behavioral and psychological correlates outside the domain of terrorism; perceptions of fearful and angry faces differentially predict approach and avoidance behaviors [7], and there is mounting evidence that the experience of fear and anger are respectively related to avoidance and approach motivations [8]. Consequently, the public's experiences of both fear and anger in the wake of a terror attack likely play important and distinct roles in determining the attack's overall effect on national discussions of policy.

Of course, all expressions of fear and anger in the wake of terrorism are not created equal, and given their relevance to risk judgments and policy preferences, it is worth understanding how such emotional reactions are situated within time and place. Knowing how public fear and anger increase or decrease over time after an attack can inform predictions on how public discourse surrounding the attack will take shape, especially if one emotion proves to be more temporally stable than the other. Similarly, understanding how expressions of fear and anger depend on one's distance from the attack can shed light on regional differences in risk perceptions. Such gradients of fear and anger are especially important to understand at the national level given that terror attacks are becoming increasingly local in nature, with a greater emphasis on small, ground-based, "soft-target" attacks [9]. Such small-scale attacks may not produce the kinds of far-reaching, long-lasting emotional responses that characterized the September 11th attacks, which fueled so much of the seminal research on public risk perception in the U.S. Thus, the degree to which one's expressions of fear and anger after an attack depends on their temporal and geographic proximity is an important

empirical question, and one that this study seeks to address.

1.1 Emotion, Geography, and Time: Existing Research

Temporal and geographic proximity to terror attacks have already been studied as meaningful predictors of public terror reactions. After the September 11th, 2001 attacks, posttraumatic stress, driving fatalities (assumed to reflect increased road traffic due to an aversion to flying), and estimates of future terrorism risk were greatest among individuals living close to New York City [10, 11, 12]. Furthermore, recent work [13] has found that social media expressions of fear and anger in the Paris metropolitan area sharply increased after a series of shootings in November of 2015, then sharply decreased in the days following. Yet it is still unclear how post-attack expressions of fear and anger vary over time *and* geographic proximity in a country-wide sample; [13] limited their analyses to the Paris metropolitan area (a necessary characteristic of their method), while other studies [10, 11, 12] focused mainly on fear- or stress-based reactions. It is not yet known whether public fear and anger “behave” similarly or differently in the aftermath of a crisis when spatial distance from the event is considered, a finding that could help better clarify the roles that these emotions play in public disaster response.

We expect that, in the aftermath of a disaster event (specifically, a terrorist attack for the purposes of this study), public expressions of both fear and anger will be strongest in the event’s immediate aftermath and decrease over time. This expectation aligns with previous research on emotional expressions in the aftermath of terror attacks [13], and intuitively aligns with the notion that one’s emotional reaction to an event becomes less severe as the event fades from immediate memory.

We also expect that expressions of fear will be greatest at locations close to a disaster event of interest, as suggested by research on regional variations in fear following the 9/11 attacks. The relationship between anger and geographic distance, however, is more theoretically complex. Just as fear reactions are strongest among individuals who reside near where a disaster event occurred [10, 11, 12, 14], the same may be true of anger; the psychological closeness of the event may simply amplify its emotional intensity across all negative emotions, including anger. Furthermore, anger is a moral emotion that is often brought on by perceptions of suffering [15], which are likely strongest at locations close to a disaster event [14]. However, fear and

anger are served by different cognitive appraisals [16], with fear arising from appraisals of uncertainty and uncontrollability, and anger arising from appraisals of certainty and controllability. It may be that individuals closer to a disaster event experience heightened levels of fear, but that the cognitive uncertainty producing such fear inhibits the expression of anger to the same degree as individuals who are distant from the attack. Thus, the unclear relationship between geographic distance and anger is one of the primary contributions of this work.

Lastly, we investigate whether the hypothesized decreases in fear and anger over time depend on an individual’s distance from the attack. To inform this research question, we again draw on [13] which showed that decreases in fear and anger over time were well-modeled by a Weibull survival model. This finding suggests that fear and anger did not decrease linearly in the sample, but that the magnitude of a given day’s decrease was related to the magnitude of emotional expression on the previous day. Based on this result, we expect that locations with the highest levels of expressed fear and anger will exhibit sharper decreases than locations with lower levels of emotional expression. Thus, the effect of time on fear and anger will be strongest at the geographic distance corresponding to the highest level of their expression.

1.2 Present Study: Social Media Responses to the Orlando Nightclub Shooting

To investigate gradients of fear and anger responses to terrorism at a national level, we focus our analysis on the 2016 shooting at the Pulse Nightclub in Orlando, Florida. The 2016 Orlando nightclub shooting was the deadliest mass shooting in the United States at the time of its occurrence, and is prototypical of the kinds of soft-target terror attacks that have come to dominate the terrorism landscape in recent years [9, 17], making it a useful case study on post-terror attack discourse across the United States. Note that we classify the Orlando shooting as a *terror attack* for the purposes of this study given that the shooter had personally pledged allegiance to the Islamic State in Iraq and Syria before carrying out the assault.

To assess individual reactions to the attack in a naturalistic setting, we focus on the social media response following the shooting. In the aftermath of highly-publicized crisis events, social media can serve as a platform for collective information sharing [18, 19], partly because the desire to obtain current information is a strong motivation for social media use [19]. Social media has even been shown to serve

different functions for those at varying distances from a mass emergency, with users immediately affected by an event more likely to share locally-relevant information and those further away more likely to engage in generic commentary [20].

Of course, behavior on social media platforms is driven by a host of factors that can skew the quality of the information shared, such as users' reputations [19]. Still, social media has served as a useful data source in other investigations of public terror reactions [13, 14], and it carries the benefit of allowing for unobtrusive measurement of individuals' expressed thoughts and opinions. Thus, we argue that it still holds relevance for theoretically-motivated questions if interpreted with caution.

1.3 Variables of Interest and Hypotheses

Our outcome variables of interest are the use of fearful and angry language in Twitter posts discussing the 2016 Orlando nightclub shooting, with each user's geographic proximity to the shooting's location and the elapsed time between the attack and authorship of their Twitter post as the primary predictor variables. Regarding social media expressions of fear in the aftermath of the attack, we hypothesize the following:

- 1) Use of fear-related language will negatively correlate with the elapsed time between the attack and each Twitter post (i.e., decreases over time, as suggested by [13]).
- 2) Use of fear-related language will negatively correlate with users' distance from the shooting's location in Orlando, FL (as suggested by research on regional variations in reactions to 9/11).
- 3) Distance will moderate the effect of time on fear-related language, such that the (hypothesized) decrease in fear over time will be strongest at locations closer to the attack.

Regarding social media expressions of anger, we hypothesize the following:

- 1) Use of anger-related language will negatively correlate with the elapsed time between the attack and each Twitter post (i.e., decreases over time, as suggested by [13]).
- 2) Use of anger-related language will depend on geographic distance from Orlando, FL, though we do not hypothesize the direction of this effect (given the potential theoretical justifications for both directions).
- 3) Distance will moderate the effect of time on anger-related language, such that the (hypothesized) decrease in anger over time will

be strongest at whichever distance is related to higher initial levels of anger.

2. Method

2.1 Sample

We obtained a sample of Twitter posts made between June 11, 2016 and June 19, 2016 (the week following the Orlando nightclub shooting) that included one or more of the hashtags "#OrlandoShooting", "#Orlando", or "#pulseshooting" (an initial web search suggested that these were the most common hashtags used on social media to refer to the event). The initial dataset yielded over 4 million posts, from which we excluded all retweets (posts written by one user and re-posted by another) and posts that only contained hashtags or web address links.

Posts had to be in English (see automatic language detection function in R package "cld2"; [21]), authored by non-verified Twitter accounts (where *verified* refers to official accounts for organizations or celebrities), and posted by users in the United States for whom location data (at the city level) was available. Where there were multiple posts written by the same user, we include only their earliest post, and we eliminated all posts from before the onset of the shooting (defined as 2:06 a.m. Eastern Time, June 12, 2016 [22]). Further inspection of the data revealed some tweets posted by news sites (rather than individuals) that were not screened out with the original criteria, which were subsequently removed. Filtering the dataset by these criteria yielded a final sample of 36,259 posts. Note that a sample of this size would allow us to detect bivariate correlations of 0.019 with 95% power, and while our analyses do not primarily rely on null hypothesis significance testing or bivariate correlations, this serves as an example of the sufficiency of our sample size for our research purposes.

2.2 Procedure

2.2.1. Text preprocessing. We implemented multiple cleaning steps to convert each Twitter post into an analyzable text object for further analysis. Each post was stripped of all non-punctuation/non-alphanumeric characters (which removes special characters such as Emojis), as well as all links to other content (such as webpages or pictures). Because some users often use hashtags as parts of their post's syntax (e.g., "Our thoughts are with the #pulseshooting victims"), we chose to retain all hashtags that were directly followed by non-hashtag

words while removing all others. Thus, hashtags embedded in the middle of sentences are assumed to serve some grammatical function and are kept as part of the post's content, while those that appear at the end (where users often place multiple hashtags in a row) are removed. For example, the tweet "Our thoughts are with the #pulseshooting victims #Orlando #OrlandoShooting" would be shortened to "Our thoughts are with the pulseshooting victims." While these procedures cannot guarantee that each post will perfectly reflect the semantic content intended by the author, it reduces much of the noise introduced by social media conventions.

2.2.2. Time and distance from attack. For each post, temporal distance from the Orlando attack was measured as the number of days (including partial days) between 2:06 a.m. ET on June 12, 2016 (the time at which police were notified of the Orlando shooting; [17]) and the date/time at which the post was created.

Geographic proximity was measured as the number of miles between the latitude and longitude coordinates of the Pulse nightclub and the central latitude/longitude coordinates of each user's nearest ZIP code, which was retrieved through the Bing Maps Application Programming Interface (API). Because not all users choose to report their city of residence in their profile (or might simply include broader location information, such as state or country), location data was only retained for users in the United States where the API could identify a single location profile at the city level. We identified the closest U.S. ZIP code to each city center, and calculated the distance between this ZIP code and the Pulse Nightclub using the "Imap" R package [23] and ZIP code location data from the United States Census Bureau [24]. Note that distance was measured "as the crow flies," rather than based on driving distance, which was necessary given that some users resided in Hawaii.

2.2.3. Covariates. One of the challenges of measuring the effect of geographic proximity on emotional expression is its confounds with other potentially relevant variables. Thus, we also account for the following covariates in our analysis:

It is possible that individual reactions to terror attacks depend partly on one's residence in an urban or rural area (given that many high-profile terror attacks target dense, urban locations), and we thus control for population density (given that some of the least densely populated areas of the United States are in western states and Alaska, relatively distant from Orlando). As aforementioned, we employed the Bing

Maps API to match each user's self-described location with a U.S. city; those that could be correctly matched were cross-referenced with data from the 2010 Census [25] to produce city-level population density for each user.

Note that we conduct all analysis with the logarithm of population density; this transformation reduced substantial positive skewness, and reflects the theoretical assumption that differences in population density on the lower end (distinguishing urban from rural areas) likely matter more than similar-sized variations at the high end (distinguishing urban areas of different density). As an example, using raw population densities for Galena, AK (a small town with population 5,700), Dallas, TX, and New York, NY leads to a difference in "urban-ness" between Dallas and New York (both major cities) that is roughly 7.5 times the difference in "urban-ness" between Dallas and Galena (one of which is considered a large city, the other a small town). Meanwhile, employing log-scaled population density sets these differences as approximately equal, which seemed to better represent the distinction between urban and rural areas.

Many of the lowest-earning states in the U.S. are in or near the American South, according to recent census data (e.g., Mississippi, West Virginia, Arkansas, Alabama, Kentucky), whereas some of the highest earning states are in the northeast (New Jersey, Connecticut, Massachusetts, New Hampshire) or even outside the continental U.S. (Alaska, Hawaii). To ensure that any effect of geographic distance was unrelated to any regional disparities in wealth, we estimated users' income based on the Twitter accounts they followed. This method comes from [26], which identified the Twitter accounts that best predicted the annual income of the users that followed them. We apply the regression model from their study to each of our user's friends list to estimate which of four income categories (\$0-\$50,000; \$50,000-\$100,000; \$100,000-\$150,000; \$150,000+) the user most likely belongs to, and defined their estimated income as the lower bound of their most likely income bracket. See [26] for a list of the Twitter accounts used to estimate income, along with their respective regression weights. Users' scores for each income bracket were calculated as the sum of the weights for the accounts they followed that corresponded to that income bracket, and they were assigned to the income bracket for which they had the highest score.

Lastly, we sought to ensure that any effects of geographic distance were not confounded with any regional differences in political orientation. While political orientation is not directly available from

Twitter profiles, we employ a method introduced by [27] to estimate the political sentiments of each user. In this procedure, [27] estimated the average political ideology of the users who followed various Twitter news accounts (e.g., @foxnews, @latimes), based on the Congress members that their audiences also followed. We use these ideology estimates for 20 news accounts to estimate each participants' political orientation as the average ideology score of the news accounts they follow; see [27] for a list of the news accounts used for this calculation.

2.2.4. Anger- and fear-related language. Anger and fear expressions in each post were defined as the number of words from the *anxiety* and *anger* dictionaries from Linguistic Inquiry and Word Count [28], a widely-used and well-validated collections of terms representing psychological constructs that has previously been used in social media analyses of public terrorism response [13]. The anxiety dictionary contains terms such as *scared*, *vulnerable*, *stunned*, and *uneasy*, while the anger dictionary contains terms such as *angry*, *evil*, *mad*, and *hate*. We specifically removed the term *terror** from the anxiety dictionary (as users who referenced *terrorism* may not necessarily be expressing fear), and removed the terms *kill* and *victim* from the anger dictionary, as users employing these terms may have simply been describing the attack.

2.3 Analysis

We rely on Bayesian Poisson regression to test for relationships between each of our psychological distance predictors and our word-count variables of interest (fear expressions, anger expressions, and concrete term use), and communicate all regression effects using 95% Highest Posterior Density (HPD) intervals (i.e., the shortest interval containing 95% of the parameter's posterior density). Poisson regression is a method for analyzing count data when the variance is approximately equal to the mean, which was true for word counts of fear ($\bar{x}=0.068$, $s^2=0.072$) and anger ($\bar{x}=0.388$, $s^2=0.401$), and both word count variables were well-approximated by Poisson distributions. Note that, because we use raw word counts rather than term frequencies (i.e., raw word counts divided by the text's length), we include Tweet word count as a covariate in all analyses, such that all reported effects are independent of the length of each Twitter post.

Given that estimating users' income and political orientation required that they follow certain accounts (and also required that their list of followed accounts was made publicly available), only $n=16,492$ users

had non-missing estimates for political orientation and income. Thus, in all models, we first test for the effect of temporal and geographic distance on each user's fear and anger term counts, then confirm that such effects remain reliable when including population density (in the full sample) and income and political orientation (in the $n=16,492$ subsample of users with non-missing data). Note that fear and anger term counts did not appreciably differ between the subsamples of users with and without missing political orientation and income estimates ($|ds|<0.05$).

3. Results

3.1 Descriptive statistics

Figure 1 presents kernel density plots for each predictor, and shows how most variables exhibited at least some degree of skew, with some bimodality present in the population density variable (largely due to the high calculated population density of New York City, which comprises the righthand mode of the distribution). As aforementioned, we conduct all analyses using the *logarithm* of population density rather than its raw value (also presented in Figure 1).

Time and geographic proximity were virtually uncorrelated with each other (Spearman's $\rho=-0.019$), as was geographic proximity with income estimates ($\rho=-0.059$) and political orientation ($\rho=-0.050$). Distance and population density were correlated at $\rho=0.262$, suggesting that users who lived further from Orlando tended to occupy more densely populated areas. Lastly, counts of fear- and anger-related terms were relatively uncorrelated ($\rho=0.025$).

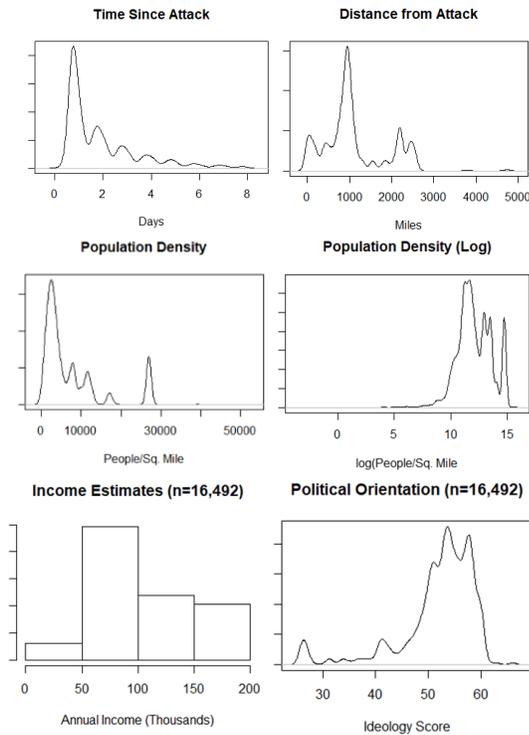


Figure 1. Density plots for predictors and covariates

3.2 Predicting Expressions of Fear

Expressions of fear on Twitter decreased as a function of time since the attack ($e^b=0.816$, 95% HPD=0.770, 0.862); interpreting the exponentiated regression coefficient suggests that each passing day corresponded to an 18.4% decrease in the prevalence of fear-related terms. There was no reliable main effect of geographic distance ($e^b=0.969$, 95% HPD=0.892, 1.053), though time and distance did interact in predicting fear-related language ($e^b=1.061$, 95% HPD=1.021, 1.102), with the effect of time on fear-related term use decreasing as distance from the attack increased. Each passing day corresponded to an 18.4% decrease in term use at 0 miles from the attack, a 13.4% decrease at 1000 miles from the attack, an 8.2% decrease at 200 miles from the attack, and a 2.6% decrease at 3000 miles from the attack.

Note that the magnitudes of these effects were unchanged when including population density in the model ($e^b_{time}=0.815$; $e^b_{dist}=0.962$; $e^b_{t*d}=1.061$). Effect sizes changed slightly when including income and political orientation (for the $n=16,492$ subset with non-missing values; $e^b_{time}=0.846$; $e^b_{dist}=0.979$; $e^b_{t*d}=1.050$), though were comparable to the effects

estimated in the $n=16,492$ subset without covariates included ($e^b_{time}=0.842$; $e^b_{dist}=0.977$; $e^b_{t*d}=1.051$), suggesting that the inclusion of income, political orientation, and population density did not attenuate the effects of time and proximity on expressions of fear. While not relevant to our hypotheses, note that political orientation (but not income) was reliably predictive of fear-related language ($e^b=1.010$, 95% HPD=1.001, 1.018), with a one SD change towards more liberal political orientation predicting a 7.5% increase in fear-related term use.

Figure 2 shows the gradient of predicted fear term use at different values of time and geographic distance. Note that the effect of time is strongly negative at close distances, but attenuates towards zero at more distant locations. Interestingly, the effect of distance is almost nonexistent at the time of the attack's occurrence, but becomes increasingly positive (with greater geographic distance predicting more expressions of fear) as time passes.

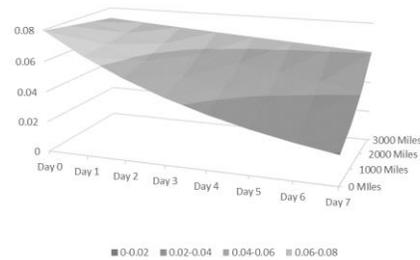


Figure 2. Gradient for predicted fear terms

3.3 Predicting Expressions of Anger

Anger-related terms also reliably decreased over time ($e^b=0.961$, HPD=0.942, 0.979), though to a lesser degree than fear-related language, with the average number of anger terms decreasing by 3.9% with each passing day. Unlike with fear, the main effect of distance was reliable ($e^b=1.061$, HPD=1.025, 1.095) and suggested that the average number of anger-related words increased by 6.1% for each 1000 miles of distance from Orlando. However, the use of anger-related language did not depend on an interaction between time and geographic distance ($e^b=1.005$, HPD=0.991, 1.019). These effects did not appreciably change when including population density ($e^b_{time}=0.961$; $e^b_{dist}=1.059$; $e^b_{t*d}=1.004$) or income and population density ($n=16,492$ subsample; $e^b_{time}=0.965$; $e^b_{dist}=1.045$; $e^b_{t*d}=1.005$). Note that estimates of political orientation (but not income) were slightly predictive of anger-related language

($e^b=0.992$, 95% HPD= $0.989, 0.996$), with a one SD change towards more liberal political orientation predicting a 5.4% decrease in anger-related term use.

Figure 3 shows the gradient of predicted anger-related term use at various times and distances from the attack. Note that the gradient is generally flatter than that of fear, largely reflecting the smaller role of time in predicting anger.

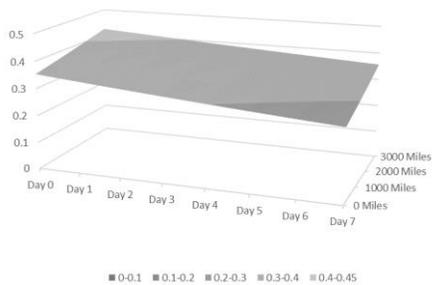


Figure 3. Gradient for predicted anger terms

3.4 Base Rates of Fear and Anger

To better contextualize the degrees of fear and anger expressed in each user’s post, we collected up to 200 of each user’s most recent posts made before the Orlando shooting, data which was publicly available for $n=21,811$ users, and identified the average number of fear- and anger-related terms that appeared in each post prior to the Orlando attack (note that this subsample did not severely differ from the other users in terms of fear or anger expressions in their posts; $|ds|<0.05$). On average, we collected 163 posts per user, which contained an average of 0.027 fear-related words per post (compared to 0.068 across the Orlando-focused tweets) and 0.060 anger-related words per post (compared to 0.388 across the Orlando-focused tweets).

Taking these values as base rate estimates of fear and anger term usage suggests that expressions of fear and anger were understandably more prevalent in our sample of Tweets than would be expected during “normal” Twitter activity by the same group of users, with the pre/post-attack change in anger language (a 547% increase) more pronounced than for fear (a 153% increase). Furthermore, examining these base rates in relation to the prediction gradients in Figures 2 and 3 suggests that anger-related language remained heightened through the week for users at all locations. Fear-related language remained heightened only for users distant from Orlando; those who lived near Orlando were predicted to return to pre-attack levels of fear language by roughly the fifth day after the attack.

4. Discussion

The data partially supported our hypotheses regarding the effects of time and geographic distance on social media sentiment following the Orlando nightclub shooting. As hypothesized, both anger-related and fear-related language in Tweets discussing the attack decreased over time, though expressions of fear decayed much more quickly than expressions of anger. As hypothesized, time and distance interacted in the model predicting fear (but not anger), suggesting that the time decay of fear-related language was most pronounced at locations closest to the Orlando shooting; however, this interaction seemed to increase, rather than decrease, the magnitude of the distance-fear association over time, given distance’s null effect in the attack’s immediate aftermath. Regional differences in anger were not as pronounced, but our model suggested that anger-related language was most common at distant, rather than proximal, locations (though distance and time did not interact in their effects on anger expressions, as hypothesized). Note that all estimated effects controlled for any potential influence of population density, estimated political orientation, and estimated user income.

Taken together, these results confirm past findings on the decline in negative sentiment in the days following terror attacks [13], while also suggesting potentially important differences in the public’s experience of the two emotions. Future studies should attempt to replicate our finding that linguistic markers of anger were more temporally stable than fear (and were less likely to return to pre-attack levels), as any real difference in the public’s tendency to “hold on” to one emotion over the other could have important implications for public discourse about terrorism. In general, spatial variation in both emotions suggests that users far from the attack remained relatively fearful *and* angry about the attack in the week following it, whereas users closer to the attack seemed to remain angry while reducing their expressions of fear to pre-attack levels relatively quickly. One possible explanation is that anger was related to users’ frustrations about societal or policy issues (e.g., recurrence of mass shootings, feelings towards terrorism, gun control laws, etc.) and was thus relatively stable across time and space. Yet as time passed, users close to the attack may have transitioned from expressing fear to expressing other emotions such as sympathy or solidarity (which have been shown to increase with geographic proximity to an event [13]), while distant users may have felt these

emotions to a lesser degree and remained generally fearful about the prospect of a future attack.

Perhaps one of the most applicable conclusions to be drawn here is the value of studying public anger following a highly publicized act of terror. Many studies on public terror reactions focus primarily on fear [29, 30, 31] or fear-related constructs, such as perceptions of risk. Yet while a handful of studies have acknowledged the diverging effects of fear and anger on relevant constructs such as policy attitudes, anger is still far from the focus of attention as a dependent measure in terrorism research. Our results suggest that anger, compared to fear, 1) increased more sharply in response to the Orlando shooting, 2) remained at elevated levels for longer, and 3) exhibited fewer regional variations (whereas fear decayed at different rates at different distances from Orlando), suggesting its potential usefulness as a central variable of interest in research on the public's response to terror. Furthermore, given that emotions can exhibit a contagion effect on social media [32, 33], the relative temporal and geographic stability of anger may suggest a particularly strong ability for it to propagate through social networks (compared to fear, which exhibited more regional variation and diminished more quickly over time), though this assertion should be tested in future research. If this is truly the case, then specifically monitoring angry social media posts in an event's aftermath can help officials better predict which sentiments and ideas may propagate to others and remain in circulation the longest.

There are, of course, many limitations of any psychological investigation involving social media. Participants were inherently selected by outcome, since they were only included if they had a Twitter account and specifically chose to respond to the Orlando attack. Many potential covariates of interest (e.g., gender and other demographic characteristics) were unable to be studied due to the difficult nature of ascertaining individual-level data from social media. Thus, our inclusion of income, population density, and political orientation merely represents an attempt to estimate variables that might have been correlated with physical distance, rather than an exhaustive set of relevant individual difference variables. Another limitation of text-analytic studies is small effect sizes, given the noise inherent to studies of natural language, and it is important to emphasize that such effects are much more valuable from a theory-building standpoint than a predictive standpoint. As aforementioned, sentiment expressed on social media should not necessarily be taken as an unbiased measure of an individual's true emotional experience, given the many motivations that can

drive individual social media behavior (desire to impress followers, conformity to the behavior of one's social network, etc.). Furthermore, we selected an event that victimized a specific minority community (LGBTQ+), and levels of fear and anger may have been affected by this unique characteristic of the attack (e.g., more emotional intensity from users that identify with the LGBTQ+ community).

Still, this data's value lies in its ability to convey contemporaneous reactions to a highly publicized terror attack across a wide geographic area, while allowing us to compare individuals' emotional expressions to baseline activity and control for potential confounds (income, urban vs. rural residence, political orientation). Furthermore, as other researchers have pointed out [34], social media studies are valuable in their ability to generate data-driven hypotheses for future study in more controlled settings. Our results have suggested that anger may be a more predictable and stable emotional response to terror attacks than fear, a contrast that this study's focus on temporal *and* geographic distance is in a unique position to draw.

5. References

- [1] G. Blalock, V. Kadiyali, and D.H. Roberts. "Driving fatalities after 9/11: A hidden cost of terrorism", *Applied Economics*, 2009, 41(14), pp. 1717-1729.
- [2] G. Fairbrother, J. Stuber, S. Galea, A.R. Fleischman, and B. Pfefferbaum. "Posttraumatic stress reactions in New York City children after the September 11, 2001, terrorist attacks", *Ambulatory Pediatrics*, 2003, 3(6), pp. 304-311.
- [3] S. Galea, D. Vlahov, H. Resnick, J. Ahern, E. Susser, J. Gold, M. Bucuvalas, and D. Kilpatrick. "Trends of probable post-traumatic stress disorder in New York City after the September 11 terrorist attacks", *American Journal of Epidemiology*, 2003, 158(6), pp. 514-524.
- [4] R.D. Marshall, R.A. Bryant, L. Amsel, E.J. Suh, J.M. Cook, and Y. Neria. "The psychology of ongoing threat: Relative risk appraisal, the September 11 attacks, and terrorism-related fears", *American Psychology*, 2007, 62(4), pp. 304-316.
- [5] J.S. Lerner, R.M. Gonzales, D.A. Small, and B. Fischhoff. "Effects of fear and anger on perceived risks of terrorism: A national field experiment", *Psychological Science*, 2003, 14(2), pp. 144-150.
- [6] L.J. Skitka, C.W. Bauman, N.P. Aramovich and G.S. Morgan. "Confrontational and preventative policy responses to terrorism: Anger wants a fight and fear wants 'them' to go away", *Basic and Applied Social Psychology*, 2006, 28(4) pp. 375-384.

- [7] A.A. Marsh, N. Ambady, and R.E. Kleck. "The effects of fear and anger facial expressions on approach- and avoidance-related behaviors", *Emotion*, 2005, 5(1), pp. 199-224.
- [8] C.S. Carver and E. Harmon-Jones. "Anger is an approach-related affect: Evidence and implications", *Psychological Bulletin*, 2009, 135(2), pp. 183-204.
- [9] A. Barnett. "Has successful terror gone to ground?", *Risk Analysis*, 2015, 35(4), pp. 732-740.
- [10] W. Gaissmaier and G. Gigerenzer. "9/11, Act II: A fine-grained analysis of regional variations in traffic fatalities in the aftermath of the terrorist attacks", *Psychological Science*, 2012, 23(12), pp. 1449-1454.
- [11] E. Cardena, J.M. Dennis, M. Winkel, and L.J. Skitka. "A snapshot of terror: Acute posttraumatic responses to the September 11 attack", *Journal of Trauma and Dissociation*, 2005, 6(2), pp. 69-84.
- [12] B. Fischhoff, R.M. Gonzales, D.A. Small, and J.S. Lerner. "Judged terror risk and proximity to the World Trade Center", *Journal of Risk and Uncertainty*, 2003, 26(2), pp. 137-151.
- [13] Y.R. Lin, D. Margolin, and X. Wen. "Tracking and analyzing individual distress following terrorist attacks using social media streams", *Risk Analysis*, 2017, 37(8), pp. 1580-1605.
- [14] Y.R. Lin and D. Margolin. "The ripple of fear, sympathy and solidarity during the Boston bombings", *EJP Data Science*, 2014, 3(1).
- [15] K. Gray and D.M. Wegner. "Dimensions of moral emotions", *Emotion Review*, 2011, 3(3).
- [16] J.S. Lerner, D. Keltner. "Fear, anger, and risk," *Journal of Personality and Social Psychology*, 2001, 81(1).
- [17] M. Baucum, H. Rosoff, R. John, W. Burns, and P. Slovic. "Modeling public responses to soft-target transportation terror", *Environment Systems and Decisions*, 2018, 1-11.
- [18] O. Oh, M. Agrawal, and H.R. Rao. "Community intelligence and social media services: A rumor theoretic analysis of tweets during social crises", *Mis Quarterly*, 2013, 37(2).
- [19] S. Oh and S. Y. Syn. "Motivations for sharing information and social support in social media: A comparative analysis of Facebook, Twitter, Delicious, YouTube, and Flickr," *Journal of the Association for Information Science and Technology*, 2015, 66(10), 2045-2060.
- [20] K. Starbird and L. Palen. "Pass it on?: Retweeting in mass emergency," *Proceedings of the 7th Annual ISCRAM Conference*, 2010.
- [21] J. Ooms. "cld2: Google's compact language detector 2", R package version 1.1, 2017, <https://CRAN.R-project.org/package=cld2>
- [22] A.C. Stapleton. "Timeline of the Orlando nightclub shooting", CNN, 2016, retrieved from <https://www.cnn.com/2016/06/12/us/orlando-shooting-timeline/index.html>
- [23] J.R. Wallace. "Imap: Interactive mapping", R package version 1.32, 2012, <https://CRAN.R-project.org/package=Imap>
- [24] United States Census Bureau, "2017 census U.S. gazetteer files", text file, 2017, retrieved from <https://www.census.gov/geo/maps-data/data/gazetteer.html>
- [25] United States Census Bureau, "2010 census summary file 1", comma separated values file, 2010, retrieved from <https://factfinder.census.gov/faces/tableservices/jsf/pages/productview.xhtml?src=CF>
- [26] A. Culotta, N.K. Ravi and J. Cutler. "Predicting the demographics of Twitter users from website traffic data", *Proceedings of the 29th AAAI Conference on Artificial Intelligence*, 2015, AI Access Foundation, pp. 72-78.
- [27] J. Golbeck and D. Hansen, "Computing political preference among Twitter followers", *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, 2011, ACM, pp. 1105-1108.
- [28] J.W. Pennebaker, R.J. Booth, and M.E. Francis. "LIWC 2007: Linguistic inquiry and word count", Austin, TX, 2007.
- [29] J.A. Boscarino, C.R. Figley, R.E. Adams. "Fear of terrorism in New York after the September 11 terrorist attacks; Implications for emergency mental health and preparedness", *International Journal of Emergency Mental Health*, 2003, 5(4), pp. 199-209.
- [30] L. Huddy, S. Feldman, G. Lahav and C. Taber. "Fear and terrorism: Psychological reactions to 9/11", *Framing terrorism: The news media, the government, and the public*. Eds. P. Norris, M. Kenn and M. Just, Psychology Press, 2003, pp. 255-278.
- [31] A.M. Nellis and J. Savage. "Does watching the news affect fear of terrorism? The importance of media exposure on terrorism fear", *Crime and Delinquency*, 2012, 58(5), pp. 748-768.
- [32] A.D. Kramer, J.E. Guillory, and J.T. Hancock. "Experimental evidence of massive-scale emotional contagion through social networks", *Proceedings of the National Academy of Sciences*, 2014, 111(29).

[33] S. Stieglitz and L. Dang-Xuan. “Emotions and information diffusion in social media—sentiment of microblogs and sharing behavior”, *Journal of Management Information Systems*, 2013, 29(4), 217-248.

[34] J. Hoover, K. Johnson, R. Boghrati, J. Graham and M. Deghani. “Moral framing and charitable donation: Integrating exploratory social media analyses and confirmatory experimentation”, *Collabra: Psychology*, 2018, 4(1).