

How Social Ties Influence Hurricane Evacuation Behavior

DANAË METAXA-KAKAVOULI, Stanford University, USA

PAIGE MAAS, Facebook Inc.

DANIEL P. ALDRICH, Northeastern University, USA

Natural disasters carry enormous costs every year, both in terms of lives and materials. Evacuation from potentially affected areas stands out among the most critical factors that can reduce mortality and vulnerability to crisis. We know surprisingly little about the factors that drive this important and often life-saving behavior, though recent work has suggested that social capital may play a critical and previously underestimated role in disaster preparedness. Moving beyond retrospective self-reporting and vehicle count estimates, we use social media data to examine connections between levels of social capital and evacuation behavior. This work is the first of its kind, examining these phenomena across three major disasters in the United States—Hurricane Harvey, Hurricane Irma, and Hurricane Maria—using aggregated, de-identified data from over 1.5 million Facebook users. Our analysis confirms that, holding confounding factors constant, several aspects of social capital are correlated with whether or not an individual evacuates. Higher levels of bridging and linking social ties correlate strongly with evacuation. However, these social capital related factors are not significantly associated with the rate of return after evacuation.

CCS Concepts: • **Human-centered computing** → **Social networks**; **Empirical studies in collaborative and social computing**; **Empirical studies in HCI**; *Social media*; *Social network analysis*;

Additional Key Words and Phrases: Crisis informatics; social capital; hurricane evacuation; social networks

ACM Reference Format:

Danaë Metaxa-Kakavouli, Paige Maas, and Daniel P. Aldrich. 2018. How Social Ties Influence Hurricane Evacuation Behavior. *Proc. ACM Hum.-Comput. Interact.* 2, CSCW, Article 122 (November 2018), 16 pages. <https://doi.org/10.1145/3274391>

1 INTRODUCTION

The number and cost of disasters have been rising over the past decades [19]. Many of the deaths that occur during events such as flooding, fires, hurricanes, and mudslides could be prevented by evacuating from vulnerable areas. For example, in Montecito, California during mudslides in January of 2018, local authorities and disaster managers tried to warn “residents of an approaching storm that caused the mudslides” through a variety of channels, “including emails, social media alerts, press releases and even deputies going door to door in some areas” [36]. Local authorities knocked on 1400 doors to get the message out, but despite these efforts, not all residents evacuated and nearly two dozen lost their lives [39]. Why residents remain in vulnerable areas after receiving evacuation orders or information about imminent risks is a critical public policy question. Most research on evacuation behavior relies on data such as retrospective surveys of evacuees (or would-be evacuees) or body and vehicle counts, along with best guesses and computer simulation. More recent attempts to gather data on evacuation behavior have surveyed evacuees as they depart

Authors’ addresses: Danaë Metaxa-Kakavouli, Stanford University, Stanford, CA, USA, metaxa@stanford.edu; Paige Maas, Facebook Inc. pmaas@fb.com; Daniel P. Aldrich, Northeastern University, Boston, MA, USA, daniel.aldrich@gmail.com.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

© 2018 Copyright held by the owner/author(s). Publication rights licensed to the Association for Computing Machinery. 2573-0142/2018/11-ART122 \$15.00
<https://doi.org/10.1145/3274391>

vulnerable areas [8]. While lacking qualitative dimensions, large-scale analysis of social media data can allow for a more nuanced picture of actual decision making and evacuation behavior during crises.

In this paper, we examine evacuation in the context of three major disasters—Hurricane Harvey, Hurricane Irma, and Hurricane Maria—affecting coastal regions in the United States (including Puerto Rico) during the summer of 2017. In particular, we focus on the connection between evacuation choices and social capital, as measured by metrics describing Facebook users’ social ties. The use of Facebook ties to study disasters has been established by previous work in *crisis informatics*, a “multidisciplinary field combining computing and social science knowledge of disasters” [31].

This paper makes several contributions to the literature. It is among the first large-scale studies of social capital in the context of disasters using social network data. While scholars have suggested the possibility of using social media data to study disasters [4, 31, 37], we have done so at an unprecedented scale; rather than relying on user content production, we are able to use aggregated, de-identified data to study all Facebook users in the affected areas through the lens of their online social ties. Importantly, our work does not rely on retrospective surveys that ask survivors of crises to recall how they behaved during the event [20] nor on small samples of individuals in the process of evacuating [8]. While those studies have been standard practice as a way of collecting data on evacuation behaviors, and can provide nuanced insights by asking more open-ended questions of their participants, there are some inherent limitations of relying on the recall accuracy of respondents remembering events that happened in the past, which we are able to avoid. Further, we are also able to capture a broader sample of evacuee and non-evacuee behaviors than is often possible when conducting in-person interviews. Other work leveraging large scale analysis of post-crisis evacuation based on GPS data or user-generated Twitter content has lacked demographic and social capital data on their populations of interest [25]. Finally, in this article we study multiple disasters over space and time, moving beyond past research on evacuation, which has typically focused on a single case study, such as one hurricane or fire (e.g., [5]). By analyzing data for three hurricanes, we can better understand the consistency of the observed relationships and come to broader conclusions about behavior in the context of U.S. hurricane evacuation.

2 RELATED WORK

Prior literature motivating this work has connected social capital with community resilience in the context of disasters, and has also used social networking site data as proxies for social capital.

2.1 Social Capital and Disasters

Researchers have defined social capital as the “resources embedded in one’s social networks, resources that can be accessed or mobilized through ties in the networks” [24]. There are several central forms of social capital, including three main types of connections examined in this work: *bonding*, *bridging*, and *linking* capital [40, 45]. Bonding ties occur most often between individuals who are quite similar in terms of ethnicity, religion, language, and educational background, and often include ties with family and extended kin; bridging ties, in contrast, often occur through institutions such as schools, clubs, and workplaces, and connect people from different backgrounds [33, 40]. For example, a Parent Teacher Association, a club soccer league, or a faith-based community might provide ways for individuals from different backgrounds to create bridging ties. The final type of connection we investigate is linking social capital which connects people to local government entities and other formal sources of power and resources [45].

In recent decades, scholars studying natural disasters have argued for the importance of social capital, in addition to traditionally-considered factors like physical infrastructure, as a crucial aspect of disaster survival and recovery [11, 28]. Prior to this shift, most research focused on the

Displacement from Hurricane Irma



Fig. 1. Visualizing the exodus of Miami-area residents in the days prior to Hurricane Irma’s landfall, with each dot representing an aggregate group of people within 0.5 latitude/longitude degrees, colored by evacuees (in blue) and non-evacuees (in red). Evacuation began when Irma was first predicted to affect the area, on September 6, and continued through its landfall in Miami, on September 10.

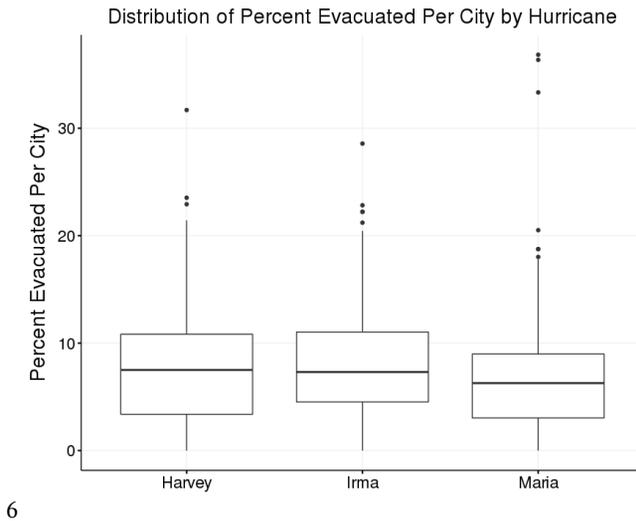
role of governmental response, policy surrounding the allocation of financial and other resources, and long-term economic trends [1]. A growing body of qualitative and quantitative evidence has demonstrated that grassroots, local neighborhood groups and community members serve as a critical element in disaster survival and recovery [6, 12, 27]. While some researchers have shown that those with stronger social networks are less likely to evacuate [8], others have shown that strong bonding ties created higher rates of evacuation in populations at risk [2]. These analyses did not consider wider characteristics of social networks nor were they done at scale.

2.2 Social Capital and Social Networks

Scholars have shown that social networking sites (SNSs) serve as effective vehicles for capturing social capital. Research using data from Nextdoor has shown that communities that were more active online also had higher levels of social capital in the offline world, as measured through organizational strength and civic engagement [30]. Facebook behavioral metrics, such as a user’s number of friends and number of messages sent, have been similarly shown to predict social tie strength and reflect real world connections [16]. Among college students, time spent on Facebook has been shown to correlate positively with greater levels of social capital, in particular, bridging ties [13]. Further, the relationship between a user’s specific communication style on Facebook and social capital metrics shows that social networking sites not only reflect users’ existing social capital but may actively aid in its development [14].

While SNS data has been found to reflect (and, in some cases, increase) all three types of social capital, such data capture bridging capital, which relies on weak ties, particularly well. Using social networks like Facebook, people are able to find and, with relatively little effort, keep in contact with geographically distant friends and acquaintances. In contrast, bonding capital relationships are often maintained in a plethora of other ways, including regular face-to-face communication, and may not be as well reflected in SNS data [7]. Civic engagement and political participation, factors associated with linking capital, have also been found to correlate positively with Facebook use [41].

In the domain of computer-supported cooperative work (CSCW), crises and disasters have been studied from many angles, including preparedness, emergency response, recovery, and mobilization [32]. There is substantial work on the technologies used by those experiencing disasters, in particular Twitter [21, 38, 43], as well as work examining the role of emergency responders and journalists



6

Fig. 2. In all three hurricanes examined, the overall rate of evacuation was less than 10%; however, the distribution of evacuation rates at the city level was skewed, with a small number of cities evacuating at much higher rates, up to nearly 40%.

[9, 17]. Taken together, this work supports the idea that social ties and digital tools can together be very effective in gathering knowledge, coordinating resources, and mobilizing resources before, during, and after disasters.

3 DATA COLLECTION

This work focuses on evacuation during severe category 4 and 5 hurricanes affecting the United States (including Puerto Rico) during the 2017 hurricane season: hurricanes Harvey, Irma, and Maria. For each hurricane, we identified all major metropolitan areas for which a mandatory evacuation or state of emergency was issued, and manually produced the following metadata: (1) minimum and maximum latitude and longitude to create a bounding box around that area, (2) the date at which it was first reported that the area would need to evacuate (the *evacuation start date*), and (3) the date on which the hurricane first affected that area (the *disaster date*). We defined the time window of interest to range from three days before the evacuation start date (in order to serve as a baseline on user behavior before the disaster was projected to affect them) to 14 days after the disaster date, a time-frame chosen to capture the majority of evacuees' return behavior while maintaining data sets of a computationally feasible size. These time windows were approximately three weeks long, depending on the amount of time between the evacuation start and disaster dates.

3.1 Identifying affected users

For each hurricane, we identified all cities in the associated bounding box, and then identified all Facebook users whose estimated home city was in this set. We chose to select people using a bounding box centered around areas with evacuation orders rather than specifically selecting people under such evacuation orders, as previous literature has suggested that evacuation orders themselves, while useful for identifying affected areas, have relatively low impact on people's decision-making [10]. Estimated home city is determined by a person's information on Facebook, including the stated city on their Facebook profile, and device and connection information.

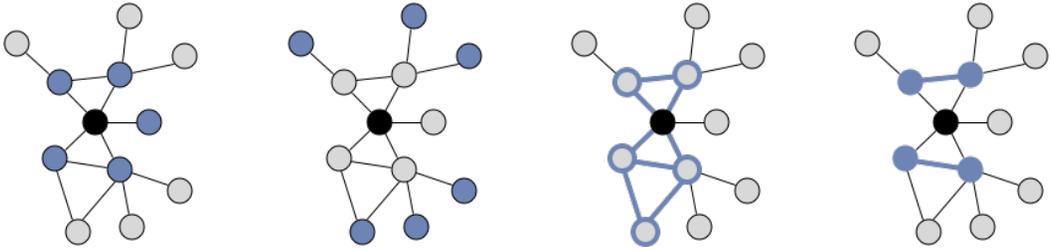


Fig. 3. These graphs illustrate a small-scale example of a social network, with each node representing a user and each edge a friendship tie between two users, with the central user whose network we are examining colored in black. From left to right: (1) The first graph highlights in blue the person's first degree network (their number of friends); (2) the second graph highlights second-degree network size (friends of friends); (3) the third graph highlights two bi-connected components (two interconnected groups of friends which the person is part of); (4) the fourth graph highlights all pairs of the person's friends who are friends with each other, from which clustering coefficient is computed.

Having identified these users, we filtered to include only those whose daily location in the three days before the evacuation start date indicated that they were within an average of 50 miles away—an approximate threshold for more than an hour away by car—from their home city. A person's daily location was the city or town in which they spent most of their time that day. This was done to exclude anyone whose estimated home city was incorrect, or who was for any reason not present in the affected area prior to the disaster. For each of the affected users, we obtained the person's daily location during each day in the time window of interest. People who did not log into Facebook in this period lack up-to-date location information, so we filtered to only include users who were active on Facebook for at least 2/3 of the time window of interest.

3.2 Identifying evacuation behavior

Having identified people affected by these disasters, we examined each person's location in the days between the evacuation start date and disaster date, during which evacuation would have occurred. We categorized a person as an evacuee if their location was an average of at least 50 miles from their home city in this time range. All other people are marked as non-evacuees. For evacuees, we marked the first day following the disaster date on which they were within 50 miles of their home city to be the date at which the evacuee returned from evacuation. For example, Figure 1 visualizes the exodus of Miami-area residents in the days prior to Hurricane Irma.

3.3 Independent variables

3.3.1 Social capital metrics. We joined this data with metrics relating to each user's social ties on the site, chosen as proxies for social capital metrics. The choice of metrics was informed by existing literature [7, 22, 41]. These metrics, also summarized in Table 1, were:

- **First degree network size.** This metric quantifies a person's number of friends, with stratification into eight equally sized groups, or *octiles*. This metric has traditionally been conceived of as a measure of bridging capital, as it includes a person's weaker acquaintance ties [13]. However, other literature has found no effect of number of Facebook friends on an individual's bridging capital, and has instead found a positive effect on bonding capital [42], possibly because those with more friends are more likely to have the emotionally close relations comprising bonding capital, or to have such relations in their Facebook network, so in this work we consider it a form of bonding capital.

Social capital type	Definition	Operationalization in Facebook metrics
Bonding	Strong connections between close friends and family [33, 40]	First degree network size; clustering coefficient
Bridging	"Weaker" ties which connect friends of friends, people with a single interest in common, etc. [18, 33]	Second degree network size; biconnected components
Linking	Ties between people and those in positions of legitimacy and authority, such as politicians [45]	Politicians followed

Table 1. A summary of the three main types of social capital examined in this paper, along with the way we operationalize them in this work.

- **Clustering coefficient.** For every pair of a person’s Facebook friends, the clustering coefficient quantifies the proportion of those who are friends with each other on Facebook (see Figure 3). This metric was also bucketed into octiles. This metric measures the cohesiveness of a network [44], and has been shown to correlate with better collaboration between connected parties in a social network [29]. We anticipated this metric would behave consistently with bonding social capital, since it describes the density of a person’s friend group—people with many connected friends are more likely to be embedded in a network of strong, close ties.
- **Second degree network size.** This attribute of a user’s network measures their number of unique friends-of-friends, bucketed into octiles. While similar to the first degree network size, prior work suggests that variety (or lack thereof) among friends-of-friends can provide fundamentally different and valuable information about people’s social networks [3]. This measure examines heterogeneity in a person’s extended network, as such is a proxy for bridging capital which relies on weaker, more diverse ties [40].
- **Biconnected components.** When considering a person’s friendship graph, this measures the number of maximal biconnected subgraphs that exist in it—in other words, as visualized in Figure 3, this quantifies how many separate groups of friends the person is part of, again bucketed into octiles. This measure was introduced as a measure of the cohesion of a network, since “biconnected components (loops) allow information to flow freely throughout the community” [26]. This kind of wider group cohesion that relies on intermediaries aligns with bridging capital.
- **Politicians followed.** For each Facebook user, this measures the number of politicians’ Facebook pages the person followed on Facebook prior to the disaster in question. Since the distribution of politician follows is very skewed, with many people following no politicians and a long tail of few people following many politicians, we bucketed this metric into three categorical variables (0, 1-2, or more than 2 politicians followed). We operationalize this metric as a proxy for linking capital.

Rather than treat variables as continuous, we bucketed them to avoid making assumptions about the linearity of these measures. We chose octiles in order to balance granularity with having enough users in each bucket to identify trends in the data.

Disaster	Dates	Number cities	Number affected users	Number evacuees	Percent evacuees	Gender breakdown	Median age
Hurricane Harvey	August 17–September 2, 2017	79	103,388	8,788	8.50%	59% women, 40.4% men, 0.6% other	45-54yo
Hurricane Irma	August 30–September 16, 2017	178	779,636	60,032	7.70%	56.9% women, 42.4% men, 0.8% other	35-44yo
Hurricane Maria	September 16–October 3, 2017	76	35,492	2,496	7.03%	48.3% women, 50.5% men, 1.3% other	35-44yo

Table 2. A summary of descriptive statistics for the final data set, including the name of each disaster as well as the number of affected cities in our bounding box, number of affected users identified in that region, number of those affected users who evacuated, a percentage representing the share of affected users who evacuated, gender of affected users, and affected user median age.

We considered first degree network size and clustering coefficients to be proxies for bonding capital, as they describe the people’s immediate network—its size and its density, respectively—based on prior literature [33, 42]. Second degree network size serves as a proxy for bridging capital, as it identifies the breadth of reach of a person’s network, and the weaker, indirect resources a user could mobilize through their network [18]. We use the number of biconnected components in a user’s friendship graph as a measure of bridging capital, since it describes how many separate circles of friends the person is part of (heterogeneity is an important characteristic of bridging ties) [26]. Linking capital involves vertical ties between residents of affected areas and people in positions of power [23, 40]. We adopt the number of politicians followed as a measure of linking capital, as it describes connections between users and their representatives, formal sources of power who are also likely to have social media presences. While these follows are not reciprocal, meaning a politician in question may not follow a user who follows her, people who follow more politicians may be more politically engaged and invested in reaching their elected officials.

3.4 Descriptive statistics

Our analysis of several natural disasters at such large scale is the first of its kind; as such, the following descriptive statistics present an important addition to the literature in crisis informatics, and provide insight into the people affected by these events. These data are described below, and summarized in Table 2.

In collecting data about Hurricane Irma, bounding boxes were drawn around four metropolitan areas: Charleston, South Carolina; Jacksonville, Florida; Miami, Florida; and Savannah, Georgia. These are all coastal areas affected by the hurricane, and after filtering for people in those areas during the days before the disaster (around August 25, 2017), our data set contained 779,636 distinct people from 178 towns. Of these people, 7.7% evacuated the disaster area ($M = 8.1$, $SD = 5.1$). The median age in this group was 35-44 years old, which skewed slightly (56.9% women, 42.4% men, 0.8% other or unknown).

Two major metropolitan areas were affected by Hurricane Harvey and faced mandatory evacuation orders in mid August of 2017: Galveston, Texas and Corpus Christi, Texas. The bounding box around these locations included a total of 79 towns, and 103,388 unique affected users. Of these,

Field	Description
Age bracket	Ordinal; categorizing users by age bracket (e.g., 35–44, 65+, etc.)
Gender	Categorical; identifying user gender (man, woman, or other/unknown)
Home city identifier	Categorical; user home city was included, to control for geographic differences
First degree network size	Numeric; a user's number of friends
Clustering coefficient	Numeric (between 0 and 1); the density of a user's network
Second degree network size	Numeric; user's number of unique friends-of-friends
Biconnected components	Numeric; the number of separate groups in a user's immediate network
Politicians followed	Categorical; the number of politician accounts followed (levels: 0, 1-2, >2)
Evacuee status	Binary; coding whether or not a user evacuated prior to the disaster
Days until returned	Numeric; for evacuees, the number of days before they returned to their home city (no value for non-evacuees)

Table 3. The dataset in this work was derived from historical log data, and included 10 variables per user, including age bracket and gender demographics, five independent variables describing user social capital, and two outcome variables (evacuation status, and days until evacuee return). These variables were aggregated across users for use in our models.

8.5% were evacuees ($M = 8.0$, $SD = 6.1$), and the group as a whole had a median age of 45-54 years old. This group also skewed slightly towards women (59.0% women, 40.4% men, 0.6% other or unknown).

Hurricane Maria affected the entire island of Puerto Rico; the bounding box encompassing this entire region included 76 cities, and identified 35,492 affected users in the area during mid-September of 2017. This number of people is quite low relative to the entire population of Puerto Rico (over 3 million), perhaps because of connectivity loss (people with location data on fewer than 2/3 of the days in our time window of interest were excluded). However, evacuation rates were similar to the other disasters, at 7% ($M = 7.7$, $SD = 7.4$). This group skewed slightly towards men (48.4% women, 50.5% men, 1.3% other or unknown), and had an median age of 35-44 years old.

Notably, overall evacuation rates across all three disasters were low, under 10%. As is shown in Figure 2, this distribution was skewed, with most cities displaying low evacuation rates, but with a small number of presumably higher-risk cities showing much higher rates of evacuation, up to 40%.

3.5 Data Ethics and Privacy

The ethics of this research were carefully considered throughout the research process, and considerable steps were taken to handle the data ethically and preserve user privacy. The data set is comprised of historical log data for people using the Facebook application with Location Services turned on, and the analysis did not impact anyone's experience on Facebook. All data was aggregated and de-identified to ensure no individual names or text content were handled. Data was handled securely on Facebook servers, and in accordance with Facebook's data policy. The research passed a rigorous internal methodological review by a cross-functional team including data science, legal, privacy, and security experts before any analysis was performed. All variables included in the final data set per user are recorded in Table 3. Methodological decisions were also made in the interest of preserving privacy. For example, demographic variables included in our data set and models were in broad categories (age bucketed into age brackets, location no more specific than the city level). There are many benefits of research allowing for a better understanding of disaster evacuation behavior, including preventative strategies such as the identification of vulnerable populations and more effective targeting of evacuation orders where needed. Data were handled

and analyzed to maximize these potential benefits for the vulnerable population studied while preserving people's privacy.

4 MODELS & FINDINGS

4.1 Hypotheses

We made several hypotheses as to the effect of our social tie metrics on evacuation behavior.

- H1. We expected that deeper reservoirs of bonding social capital would decrease the likelihood of a user evacuating, as a stronger local network may encourage sheltering in place alongside one's community.
- H2. Conversely, we expected that having higher levels of bridging capital would increase the likelihood of evacuation, as diverse bridging ties may provide people with the ability to mobilize geographically diverse resources outside the affected area, and ease the high financial and emotional costs of evacuation. Bridging ties also bring diverse types of information and may help counter opinions which would dismiss authorities' warnings.
- H3. Similarly to H2, we expected that having more linking capital would increase the likelihood of evacuation, as mandated evacuations are directed through governments and other formal power structures, and users more proximal to those messages may be more likely to follow them; being connected to the governor's office or the page of a politician may cause a user to more strongly consider or trust the recommendation to evacuate.

We find support for all three hypotheses in the context of users' likelihood of evacuation. However, we do not find any significant associations between the social ties we examined and evacuees' rate of return, perhaps because our data set is limited to information from a two-week window after each disaster, cutting off the long tail of evacuee return.

4.2 Logistic regressions predicting evacuation

For each of the three hurricanes, we fit a single logistic regression to predict a binary variable encoding whether each user evacuated (1 for yes, 0 for no). As inputs to the models, we used the four social capital metrics along with the user's gender, Facebook activity level, age, and home city to control for individual and city-level differences such as city population size, quality of physical infrastructure, and local disaster severity. Figure 4 presents the odds ratios for the features in each model, showing similar associations for Hurricane Irma and Hurricane Harvey, with less clear associations for Hurricane Maria.

4.2.1 Hurricane Irma and Hurricane Harvey. We find that Hurricane Irma and Hurricane Harvey show largely consistent results with one another.

Bonding Ties. For these two events, we find that having more friends on Facebook is associated with a lower likelihood of evacuation. In particular, the odds of evacuation for people in the category with most friend connections were approximately 60% of the odds of evacuation for people in the group with fewest friends, controlling for all other factors. Our other measure of bonding capital was also associated with lower likelihood of evacuation for Hurricane Irma: the odds of evacuation for people in the top five of the eight categories of clustering coefficient were 85% of those in the three categories with lowest clustering coefficient.

Bridging Ties. In contrast, a user's second-degree network size (distinct friends-of-friends), our measure of bridging capital, was strongly associated with a greater likelihood of evacuation, with people in the highest category being twice as likely to evacuate as those in the lowest category for Hurricane Harvey, and 3.5 times more likely for Hurricane Irma. Our second measure of bridging capital, the number of biconnected components, showed different associations with evacuation

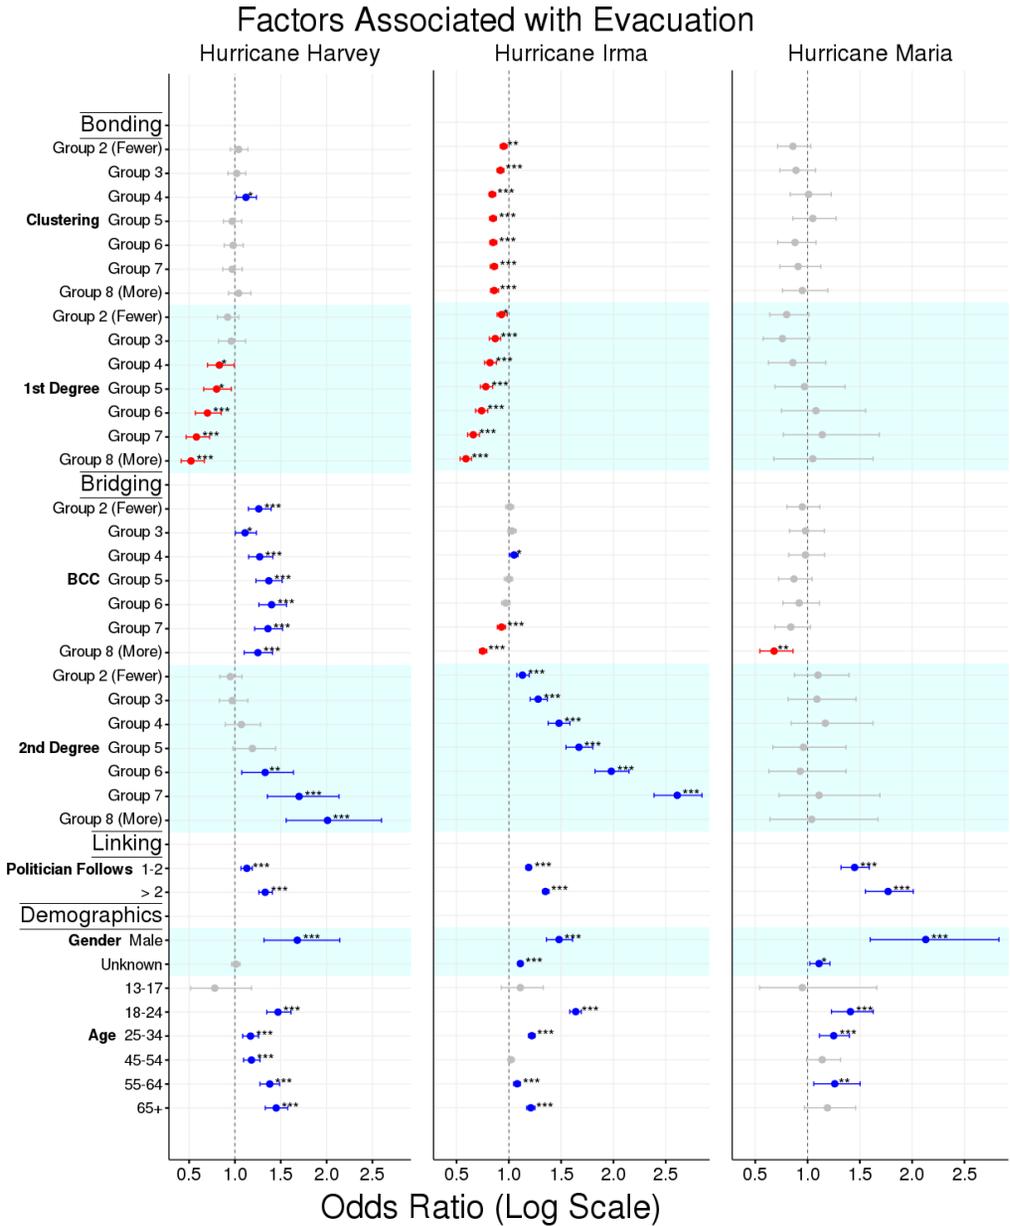


Fig. 4. Plots of the odds ratios for logistic regressions modeling evacuation, with statistically significant differences labeled with asterisks (* for $p < 0.05$, ** for $p < 0.01$, *** for $p < 0.001$), and color-coded by the direction of the effect (red for association with lower likelihood of evacuation, and blue for higher likelihood to evacuate).

outcome in each crisis. For users affected by Hurricane Irma, people with the most distinct friend groups (highest category of biconnected components) were much less likely to evacuate relative to

those with few different circles of friends, while during Hurricane Harvey the middle categories were more likely to evacuate than users at either extreme.

Linking Ties. Users that followed 1-2 or more than 2 politicians on Facebook were also more likely to evacuate, controlling for all other factors including Facebook activity level. This effect was smaller, with people following more than two politicians only 33-35% more likely to evacuate after controlling for other factors; however, convincingly this effect showed very little variance and was consistent in magnitude and dose response across all three hurricanes.

With regard to the other variables included in our models such as gender, age, or hometown, we found that men and women evacuated at roughly similar rates, and that both younger and older residents (those younger than 34 and older than 55) were more likely to evacuate than those in the middle age range (35-44). There was one exception to this pattern: children aged 13-17 were not significantly more or less likely to evacuate than middle aged users, consistent with the idea that families make evacuation decisions as a unit.

4.2.2 Hurricane Maria. In the case of Hurricane Maria the relatively small sample size, likely due to the small population of people with consistent data, as well as the difficulty of evacuating from an island, resulted in data with too much variance to identify clear trends in most cases. As with the other two hurricanes, however, we found that politician follows were strongly correlated with likelihood to evacuate, and that young adults (aged 18-24) and elderly people (aged 55+) were more likely to evacuate relative to children (aged 13-17) and middle-aged individuals (aged 35-44). Additionally, as with with Hurricane Irma, the negative correlation of biconnected components with likelihood of evacuation was weakly observed, but was only significant when comparing people in the highest category to those in the lowest.

Notably, the odds ratios for Hurricane Maria had much larger standard errors. This may be due to the relatively small number of evacuees in our dataset. The primary reasons for this smaller sample size were the massive power outages and loss of network connectivity that plagued Puerto Rico in the aftermath of the hurricane, causing larger gaps in our social media-generated data set. Contrasting Hurricanes Irma and Harvey with Hurricane Maria, whose model gave more mixed results, it is clear that a large sample size is critical for studying a relatively rare behavior like hurricane evacuation (our dataset contained fewer than 2,500 evacuees for Hurricane Maria). Additionally, this raises the issue of geographical differences in disaster response—the decision and ability to evacuate from an island like Puerto Rico may be a very different decision than evacuating from less geographically isolated locations. In addition to the extra cost and difficulty of needing to evacuate by air, the vast majority of a Puerto Rican user’s social ties may have been on the island, making it difficult to find support from friends outside the affected region.

4.3 Cox proportional-hazards modeling time to return

Using the same social capital and demographics metrics, we fit Cox proportional hazards models to estimate the association between these features and the time at which evacuating users returned to the affected region. As shown in Figure 5, small differences were observed by gender (with women slightly slower to return than men), consistent with prior literature suggesting that women are more likely to act conservatively in response to perceived risk in the context of natural disaster evacuation [5]. However, we did not find any significant effects of social capital on evacuee return time.

When tracking evacuee return over the two-week window post-disaster, in the cases of hurricanes Irma and Harvey, a large majority—around 85% and 75% respectively—had returned within a two-week window. The rate of user return began with a sharp drop then evened out, suggesting that the remaining evacuees are part of a long tail returning more slowly to the area. In the case of

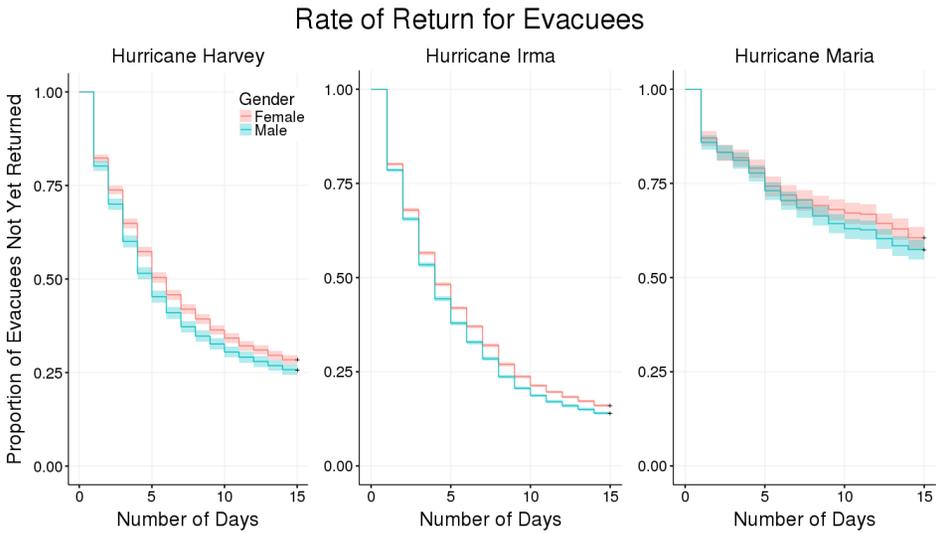


Fig. 5. Kaplan-Meier survival plots showing the proportion of evacuees who had returned to their home region in the two weeks following each hurricane, by gender.

Hurricane Maria, however, the picture is a very different one, with not even half of evacuees having returned after two weeks, again suggesting that Hurricane Maria had a substantively different impact than the other two hurricanes.

5 DISCUSSION

We find that several social capital proxies derived from social media use are associated with the decision to evacuate. Our results support research literature that argues for the influence of social capital in disaster resilience [2, 15], and that social ties may be a core factor driving evacuation decisions [10, 20]. Scholars, for example, have argued that “the dependability dimension of social connections has a significant effect on the decision to evacuate, with those who have more perceived functional support tending to not evacuate, as they feel comfortable and receive support with their local communities” [8]. We found support for this, in particular observing that ties associated with bonding capital, such as a user’s number of friends (first degree network size) and the density of their network (clustering coefficient), are related with lower likelihood of evacuation (the former more consistently than the latter). We further observe that ties associated with bridging and linking capital, measured by second-degree network size and the number of politicians followed, are associated with a greater likelihood of evacuation. We do not find strong effects for a user’s number of distinct friend groups (biconnected components), nor do we see strong associations with evacuees’ rate of return.

The strong positive correlation between second-degree network indicates that when people’s extended networks are broader-reaching, they more commonly evacuate. A possible explanation of this behavior is that people who are connected to more communities outside their immediate city of residence are more comfortable or able to temporarily relocate. Perhaps they know they can find support in farther-away towns, or this wider network is an indication that they have traveled more and are therefore more familiar with other locations to which they might evacuate. Broader bridging networks also provide more diverse types of information and may help users be more exposed to government messages regarding oncoming natural disasters and recommendations

around evacuation. A friend outside the area may reference a disaster that they experienced and encourage the individual to depart, for example. Similarly, the strong relationship between number of politicians followed and likelihood of evacuation suggests that individuals who are more closely connected to governmental sources of information may be more willing to trust and follow governmental instructions to evacuate in the days leading up to a hurricane. Meanwhile, the strong negative correlation between evacuation and number of friends suggests that when users' bonding ties are strongest they may choose to remain in the threatened area since they perceive themselves as having a strong support network in that area, and are therefore less inclined to leave.

We also see clear trends in evacuation behavior as it relates to age. Middle-aged people were less likely to evacuate than nearly every other group, with the exception of children (ages 13-17). This suggests both that people are likely to make an evacuation decision together with their whole family unit, and that those with families are less eager to evacuate perhaps due to the many logistical challenges of evacuating with children. Furthermore, even controlling for social capital metrics, young adults and older users are more likely to evacuate. This is encouraging, as the elderly may be particularly at risk in the case of a natural disaster, though it may be a function of the types of users on Facebook (perhaps elderly users with more technological literacy are more able or likely than other elders to evacuate).

We did not see consistent trends relating to two of our social capital metrics: biconnected components, and clustering coefficient. This might point to regional differences, as these metrics did appear to correlate significantly with evacuation in some disasters but not others, as shown in Figure 4. It may also suggest, however, that these metrics are less effective proxies of social capital than the other metrics in our models. We also did not find any statistically significant effects on evacuee rate of return. The lack of effects seen in these models could be due to a number of factors; for instance, although the number of people in these studies is large, we observe evacuation to be an uncommon behavior, and so the relatively small number of evacuees in our sample impedes sensitivity to small effect sizes.

Additionally, Hurricanes Harvey and Irma, while major disasters, were manageable enough that three quarters of evacuees had returned within two weeks. In these cases, where nearly all evacuees return home to rebuild, the impact of social capital may be less important than in cases where the decision to return is uncertain. In contrast, Hurricane Maria caused massive, sustained damage to Puerto Rico—it took over eleven months for power to be restored on the island [35]. The severity of the situation is reflected in our data, where two weeks after the disaster, not even half of evacuees had returned to the area. With widespread connectivity losses, a geographically bounded area, and so few evacuees returning in the first weeks following the disaster, modeling time to return is more challenging and more complex; this suggests that longer-term study or other approaches are needed in such cases. These results underscore prior findings that social influence and access to resources are critically important in evacuation scenarios [34]. In particular, comparing results from Hurricane Maria to Irma and Harvey reflects the need for resource access, as undertaking evacuation from an island is a much more resource-intensive process than elsewhere, putting the possibility of evacuation out of reach for many. These insights also convey that evacuee rate of return is best studied in the cases of disasters catastrophic enough that return is not inevitable, where geography allows for evacuation, and using a timeline of months, rather than weeks.

5.1 Limitations

Though the analysis we present here is much larger in scale than previous work in this domain, the sample studied is restricted to Facebook users who were relatively frequently active during three specific natural disasters in the United States and Puerto Rico; while this work is made possible by online log data, it is not a representative sample of all evacuees, nor does it capture the entirety

of evacuation behaviors. Facebook penetration is high in the contiguous United States, but it is substantially lower in Puerto Rico, and our data include no users outside these two regions; these findings must be interpreted through the lens of the population studied.

While this work presents an analysis at a scale and granularity unique in the disaster literature, our analyses do not directly include or control for features hypothesized in the disaster literature to influence evacuation behavior, such as the severity of the disaster, demographics like socioeconomic status, the number of friends known to be outside of the evacuation zone, or the actual evacuation behavior of one's friends. Including such features holds promise for future work which could combine richer social network data with information from census sources, surveys, and interviews. Similarly, this work examines longer-distance evacuation (evacuees relocating to more than 50 miles from home). Evacuation behavior can also include relocating to higher ground or better-protected places within one's local geography; due to limitations in the granularity of our data, which located a user's daily city location but not precise GPS location, we were not able to study this type of evacuation.

Due to the nature of these disasters and timing of our analysis, we study relatively short-term evacuation behavior, which may differ for disasters so catastrophic as to preclude resettlement in the area in a much longer time frame, and also does not examine phenomena like the impact of previous disaster exposure [10]. Such disasters may differ substantially from those examined here, particularly with regards to the effect of social ties on evacuee rate of return. Finally, while our data and scientific methodology offer many insights, we acknowledge and appreciate traditional disaster work, which leverages less automated methods such as interviews or surveys, and is also critical and informative to the study of disasters and resilience.

5.2 Future Work

In addition to the possibilities laid out in the previous section, there are many opportunities for future work in this space. For instance, the analyses we present would be enriched by adding data about evacuation-related content a user consumed on social media, including ads, posts from friends, and posts from government officials, to study the influence of specific messaging on evacuation behavior, something that could help direct evacuation efforts in the future. Additionally, other types of natural disasters such as earthquakes and wildfires would benefit from the analyses presented in this work to identify disaster type-specific evacuation behaviors. Similarly, it is promising to consider the application of this method, which can be applied regardless of disaster type, to make meaningful comparisons between different types of disasters. Finally, the combination of automated at-scale methods with insights yielded from surveys, interviews, and focus groups is an exciting direction that can benefit from the strengths of these different methods.

6 CONCLUSION

Disasters affect hundreds of thousands of people annually; the better we understand the choices made during those times of crisis, the better we can help those communities prepare and recover. In particular, this work examines the issue from a novel angle, studying user-level behavior as it relates to ties captured by social media. It is the first analysis of its kind which can be consistently applied to a wide range of locations and disaster types, allowing us to understand whether patterns are consistent between different disasters, as well as facilitating comparisons of different disaster types in the future. We find that social ties, in particular the breadth of someone's extended social network, have an effect on evacuation behavior leading up to hurricanes. Our methodology and findings have implications for federal public health and scientific agencies whose work includes informing the public about natural disaster response as well as facilitating disaster recovery. This work presents a methodology for identifying vulnerable populations which may have smaller or

more geographically centered networks, with the goal of engaging these populations and ensuring that they receive sufficient information to make informed decisions and, when necessary, undertake evacuation.

7 ACKNOWLEDGEMENTS

Danaë would like to thank Facebook's Core Data Science team, especially Alex Dow, for enabling this collaboration, as well as Eduardo Laguna Muggenburg for his invaluable guidance regarding social capital in the context of Facebook data. The authors would all like to thank Chaya Nayak and Graeme Burrows for facilitating this collaboration.

REFERENCES

- [1] Daniel P Aldrich. 2012. *Building resilience: Social capital in post-disaster recovery*. University of Chicago Press.
- [2] Daniel P Aldrich and Yasuyuki Sawada. 2015. The physical and social determinants of mortality in the 3.11 tsunami. *Social Science & Medicine* 124 (2015), 66–75.
- [3] Kristen M Altenburger and Johan Ugander. 2017. Bias and variance in the social structure of gender. *arXiv preprint arXiv:1705.04774* (2017).
- [4] Dimitris Ballas, Graham Clarke, Rachel S Franklin, and Andy Newing. 2017. *GIS and the Social Sciences: Theory and Applications*. Routledge.
- [5] Julie M Bateman and Bob Edwards. 2002. Gender and evacuation: A closer look at why women are more likely to evacuate for hurricanes. *Natural Hazards Review* 3, 3 (2002), 107–117.
- [6] John J Beggs, Valerie A Haines, and Jeanne S Hurlbert. 1996. Situational contingencies surrounding the receipt of informal support. *Social forces* 75, 1 (1996), 201–222.
- [7] Moira Burke, Robert Kraut, and Cameron Marlow. 2011. Social capital on Facebook: Differentiating uses and users. In *Proceedings of the SIGCHI conference on human factors in computing systems*. ACM, 571–580.
- [8] Jennifer Collins, Robin Ersing, and Amy Polen. 2017. Evacuation Decision-Making during Hurricane Matthew: An Assessment of the Effects of Social Connections. *Weather, Climate, and Society* 9, 4 (2017), 769–776.
- [9] Dharma Dailey and Kate Starbird. 2014. Journalists as crowdsourcerers: Responding to crisis by reporting with a crowd. *Computer Supported Cooperative Work (CSCW)* 23, 4-6 (2014), 445–481.
- [10] Kirstin Dow and Susan L. Cutter. 1998. Crying wolf: Repeat responses to hurricane evacuation orders. *Coastal Management* 26, 4 (1998), 237–252. <https://doi.org/10.1080/08920759809362356>
- [11] Russell R Dynes. 2005. Community social capital as the primary basis for resilience. (2005).
- [12] James R Elliott, Timothy J Haney, and Petrice Sams-Abiodun. 2010. Limits to social capital: Comparing network assistance in two New Orleans neighborhoods devastated by Hurricane Katrina. *The Sociological Quarterly* 51, 4 (2010), 624–648.
- [13] Nicole B Ellison, Charles Steinfield, and Cliff Lampe. 2007. The benefits of Facebook “friends.” Social capital and college students’ use of online social network sites. *Journal of computer-mediated communication* 12, 4 (2007), 1143–1168.
- [14] Nicole B Ellison, Charles Steinfield, and Cliff Lampe. 2011. Connection strategies: Social capital implications of Facebook-enabled communication practices. *New media & society* 13, 6 (2011), 873–892.
- [15] Markus Friberg and Michael Hjelm. 2015. Mass evacuation-human behavior and crowd dynamics-What do we know? *LUTVDG/TVBB* (2015).
- [16] Eric Gilbert and Karrie Karahalios. 2009. Predicting tie strength with social media. In *Proceedings of the SIGCHI conference on human factors in computing systems*. ACM, 211–220.
- [17] Athula Ginige, Luca Paolino, Marco Romano, Monica Sebillio, Genoveffa Tortora, and Giuliana Vitiello. 2014. Information sharing among disaster responders-an interactive spreadsheet-based collaboration approach. *Computer Supported Cooperative Work (CSCW)* 23, 4-6 (2014), 547–583.
- [18] Mark Granovetter. 1983. The strength of weak ties: A network theory revisited. *Sociological theory* (1983), 201–233.
- [19] D Guha-Sapir, P Hoyois, P Wallemacq, and R Below. 2017. Annual Disaster Statistical Review. *EMDAT* (2017). http://emdat.be/sites/default/files/adsr_2016.pdf
- [20] Samiul Hasan, Satish Ukkusuri, Hugh Gladwin, and Pamela Murray-Tuite. 2010. Behavioral model to understand household-level hurricane evacuation decision making. *Journal of Transportation Engineering* 137, 5 (2010), 341–348.
- [21] Amanda Lee Hughes and Leysia Palen. 2009. Twitter adoption and use in mass convergence and emergency events. *International Journal of Emergency Management* 6, 3-4 (2009), 248–260.
- [22] Matthew Jackson. 2017. A Typology of Social Capital and Associated Network Measures. (2017).
- [23] Ichiro Kawachi, Daniel Kim, Adam Coutts, and SV Subramanian. 2004. Commentary: Reconciling the three accounts of social capital. *International Journal of Epidemiology* 33, 4 (2004), 682–690.

- [24] Nan Lin. 1999. Building a network theory of social capital. *Connections* 22, 1 (1999), 28–51.
- [25] Yago Martín, Zhenlong Li, and Susan L Cutter. 2017. Leveraging Twitter to gauge evacuation compliance: Spatiotemporal analysis of Hurricane Matthew. *PLoS one* 12, 7 (2017), e0181701.
- [26] James Moody and Douglas R White. 2003. Structural cohesion and embeddedness: A hierarchical concept of social groups. *American Sociological Review* (2003), 103–127.
- [27] Brenda L Murphy. 2007. Locating social capital in resilient community-level emergency management. *Natural Hazards* 41, 2 (2007), 297–315.
- [28] Yuko Nakagawa and Rajib Shaw. 2004. Social capital: A missing link to disaster recovery. *International Journal of Mass Emergencies and Disasters* 22, 1 (2004), 5–34.
- [29] Keiichi Nemoto, Peter Gloor, and Robert Laubacher. 2011. Social capital increases efficiency of collaboration among Wikipedia editors. In *Proceedings of the 22nd ACM conference on Hypertext and hypermedia*. ACM, 231–240.
- [30] Courtney M Page. 2017. Estimating Effects of Social Capital on Online Social Engagement: The 2014 Napa Valley Earthquake. In *Conference, Midwest Political Science Association Annual*.
- [31] Leysia Palen and Kenneth M Anderson. 2016. Crisis informatics—New data for extraordinary times. *Science* 353, 6296 (2016), 224–225.
- [32] Volkmar Pipek, Sophia B. Liu, and Andruid Kerne. 2014. Crisis Informatics and Collaboration: A Brief Introduction. *Computer Supported Cooperative Work (CSCW)* 23, 4 (01 Dec 2014), 339–345. <https://doi.org/10.1007/s10606-014-9211-4>
- [33] Robert D Putnam. 2001. *Bowling alone: The collapse and revival of American community*. Simon and Schuster.
- [34] Jasmin K Riad, Fran H Norris, and R Barry Ruback. 1999. Predicting evacuation in two major disasters: Risk perception, social influence, and access to resources 1. *Journal of Applied Social Psychology* 29, 5 (1999), 918–934.
- [35] Leyla Santiago and Natalie Gallón. 2018. Puerto Rico says power restoration after Hurricane Maria is complete, but that’s not quite right. *CNN* (14 August 2018). <https://www.cnn.com/2018/08/14/us/puerto-rico-power-restored/index.html>
- [36] Joseph Serna. 2018. County issued conflicting evacuation warnings before deadly Montecito mudslides. *Los Angeles Times* (23 January 2018). <http://www.latimes.com/local/lanow/la-me-ln-montecito-mud-20180123-story.html>
- [37] Yuya Shibuya. 2017. Mining social media for disaster management: Leveraging social media data for community recovery. In *Big Data (Big Data), 2017 IEEE International Conference on*. IEEE, 3111–3118.
- [38] Kate Starbird and Leysia Palen. 2011. "Voluntweeters": Self-organizing by Digital Volunteers in Times of Crisis. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '11)*. ACM, New York, NY, USA, 1071–1080. <https://doi.org/10.1145/1978942.1979102>
- [39] Joe Sterling. 2018. 3 remain missing in deadly California mudslides. *CNN* (15 January 2018). <https://www.cnn.com/2018/01/15/us/southern-california-mudslides/index.html>
- [40] Simon Szreter and Michael Woolcock. 2004. Health by association? Social capital, social theory, and the political economy of public health. *International journal of epidemiology* 33, 4 (2004), 650–667.
- [41] Sebastián Valenzuela, Namsu Park, and Kerk F Kee. 2009. Is there social capital in a social network site?: Facebook use and college students’ life satisfaction, trust, and participation. *Journal of computer-mediated communication* 14, 4 (2009), 875–901.
- [42] Jayant Venkatanathan, Evangelos Karapanos, Vassilis Kostakos, and Jorge Gonçalves. 2012. Network, personality and social capital. In *Proceedings of the 4th Annual ACM Web Science Conference*. ACM, 326–329.
- [43] Sarah Vieweg, Amanda L. Hughes, Kate Starbird, and Leysia Palen. 2010. Microblogging During Two Natural Hazards Events: What Twitter May Contribute to Situational Awareness. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '10)*. ACM, New York, NY, USA, 1079–1088. <https://doi.org/10.1145/1753326.1753486>
- [44] Duncan J Watts and Steven H Strogatz. 1998. Collective dynamics of ‘small-world’ networks. *nature* 393, 6684 (1998), 440.
- [45] Michael Woolcock et al. 2001. The place of social capital in understanding social and economic outcomes. *Canadian journal of policy research* 2, 1 (2001), 11–17.