Adaptation to Exogenous Shocks: Community, Team, and Individual Perspectives on the Social Web

by

Danaja Maldeniya

A dissertation submitted in partial fulfillment
of the requirements for the degree of
Doctor of Philosophy
(Information)
in the University of Michigan
2023

Doctoral Committee:

Associate Professor Daniel M. Romero, Chair
Associate Professor Elizabeth E. Bruch
Professor Yan Chen
Assistant Professor Abigail Z. Jacobs

Danaja Maldeniya
dmal@umich.edu
ORCID iD: 0000-0002-3184-1529

DEDICATION

This dissertation is dedicated to Nirmala, the light of my life, and her parents who took me in as their own.

ACKNOWLEDGMENTS

This dissertation and my seven year journey to its completion was made possible by a community of people that supported me every step of the way. First, I would like to express my gratitude and apprepication for my advisor, Daniel Romero, whose guidance and support has been invaluable in ensuring that I reached the finish line. Under his mentorship, I had the freedom to pursue research that interested me while being challenged to articulate ideas clearly and effectively as well as develop and execute on rigorous study designs to evaluate them effectively. I also had my first teaching experience under Daniel. His everpresent committment to providing students great learning experience was an inspiration and a standard that I hope to emulate. Above all I am grateful for his patience with my idioscyncracies as well as his willingness to always listen and consider suggestions and feedback.

I would also like to thank the members of my dissertation committee, Elizbeth Bruch, Yan Chen, and Abigail Jacobs for their support throughout the dissertation process and the valuable feedback they provided.

Next, I would like to acknowledge my research collaborators without whom many of this projects would not have nearly as successful. In particular, I am grateful to Ceren Budak and Lionel Roberts for their time and support with the research project that to date remains the one I'm most proud of. I had a similarly wonderful experience collaborating with Julie Jiang, under the mentorship of Emilio Fererra and Kristina Lerman, whose passion allowed me to explore a new avenue of research that I would otherwise have not had the time for. Last, but not least, I am thankful to Munmun De Choudhury and David Garcia for their unfaltering support and understanding during the numerous setbacks and rejections in six long years that it took to for our collaboration to reach fruition.

Working on a PhD is a challenging experience. Being thousands of miles from the only home I'd ever known and having left behind my friends and family for an extended period, I, sometimes, found it overwhelming. I owe my mental and physical wellbeing as well as professional accomplishments to the friends I made along the way who supported throughout my time in the PhD program. They provided valuable research feedback and logistical support for my work as well as emotional support when I needed it. A number of them also made a concerted effort to involve me in their social lives which contributed to some of the happiest experiences of my life to date.

In particular, I am grateful to my friends from my Dungeons & Dragons group, rock climbing group, and my travel buddies. In alphabetical order, I thank Alexis Alvarez, Aparna Ananthasubramaniam, Lia Bozarth, Minje Choi, Carolina Chung, Natacha Comandante, Sangmi Kim, Jessie Marie, Abraham Mhaidli, Hao Peng, Edward Platt, Christopher Quarles, Ellen Quarles, Woosuk Seo, Jeremy York, Yulin Yu, and Yixin Zou.

Finally, I want to thank my wife, Nirmala Chandrasiri, who left her career and life back home and has been here with me through all the ups and downs of this journey. You are, forever, the earth beneath my feet and the north star that guides me.

TABLE OF CONTENTS

DEDICATION ii

ACKNOWLEDGMENTS iii

LIST OF FIGURES viii

LIST OF TABLES x

LIST OF APPENDICES xi

ABSTRACT xii

CHAPTER

1 Introduction 1

1.1 Background 2
1.1.1 Quasi-Experimental Design 2
1.1.2 Quasi-Experimental Studies of Online Behavior 4
1.1.3 Adaptation to Exogenous Shocks 6
1.2 Motivation 11
1.3 Overview and Summary of Contributions 12
1.3.1 Social Media Responses of Communities to Natural Disasters follow Prototypical Trajectories 13
1.3.2 Social Media Social Capital Explains Individual Relocation Choice in the Aftermath of Disasters 14
1.3.3 Open Source Developer Communities Centralize Coordination and Decentralize Work in Response to Increased Engagement After an Attention Shock 15

2 Social Media Trajectories of Recovery in Disaster Stricken Communities 18

2.1 Introduction 18
2.2 Background and Related Work 22
2.2.1 Dimensions of Disaster Response 22
2.2.2 Trajectories of Community Response and Recovery 25
2.2.3 Social Media Based Disaster Research 26
2.3 Disaster Dataset 27
2.3.1 Spatiotemporal Boundaries 27

2.3.2 Social Media Data . . . 29
2.4 Methods . . . 30
2.4.1 Length of the Study Period . . . 31
2.4.2 Linguistic Dimensions of Disaster Response . . . 31
2.4.3 Disaster Response Trajectories . . . 32
2.4.4 Clustering Response Trajectories . . . 34
2.5 Results . . . 34
2.5.1 RQ1: Universal Patterns in Disaster Recovery Trajectories . . . 35
2.5.2 RQ2: Variations in Community Disaster Response . . . 39
2.6 Limitations . . . 44
2.7 Discussion . . . 45
2.7.1 Implications Disaster Relief and Recovery Assistance . . . 45
2.7.2 Implications for Disaster Research . . . 47

3 Influence of Online Social Capital on Relocation Choice after Disasters . . . 49

3.1 Introduction . . . 49
3.2 Background . . . 50
3.2.1 Forms of Social Capital . . . 51
3.2.2 Network Robustness . . . 52
3.2.3 Social Capital, Disaster Resilience, and Relocation . . . 53
3.3 Hypotheses . . . 55
3.4 Data . . . 55
3.4.1 Disasters . . . 56
3.4.2 Affected Twitter Users . . . 56
3.4.3 Affected User Ego-networks . . . 57
3.5 Methods . . . 58
3.5.1 Affected Ego Network . . . 58
3.5.2 Measures of Twitter Social Capital . . . 58
3.5.3 Disaster Properties . . . 61
3.5.4 Community Resilience Variables . . . 61
3.5.5 Identifying Relocation . . . 62
3.5.6 Regression Models . . . 63
3.5.7 Robustness Checks . . . 64
3.6 Results . . . 64
3.6.1 Bonding Capital (H1) . . . 64
3.6.2 Bridging Capital (H2) . . . 66
3.6.3 Network Robustness and the Effect of Social Capital (H3) . . . 66
3.6.4 Effects of Disaster Attributes and Community Resilience . . . 67
3.6.5 Geography of Social Capital and Relocation . . . 67
3.6.6 Prevalance of Relationship Types and Relocation . . . 68
3.6.7 Robustness Checks . . . 70
3.7 Limitations . . . 70
3.8 Discussion . . . 71

4 How GitHub Projects Respond to Increased Attention . . . 75

4.1 Introduction . . . 75
4.2 Related Work . . . 76
4.3 Hypotheses Development . . . 78
4.3.1 Community Engagement . . . 78
4.3.2 Growing Pains . . . 79
4.3.3 Adaptation . . . 79
4.4 Background on GitHub and Data . . . 80
4.4.1 GitHub Trending Data . . . 81
4.4.2 GitHub Projects Trace Data . . . 82
4.4.3 Preprocessing . . . 82
4.5 Measurement . . . 83
4.5.1 External Engagement . . . 83
4.5.2 Backlogs and Productivity . . . 84
4.5.3 Member Administrative Orientation . . . 84
4.5.4 Coordination Processes . . . 84
4.6 Methods . . . 86
4.6.1 Propensity Score Matching . . . 87
4.6.2 Hypothesis Evaluation . . . 91
4.7 Results . . . 91
4.7.1 H1: External Engagement and Contribution . . . 92
4.7.2 H2: Backlogs . . . 92
4.7.3 H3: Member Administrative Orientation . . . 93
4.7.4 H4: Coordination Processes and Decentralization . . . 93
4.8 Limitations . . . 97
4.9 Discussion . . . 98
5 Conclusions . . . 99
5.1 Summary of Contributions . . . 99
5.1.1 Overview . . . 99
5.1.2 Adaptaton to Exogenous Shocks . . . 99
5.2 Future Directions . . . 101
5.2.1 Fidelity of ”Social Web” observations as reflections of offline behavior . . 102
5.2.2 Resilience of adaptive responses . . . 103

APPENDICES . . . 104

BIBLIOGRAPHY . . . 128

LIST OF FIGURES

FIGURE

2.1 Histogram of the size of the disaster in terms of the number of associated Twitter users 30

2.2 Data volume maps in the contiguous US. Left: Number of incidents per state in our dataset. Right: Aggregated number of users appearing in all disasters in each state . . 30

2.3 Scale and shape difference for *anger* between clusters 1 and 2. The difference in scale is measured as the absolute difference in the root mean sum of squares of magnitudes from the average trajectories of the two clusters. The shape contribution is measured as the DTW distance between magnitude-normalized average trajectories 36

2.4 Scale vs. shape contributions to differences in dimensions. The values in both axes are scaled to the range of [0,1] 37

2.5 Dimensions with universal trajectories across disasters in the dataset. All figures in a column have the same intensity (Y) scale and the shaded regions are 95% confidence intervals 40

2.6 Categories with noticeable differences across the two clusters. All figures in a column have the same intensity (Y) scale. The shaded regions are 95% confidence intervals . . 41

3.1 Co-variate Effect Sizes on the Likelihood of Relocation after a Disaster (**M5**). The error bars show 95% confidence intervals. The stars show significance (*** p-val $<$ 0.001, ** p-val $<$ 0.01, * p-val $<$ 0.05). 74

4.1 Propensity Score Distributions. (a). Shocks vs. All Non Shocked. (b). Shocks vs. Matches 88

4.3 External Engagement (H1). The bars show the effect sizes for the Diff-in-Diff regression, which represents the effect of the shock on each measure relative to the control group as a percentage, as well as the difference in medians before and after the shock for shocked repositories only. The error bars show 95% confidence intervals. The stars show significance (*** p-val $<$ 0.001, ** p-val $<$ 0.01, * p-val $<$ 0.05). Other bar plots in this section follow the same format. 94

4.4 Backlogs (H2) 95

4.5 Member Admin. Orientation (H3) 95

4.6 Coordination Processes and Decentralization (H4). Missing bars indicate that all values for the corresponding measurement are zero. 96

A.1 Histogram of the number of lexicon dimensions in terms of the percentage of disasters for which we cannot reject the hypothesis of normality ($\alpha = 0.05$) 104

A.2 Silhuette coefficient, Dunn Index, and Connectivity for $N \in 1:10$. Higher Silhuette coefficient and Dunn index values, as well as lower Connectivity, indicate higher quality clustering . . . 105
A.3 Top 50 overall proportional wordshifts for **Risk** . . . 106
A.4 Top 50 overall proportional wordshifts for **Achieve** . . . 107
A.5 Top 50 overall proportional wordshifts for **Reward** . . . 108
A.6 Top 50 overall proportional wordshifts for **Social** . . . 109
A.7 Top 50 overall proportional wordshifts for **Health** . . . 110
A.8 Top 50 overall proportional wordshifts for **Money** . . . 111
A.9 Top 50 overall proportional wordshifts for **Work** . . . 112
A.10 Top 50 overall proportional wordshifts for **Home** . . . 113
A.11 Top 50 overall proportional wordshifts for **Ingest** . . . 114
A.12 Proportional wordshifts for **Sad** for the two clusters of disasters . . . 115
A.13 Proportional wordshifts for **Anger** for the two clusters of disasters . . . 116
A.14 Proportional wordshifts for **Fear** for the two clusters of disasters . . . 117
A.15 Comparison of forecast performance between the predicted cluster trajectory and the baseline forecasting method. The heights of the bars indicate the percentage of times each method provided a better forecast of community response to a particular event for different lengths of observed segments . . . 123
A.16 Three-way comparison of forecasting performance across the predicted cluster trajectory, the alternative cluster trajectory (i.e., the one that was not chosen), and the baseline forecasting method. The heights of the bars indicate the percentage of times each method provided a better forecast of the community response to a particular event for different lengths of observed segments . . . 123
A.17 Community response patterns using shuffled observations for dimensions with universal trajectories across disasters in the original data. All figures in a column have the same intensity (Y) scales which are identical to those of Figure 2.5. The shaded areas are 95% bootstrap confidence intervals. If a confidence interval is not visible for a particular dimension, it is narrower than width of the line showing the mean trajectory. 125

B.1 Effects of social capital on post-disaster relocation across different measurements of relocation (boolean, log distance, and log rank distance) over different measurement periods (12, 24, and 36 weeks). For boolean measurement of relocation (i.e., **Relocated**), the units of the effects is change in odds, while for log distance and log rank distance the units are corresponding magnitudes (i.e., meters (log), change in rank respectively). The error bars show 95% confidence intervals. The stars show significance (*** p-val < 0.001, ** p-val < 0.01, * p-val < 0.05). . . . 127

LIST OF TABLES

TABLE

2.2 Selected Lexicon Dimensions. All dimensions except Fear, for which used the NRC lexicon, are sourced from LIWC 33
2.3 Distribution of events by the type of hazard in cluster 1 and cluster 2 35
2.4 Odds ratios by covariate from logistic regressions predicting the cluster associated with a disaster ($^*p < 0.05$) 42

3.1 Disaster Statistics 56
3.2 Community resilience indicators from the FEMA Resilience Analysis and Planning Program (RAPT) used as controls 62
3.3 Effects of Social Capital on Relocation Choice. Values for all covariates except for the intercept are percentage changes from the base rate of relocation (i.e., Intercept). The stars show statistical significance (*** p-val < 0.001, ** p-val < 0.01, * p-val < 0.05). 65
3.4 Effect the prevalence of different relationship types on relocation choice during the first 12 weeks after the onset of disaster. Values are the changes in likelihood of relocation for a unit (SD) increase in the prevalence of each relationship type. The stars show statistical significance (*** p-val < 0.001, ** p-val < 0.01, * p-val < 0.05). 70

4.1 Measurement of GitHub project behavior used to test our hypotheses described in section 4.5 85
4.2 Propensity Score Block Statistics 90

B.1 Number of people who newly relocated from or returned to their community during different post-disaster periods 126

LIST OF APPENDICES

A Supplemental Materials for Chapter 2 . . . . **104**

B Supplemental Materials for Chapter 4 . . . . **126**

ABSTRACT

The advent of the "Social Web" has provided a rich testbed for studying numerous human behavioral phenomena at scale for social scientists. Data incidentally collected by online social systems and platforms as people use them to work, connect with others, express opinions and emotions, etc., are temporally fine-grained, rich in detail (e.g., networks of interaction, language, etc.), and importantly relatively easy to access for the motivated researcher. However, as with other, more traditional forms of observational data, online digital traces are generally not amenable to studying causal effects. In this dissertation, we use quasi-experimental study designs which can extract causal narratives from observational data under the right conditions, in order to study how people respond to challenges brought on by external shocks. Historically, the wide-use of quasi-experimental methods have been hampered by the limited availability of suitable data, a limitation that is alleviated by the affordances of online digital traces.

It is challenging to use experiments to study how people respond to external shocks as well as how different intrinsic and environmental factors influence these adaptations for ethical and practical reasons. This dissertation demonstrates how, in this vacuum, quasi-experimental studies using online data can make meaningful contributions to the causal understanding of this domain. It also highlights the limits of this approach, especially in the context of extending inferences to offline contexts, and discusses potential strategies for how these issues may be mitigated with additional work. More specifically, we study individual and collective adaptive responses to unexpected challenges in two different settings, natural disasters and attention shocks in online crowd environments. In the context of natural disasters, we use novel computational methods to characterize the response of the community in the aftermath and demonstrate that community responses follow several prototypical patterns in different disasters and the intrinsic properties of disasters can explain some differences in community responses. In the same setting, we investigate how different forms of an individual's online social capital can influence their choice to move away from the affected area and consider their differing implications for individuals and communities. Finally, in a markedly different setting, we investigate how online peer-production teams manage an influx of users and contributors after receiving an attention shock with implications for the sustainability and growth of such virtual teams. In sum, we demonstrate the potential for this data and methodology combination to infer fine-grained and nuanced narratives regarding how people

adapt to unforeseen challenges and draw attention to their implications individuals and collectives vulnerable to shocks as well as to policy makers interested in building resilience in communities.

CHAPTER 1

Introduction

Over the last few decades, the World Wide Web has grown and evolved rapidly to the point that today a large fraction of the world population uses some form of it across many spheres of life such as for work, entertainment, shopping, information search, socializing, etc. In particular, the ”Social Web” elements that began to be introduced at the start of the century have revolutionazed the ways in which people share their emotions, perspectives, and other thoughts while engaging with others on the Web. Incidentally, these innovations and their widespread adoption have created an opportunity for social scientists to study a wide range of behavioral phenonomena. Organizations that operate the many ”Social Web” platforms and systems (e.g., social media, peer-production platforms, commenting tools, etc.) retain records (or digital traces) of the most, if not all, actions of people using them that are uniquely attributed to their online persona. Overtime, this practice produces rich behavioral datasets for large populations, which have been used to great effect in numerous studies of behavioral phenomena under the umbrella of the burgeoning ”Computaional Social Science” field (Lazer et al., 2009; Edelmann et al., 2020).

However, these “Social Web” digital trace data are not a panacea for all challenges faced by social scientists when studying human behavior. These data present a number of challenges that limit their use and that need to be carefully managed to extract robust insights. One of the most significant such limitations is that Social Web data are a form of **observational data**. This means that unlike experimental data, which include counterfactual measurements by design, it is more challenging to use such digital traces to investigate causal relationships associated with the captured human behaviors. Prior work has developed different methods for addressing or at least mitigating this limitation, which are generally known as quasi-experimental designs[1] (Remler and G., 2014).

We combine quasi-experimental study design with large-scale trace data from the Social Web, spanning different populations and settings, to study how people, individually and collectively,

[1]Some prior work makes a distinction between natural experiments and quasi experiments while others consider natural experiments under the umbrella of quasi-experimental design. In the former, natural and quasi experiments are distinguished by the fact that the former rely on naturally occurring unplanned treatments and the latter are based on consciously applied (non-random) treatments. In this dissertation, we consider natural experiments as a form of quasi-experimental design.

respond to the stress and challenges arising from exogenous shocks, as well as the factors, environmental and intrinsic to the focal individual or population, that influence those responses. Due to the dramatic nature of the challenges posed by many types of exogenous shocks of interest (e.g., natural disasters, economic crises), it is ethically and practically infeasible to study these questions experimentally, which makes quasi-experimental designs the next best alternative for developing causal narratives in this space. This strategy is particularly suited to studying these phenomena for two reasons. First, prior work has established the effectiveness of quasi-experimental designs in studying these phenomena. However, until recently, these studies relied on traditionally collected observational data (e.g., surveys, interviews, etc.), which has limited the number of such studies and their coverage of the problems of interest in the domain. In general, only a small fraction of observational data collected by traditional means (for any purpose) incidentally captures changes in behaviors of interest due to exogenous shocks. In contrast, digital trace data of people's behavior over the Social Web can be leveraged relatively easily to explore a range of behavioral dimensions (e.g., emotional, social, mobility, etc.). Further, due to the continuous nature of data capture on the Web, people's responses to shocks can be studied simultaneously across many similar events.

In conclusion, this dissertation makes the following high-level contributions. First, individual studies contribute new behavioral findings to the domain of individual and collective adaptation to shocks in two different settings (natural disasters and online attention shocks). Second, as a whole, it demonstrates the potential of utilizing Social Web data in combination with quasi-experimental study design to advance our causal understanding of improtant behavioral phenomena, as well as establish more nuanced narratives.

1.1 Background

1.1.1 Quasi-Experimental Design

Data from randomized experiments are the gold standard for inferring causal relationships in social science. This is because scientists collect measurements of behavior affected by a consciously applied random treatment, as well as counterfactual measurements that can be used to isolate causal effects of the treatment from the former. On the other hand, observational data do not include consciously collected counterfactual measures, making it more challenging to infer causal relationships.

Quasi-Experimental design (QED) refers to a set of methodologies that allow scientists to identify causal effects from observational data that meet certain conditions. These methods differ in terms of the type of counterfactual that is estimated, the corresponding requirements that a dataset needs to satisfy, and the resulting limits on the credibility of causal inferences (Remler and G., 2014; Dunning, 2012; Cook et al., 2002). In the following sections, we briefly discuss, non-exhaustively, three broad categories of quasi-experimental design that use different strategies for estimating counterfactuals.

One Group Pretest-Posttest Design (Pre-Post Comparison). One of the simplest forms of quasi-experimental design is to measure the behavior of a population at two points in time such that one measurement is before the event that serves as the treatment and the other after it. The difference between two measurements is inferred to have been caused by treatment (Cook et al., 2002; Remler and G., 2014). This design is intuitive and due to its simplicity can be easily scaled when data is abundant as in the case of traces from the Social Web. Variations of regression discontinuity analysis adapted for within-subject analysis are frequently used to statistically estimate causal effects in pre-post comparison studies (Hill and Shaw, 2021; Wertis et al., 2023; Craig et al., 2017; Oktay et al., 2010).

The main weakness of this design is that, in estimating the causal effect of treatment, it assumes that no other environmental factors or events could have influenced the observed behaviors (Cook et al., 2002). Since observational data are collected from real environments (as opposed to controlled experimental settings), this is a strong assumption. One possible strategy to mitigate this limitation is to carefully consider potential confounding variables and include them in the design. When the data are biased with regard to these covariates, causal interpretations should be correspondingly conservative. A more sophisticated version of the pre-post comparison design, Interrupted Time Series Design, uses multiple measurements before and after the treatment event (Zhang et al., 2017b; Craig et al., 2017; Remler and G., 2014). It takes into account any trend in the measured behavior prior to treatment that can improve the estimate of the causal effect.

Post-Only Comparison with Nonequivalent Groups. In the absence of data prior to treatment, it is possible to estimate the effect of treatment on the treatment group by comparing the post-treatment data for them with data from the same period for a control group (Cook et al., 2002). This approach assumes that the post-treatment behavior of the control group is equivalent to the unobserved pre-treatment behavior of the treated group. A suitable control group is found by, first, identifying any attributes that may influence the effect of the treatment and then finding a group of untreated individuals in the data who have a similar distribution of values for these attributes. This strategy to find a control group is known as matching. Propensity score matching a popular sophisticated variant of this strategy in which the comparability of untreated individuals with treated individuals can be estimated simultaneously across multiple attributes (Imbens and Rubin, 2015; Cook et al., 2002; Choi et al., 2023).

Pre-Post Study with a Comparison Group. The concerns regarding the internal validity of pre-post comparison and post-only comparison with a control group can be mitigated by combining these designs into a more sophisticated strategy where pre-treatment and post-treatment measuremen are found for both the population that was treated and an appropriate comparison group. Then the pre-post comparison of the behaviors in the control group can be reasonably expected to account for both any trends in behavior common to both groups and other unobserved events (e.g. weather effects on behavior) that may have influenced the outcome being measured (Cook et al., 2002; Dunning, 2012; Card and Krueger, 1993). This approach is also known as the Difference-in-Difference (DiD) design. In studies that employ a DiD design, matching strategies such as propensity score-based matching are frequently used to ensure that the comparison group is as similar to the treated population as possible (Craig et al., 2017; Cook et al., 2002; Choi et al., 2023). The basic variant of this quasi-experimental design uses only one pretest and one posttest measurement each for the control and treatment groups. Therefore, in estimating the causal effect as the difference in the difference between posttest and pretest measuremnts across the groups, it assumes that the trend in behavior in the absence of the treatment would be identical in the two groups. This assumption can be explicitly tested by getting multiple pretest measurements for each group (Craig et al., 2017).

1.1.2 Quasi-Experimental Studies of Online Behavior

Recently, previous work has used quasi-experimental designs to study a range of different behavioral phenomena using "Social Web" data as well as other forms of digital trace data (Oktay et al., 2010; Zignani et al., 2014; Hill and Shaw, 2021; Choi et al., 2023; Zagheni et al., 2014; Wertis et al., 2023; Phan and Airoldi, 2015). This body of existing studies serves as an appropriate starting point to consider the utility of different study designs for investigating different behavioral phenomena, the potential of this combination of data and methodology to advance empirical computational social science, as well as the inherent limits in their use for inferring causal narratives.

Prior quasi-experimental studies of online behavior can be broadly categorized into two groups based on the context in which their contributions are positioned. One set of studies strictly focus on understanding online behavior and their interactions with the systems and platforms that make up the *Web* (Oktay et al., 2010; Hill and Shaw, 2021; Zhang et al., 2019). Despite making narrow claims regarding online behavior that may or may not generalize to offline contexts, these studies can make important contributions to the relevant behavioral domains and the field of computational social science due to two reasons. The first is that the virtual space of the Web has become an important setting, complementary to the physical world, where significant human dynamics take place, such as identity formation, communication, social interaction, etc. (Manago et al., 2012;

Kogan and Palen, 2018). Therefore, studying the aspects of these dynamics that play out on the Social Web are vauable in and of themselves, and this is particularly valid given the ubiquitous use of the Web by societies around the world. Second, there are now some types of activities and behaviors, such as online crowd work, that are native to the Web made possible by its unique affordances. In these cases, the corresponding trace data, in conjunction with quasi-experimental strategies, are ideal for studying causal effects (Hill and Shaw, 2021; Oktay et al., 2010; Zhang et al., 2017b; Malik and Pfeffer, 2016). Chapter 4 of this dissertation discusses a study of how online crowds adapt to sudden popularity that falls within this latter category.

The second line of studies that use quasi-experimental strategies with data from the Social Web consider the implications more generally to both online and offline settings (Kogan and Palen, 2018; Kogan et al., 2015; Metaxa-Kakavouli et al., 2018; Garcia and Rimé, 2019; Saha et al., 2020, 2019). Chapters 2 and 3 of this dissertation that discuss collective and individual adaptations to natural disasters fall into this category. The key challenge to the validity of these types of study, especially in the case of offline settings, is the varying degrees of selection bias associated with digital traces collected from different parts of the Social Web (Diaz et al., 2016). As an example, Twitter users, the source of data for a large fraction of this type of work, including the studies discussed in Chapters 2 and 3, are known to be generally younger and wealthier than the general public, as well as overrepesenting minorities such as Black and Hispanic populations (Walker and Matsa, 2021; Center, 2023). Further, along several other more fuzzy dimensions, such as behaviors and emotions, we currently lack a good understanding of whether people who engage with the Social Web are representative of those who do not. This non-representativeness of online data in regards significantly tempers the value of study contributions to the offline setting, especially when considered in isolation. However, there are two main arguments for suggesting that these studies can contribute to the offline setting under specific conditions. The first is that some measurements of online behavior, particularly on a population scale, have been found to be reliable signals of the corresponding offline behavior (Dunbar et al., 2015; Pellert et al., 2022). In general, we can be more confident of the validity of findings from studies that investigate such behaviors using online data to offline settings than studies for which the measured behaviors do not have such validations. Second, online trace data provide an opportunity to conduct large-scale panel studies in which behaviors can be measured multiple times over time. This type of study is generally cost prohibitive to carry out using traditional data collection methods. Furthermore, online trace data makes it possible to construct such behavioral panels to study external events ex-post, whereas it is usually infeasible to ex-ante collect such data using traditional methods (Diaz et al., 2016). Previous work has also proposed and investigated several strategies for mitigating the effects of selection bias when using trace data from the Social Web including using inferred demographic properties and continuous reweighting of observations using relevant dimensions (Diaz et al., 2016), calibrating online observations using strategically collected offline samples (Blumenstock et al., 2015), and collecting complementary data by offline means when possible (Jones et al., 2017).

1.1.3 Adaptation to Exogenous Shocks

Many aspects of human existence at the societal, communal, group, and individual levels follow a prototypical pattern of evolution over time. Relatively long periods of stability or gradual change are punctuated by sudden macroscopic changes that rapidly change the course of individual lives, institutions, and societies (Meyer, 1982). Often, these transformations are the consequence of unforeseen events in the external environment in which the person, community, or society is embedded. Some events, such as scientific discoveries, a new factory, or a windfall, are more likely to be seen as positive and lead to better outcomes for those affected. Other factors, such as economic recessions, natural disasters, or the death of a loved one, are negative and are more likely to change the lives of those affected for the worse. However, both ostensibly negative and positive external shocks share the feature that they serve as a strong impetus or stressor for the affected parties to rapidly adapt to a changed environment.

External shocks that pose dramatic challenges or threats to individuals and societies are not a new phenomenon. Consequently, there is a long history of scientific research that studies shocks and how people adapt to them in different settings (Ziegler, 1970; Worster, 2004; Brown, 1971; Dynes, 1970; Norris et al., 2008; Meyer et al., 2005). Most of this research is distributed across different domains such as disaster response, economics, and bereavement. They have established that there is substantial heterogeneity in how people (individually and collectively) respond to unexpected challenges arising from their environment (Rutter, 2012). Crucially, some individuals and groups are more successful in overcoming the challenges and thriving in their aftermath, while others can suffer major long-term negative effects due to maladaption. Due to the critical nature of possible, positive or negative consequences to affected people arising from these challenges, it is important to develop a comprehensive understanding of (i) the range of ways in which people, as individuals and collectives, respond to external challenges, (ii) the different intrinsic attributes (e.g., genetic, psychological, social, demographic, economic) of affected people that influence their specific response, and (iii) the desirability of outcomes arising from different responses. The following sections summarize current knowledge regarding how people, as individuals and collectively as teams, organizations, communities, etc., respond to the challenges posed by unexpected events.

This dissertation seeks to advance our understanding regarding the first two questions related to the how individuals and collectives respond to unexpected challenges — i.e., (i) what are different ways in which people respond? and (ii) what intrinsic and environmental factors drive the differences in the responses — by utilizing digital traces that record their behavior on the Social Web. The third question , regarding the outcomes that result from different responses and their desirability, is left for future research, which is briefly discussed in the conclusions.

Communities. As a group of people living in the same geographical region, communities collectively face challenges arising from environmental or man-made external shocks. While members of communities face unique individual challenges and respond accordingly, due to their shared identity, as well as social and economic resources, they also collectively process and adapt community resources, institutions, and relationships to the changing environment. Although environmental shocks largely consist of natural disasters, such as hurricanes and floods, and longer-term climate effects (e.g., droughts) (Cherry et al., 2021; Norris et al., 2002; Ursano et al., 2007), there is a wide range of shocks that are caused by human activity, including man-made ecological disasters, acts of violence and terrorism, changes to policies and laws, and entry (and exit) of businesses (Lin et al., 2017; Wicke and Silver, 2009; Morganstein et al., 2016; Zhang et al., 2017a; Verity and Jolley, 2008). In the short term, after an external shock, whether environmental or man-made, communities experience a wide range of negative emotional and cognitive effects. In the event of disasters, communities experience a surge of fear and anxiety due to hazards that actively threaten lives and property, sadness for the destruction already wrought, as well as anger and frustration (Bonanno et al., 2010; DeWolfe, 2000). Communities that experience other types of shock, including acts of violence and terrorism and the loss of a major employer, also experience similar emotional and cognitive effects (Garcia and Rimé, 2019; Cohn et al., 2004; Verity and Jolley, 2008; Blau, 2006). Parallel to this emotional transformation, often community members increasingly identify and organize around their shared identity to make sense of changes in their environment, weather physical, social, and psychological impacts, and coordinate a response (DeWolfe, 2000; Wicke and Silver, 2009; Lin and Margolin, 2014; Cohn et al., 2004). This increase in social cohesion helps communities, at least initially, remain responsive to the challenges (and opportunities) introduced by the shock. As time passes, if the negative effects of a shock persist and the community lacks physical and social resources, the communities continue to engage collectively in making sense of their new reality, which may be accompanied by a sense of anxiety about the long-term effects of the event (Doré et al., 2015; DeWolfe, 2000; Bonanno et al., 2010). The effects of shocks are also felt in the social networks of communities, and their collective adaptations to the challenges and opportunities posed by shocks transform them further. In the immediate aftermath of a shock such as disasters or the loss of a major employer, the social fabric of a community is significantly disrupted and whole sections, representing victims and their relationships, may be lost (Bonanno et al., 2010; Wicke and Silver, 2009; Gordon, 2004). Regardless of whether the shock directly disrupted the social network of the community, a new configuration emerges during the subsequent period of increased social identity as a collective response of the community emerges (Bagrow et al., 2011; Kogan et al., 2015; Misra et al., 2017; Islam and Walkerden, 2014). Social networks become more dense and inward-looking, with community members interacting more often with each other to share information, provide social support, and collectively assess the effects of shock (Bagrow et al., 2011; Kogan et al., 2015). Over time, they can continue to evolve with individuals or organizations that play different roles within the community becoming more central (Misra et al., 2017; Consoer and Milman, 2016). For example, in the case of disasters or acts of violence or terrorism, information about rapidly evolving circumstances is often communicated throughout the community through weak *bridging* relationships where local institutions (both community organizations and individuals) play an important role (Consoer and Milman, 2016; Misra et al., 2017)

Industries and Economic Networks. Industries and economic networks are composed of large populations of economic agents (e.g. organizations, executives). Therefore, they are similar to social communities in the sense that their behavior emerges from the actions and interactions of many people and entities. However, in general, they possess more cohesive shared identities and collective goals, although not to the extent of smaller units such as a single organization, family, or individual. As a consequence, they are capable of more effective self-organized collective responses to external shocks than communities.

Given that organizational interactions often tend to be a matter of public record, emergent collective adaptations of these networks have been more amenable to be studied in greater detail. Industries normally operate in a state of dynamic equilibrium under which the overall structure of collaborative and competitive interactions remains stable. A wide range of exogenous shocks, such as changes in market prices (Romero et al., 2016), new policies and laws (Meyer et al., 1990; Smith, 2020), and technological innovations (Schilling, 2015) can disrupt these equilibria to varying degrees. In the face of looming uncertainty that poses a threat to the status quo, but may also offer new opportunities, economic agents resort to exchanging information and collaborating more actively, forming a more dense network. For example, after the 9/11 terrorist attacks, US airlines sought more actively alliances and adopted new business practices to reduce costs and manage risk (Corbo et al., 2016). Similar patterns of increased integration have been observed among stock brokers after a market shock (Romero et al., 2016), US hospitals, and the broader health industry after major changes to health policy (Meyer et al., 1990), organized crime in Chicago during Prohibition (Smith, 2020), and the technology industry with the advent of the Internet (Schilling, 2015). In addition to increased overall connectivity, reorganization in response to external shocks tends to consolidate power, resources, and interactions within industries and economic networks, leading to a more core-peripheral structure (Corbo et al., 2016; Schilling, 2015). Although this transformation likely makes industries more efficient and resilient collectively, some actors may become more vulnerable. For example, in the case of the transformation of organized crime in Chicago in response to Prohibition, vulnerable actors, such as women, were further marginalized (Smith, 2020).

Finally, it is possible that these behaviors are temporary adaptations to manage risks and opportunities that arise from shocks rather than permanent changes. In the case of the explosion of alliances of technology companies during the advent of the web, the network had largely returned to pre-shock levels of interaction within a few years (Schilling, 2015).

Teams and Organizations. Organizations and teams, which are made up of small to large numbers of members, share several behavioral dynamics with larger collectives such as communities and industries. However, in contrast to these large-scale social and economic units, they also possess cohesive collective identities and are embedded in a network of other actors. As a consequence, teams and organizations share certain behavioral dynamics with individuals. This means that organization and team response and adaptation to external shocks share similarities with both community/industry level and individual adaptations. Online peer production communities, such as Wikipedia and open source development, are a recent phenomenon that adds further variation and nuance to shock adaptation on this scale (Zhang et al., 2017a). These decentralized social organizations differ from traditional organizations and teams by having porous boundaries, members with motives other than financial ones, flexible and emergent coordination, etc. (Benkler, 2006; Vasilescu, 2014).

Organizations and teams can experience challenges from a range of external shocks including impacts of industry or economy-wide crises (Corbo et al., 2016), loss of external partners (e.g., customer, supplier) (Pálovics et al., 2021), new technology (Schilling, 2015) or government regulations (Meyer et al., 2005), labor action (Meyer, 1982), and loss of workers/members (Zhang et al., 2017b). When experiencing these crises, organizations and teams tend to centralize coordination and production to its core members (Meyer, 1982; Zhang et al., 2017b,a; Edgerton et al., 2022; Hambrick and Crozier, 1985). Although centralization of power and coordination to key individuals, who are most familiar with the internal workings of the organization as well as their environment, is common, production is only centralized if the organizational boundary is not porous, i.e. if the organization is not able to recruit capacity to deal with the increased complexity of the situation (Hambrick and Crozier, 1985; Meyer, 1982; Zhang et al., 2019; Schilling, 2015). If recruitment is viable, organizations are likely to adopt a centralized coordination mechanism that drives decentralized production. For example, in organizations that experienced rapid growth(Hambrick and Crozier, 1985) or were faced with disruptive new regulations (Meyer et al., 2005), executives or experienced employees intensified communication with each other and the increasingly decentralized workforce. Similar observations have been made regarding peer production in communities such as Wikipedia. When a large fraction of editors from mainland China were unable to connect to Wikipedia due to censorship, coordination and editing became more centralized within the network of the remaining editors of Chinese Wikipedia in the short term (Zhang et al., 2017a). In contrast, when breaking news induced a sudden spike of attention on a Wikipedia page, a scenario in which an artificial entry barrier is not in place, new editors made significant contributions to updating the page. At the same time, existing experienced editors played an important role in vetting content and directing the new editors (Zhang et al., 2019).

The size of the organization also plays an important role in how well they are able to adapt to different types of shocks. In general, small organizations, which are more dynamic, are able to reconfigure their internal processes and external relationships more easily (Williams et al., 2017; Miklian and Hoelscher, 2022). This is more likely to make small organizations more responsive to shocks that do not result in immediate loss of resources or personnel. In the case of shocks that result in attrition, larger organizations that have greater reserves of both material and human resources are more successful in sustaining themselves in the short term, if not indefinitely (Zhang et al., 2017a; Meyer et al., 2005).

Individual. People can experience a wide range of stressful life events as a consequence of external shocks. These may be uniquely personal, such as a debilitating injury (Bonanno et al., 2012), loss of a loved one (Hobbs and Burke, 2017), loss of employment (Morris and Irwin, 1992; Burke and Kraut, 2013), or a romantic breakup (Garimella et al., 2014) or a part of a collectively experienced event such as a disaster (Bonanno et al., 2010) or an epidemic (Saha et al., 2020; Bonanno et al., 2008). People exposed to these stressful events can experience declines in their subjective well-being and develop physical, social, and psychological symptoms, including anxiety, depression, PTSD, and psychosomatic and stress-related health problems (Bonanno et al., 2010; Choi et al., 2023). Some effects of these events, such as loss of property or employment combined with psychological maladaptive behaviors, can also lead to long-term economic instability (Shirani and Henwood, 2011).

Not everyone affected by unexpected shocks demonstrate the same adaptations. Some people recover from negative effects faster, while others show no symptoms or demonstrate positive adaptations. In fact, across the range of possible life events, resilient adaptations appear to be common (Bonanno, 2004; Bonanno et al., 2010). These variations in how people adapt to shocks have been associated with access to intrinsic psychological attributes, economic resources, and social support. In particular, when individual capabilities are inadequate, which is common in the case of many of these life events, people rely on their social circle to provide material, emotional and informational support (De Choudhury and Kiciman, 2017; Verity and Jolley, 2008; Panday et al., 2021). Individuals with more social capital generally receive more support from their social network in the aftermath of the shocks and such support improves their chances of positive adaptation and quick recovery (Bonanno et al., 2010).

The type of support that individuals receive from different parts of their social network depends on the nature of the relationships. A person's family or close friends are more likely to provide immediate material aid and emotional support in the event of a shock, while weak or distant ties tend to be useful sources of information (Aldrich, 2012a; Burt, 2007). Additionally, prior work has shown that a person's extended social network and the social capital represented therein can play a crucial role in the long term when the resources of their immediate social circle have been exhausted (Cong et al., 2018). This is especially true in the case of collective shocks where a person and their immediate social circle can be affected simultaneously.

1.2 Motivation

Human society is increasingly characterized by coupled social, cultural, economic, infrastructure and technological systems that are also interconnected with global ecology. While humans, both individually and collectively, are no strangers to exogenous environmental shocks, the strong coupling among the networks (both natural and human) that govern our existence and behavior mean that shocks have become more frequent, more easily cascade across different systems, and become more disruptive in the process (Piquer-Rodríguez et al., 2023; Namatame and Komatsu, 2011; Trump et al., 2017; Kinzig et al., 2006; Yabe et al., 2022). In this new regime, top-down after the fact state interventions are less effective in ensuring that affected people and systems adapt effectively to the adverse effects of shocks (Trump et al., 2017; Yabe et al., 2022; Gaupp, 2020). In this vacuum, the observation that people, at different social levels, are often capable of successfully self-organizing in response to shocks has led to a renewed interest in the phenomenon of individual and collective resilience (Trump et al., 2017; Bonanno, 2004; Aldrich, 2012a). As a consequence, states have begun to adopt more proactive policies towards managing disruptive events through fostering the capacity for resilience at different levels of society. In this context, it has become crucial to develop a comprehensive understanding of the patterns of emergent human adaptation to shocks, the comparative desirability of the outcomes of these behaviors, as well as the antecedent resources and processes that generated them (Trump et al., 2017).

The goal of this dissertation is to investigate how people, individually and collectively, adapt to the challenges posed by unexpected environmental shocks across different settings with a particular focus on the social interactions and behaviors that emerge from them. More specifically, I consider the behaviors of communities, online peer production teams, and individuals across two types of shocks; (i) natural disasters (communities and individuals), and (ii) attention shocks (peer production teams). I combine an quasi-experimental study design with digital behavioral traces corresponding to many instances of a single type of external shock to infer causal relationships linking unexpected adversities to changes in people's behavior and those behavioral changes to their pre-existing demographic and social attributes. This strategy simultaneously resolves the limitations traditional observational studies (i.e., inability to make causal claims) and the challenge of identifying appropriate behavioral data for utilizing quasi-experimental designs at scale.

The use of digital traces produced by people on the Social Web has inherent benefits as well as complications in our context. As a medium that has become equally or more important as in-person interaction for important human dynamics, including social engagement, political communication, and working, increasingly, people's identities are constructed, and behaviors are enacted at least partially through their online presence. Therefore, while relying on relevant knowledge from the offline context both before and after the advent of the Web, we maintain the running theme of distinguishing observations made on the Social Web from offline behavior and discussing similarities and differences where relevant. The availability of fine-grained behavioral data across temporal, geographical, and other contextual dimensions (e.g., labeled task in a production setting in Chapter 3) allows our analyses to be based on nuanced behavioral constructs whose observations are evaluated for robustness across those dimensions.

The most significant limitation of using trace data generated on the web in analysis of sociological phenomena is the unavailability of numerous socio-demographic attributes associated with people and their relationships that are often collected in more traditional work. The work discussed in this dissertation is no exception in this regard. In this dissertation, where possible, we have mitigated these effects. In Study 1, our approach retains the variation in adaptive behavior, some of which is likely produced by variation in the distribution of unobserved socio-demographic properties, of groups of people affected by different shocks, but stops short of linking those variations to socio-demographic variables. In comparison, we control for observable aggregate variables that are correlated with some of these elements in Study 2, and in Study 3, compare the adaptive behavior of organizations experiencing shocks with others that have similar observable properties.

1.3 Overview and Summary of Contributions

This dissertation has the following overall structure. In Chapter 2, we use the tweets of over 2 million people located in communities affected by natural disasters to investigate the collective experience and adaptation of communities across a range of different events. In particular, we focus on identifying prototypical patterns in community response over social media that generalize across different communities and events. Then, in Chapter 3, we once again use Twitter data to understand how a person's online social capital is associated with their relocation decision in the aftermath of a disaster. Finally, in Chapter 4, we use contributor activity traces on GitHub to investigate how developer communities adapt their coordination and production patterns to increased outside engagement due to sudden popularity.

1.3.1 Social Media Responses of Communities to Natural Disasters follow Prototypical Trajectories

Prior work has shown that communities respond in different ways to physical and psychological trauma and damage to the social fabric caused by disasters (Tierney et al., 2001). Some communities quickly recover from the challenges posed by a disaster, while others experience lasting consequences on individual, social, and economic fronts (Bonanno et al., 2010). Further, beyond a community's intrinsic resilience, certain attributes of the hazard, such as the level of forewarning, its duration, severity and nature of damage, may influence the initial community response and the nature of its long-term recovery (Council, 2006; Norris et al., 2008). Prior work which investigated community response to disasters has explored different dimensions of interest, such as affect, social dynamics, and cognitive processing (Bonanno et al., 2010; Norris et al., 2008). However, few, if any studies have investigated community response through a holistic thematic framework that encompasses most, if not all, experiential and behavioral dimensions of relevance simultaneously. Further, prior work, other than meta-analyses, studying community response has been limited to a single event or a few of them (Lin and Margolin, 2014; Lin and Cromley, 2018; Olteanu et al., 2015).

In Chapter 2, we utilize social media posts (Tweets) of individuals located in communities affected by over 200 disasters of different types that affected the US to study community response in the aftermath of disasters at scale. We use computational linguistic analysis to extract aggregate multidimensional experiential and behavioral trajectories for these communities. We then use time series clustering to identify groups of communities characterized by similar responses, as well as dimensions along which community behavior is generalized across events.

Our results show that the response of disaster-stricken communities observed over social media follows prototypical trajectories along a number of dimensions. Early community response is characterized by a heightened sense of uncertainty and risk perception that is accompanied by a decrease in perceived productivity. While communities show signs of increased activity and coordination associated with disaster response, they show less engagement with spheres of life normally associated with productivity (e.g., school, work, sports). The results also show that the social spheres of community members contract to their most important relationships (e.g., family) in the immediate aftermath of a disaster. Overall social engagement gradually recovers over the next few weeks. We also observe that among dimensions that describe biological and physical processes within communities (e.g., health, work, finances, nutrition), nutrition related processes are much slower to recover and that this decline corresponds to high-level food needs and activities (e.g., satisfaction, socialization, restaurants) as opposed to basic sustenance.

In addition to the existence of prototypical response trajectories, we show that around a quarter of communities have a comparatively more challenging time adapting to the aftermath of a disaster. These communities show much higher levels of negative emotion (e.g., fear, anger, anxiety, sadness) over longer periods. In particular, fear and anxiety show a pattern of periodic recurrence. Overtime, these communities increasingly engage in cognitive-processing associated with sense-making, suggesting that they struggle to make sense of and accept post-disaster reality. Finally, we show that communities that demonstrate these more mal-adaptive responses are associated with more severe events.

Our results have implications for several aspects of disaster relief and recovery assistance which we cover in detail in Chapter 2. We suggest that our methodology, as well as the observed prototypical response trajectories in disaster-stricken communities, can inform strategic-level disaster planning, in addition to monitoring and forecasting specific community responses. Additionally, we note that some aspects of community response on social media are characterized by an absence of information (e.g., social engagement, nourishment). Given the primacy of social media as a source of news for the general public, this observation suggests that public attention, and grass-roots disaster assistance, which may rely on that attention, may become misaligned with community needs.

1.3.2 Social Media Social Capital Explains Individual Relocation Choice in the Aftermath of Disasters

Disasters can cause tremendous disruptions in the lives of individuals and families in the form of physical and psychological trauma, loss of property, loss of livelihood, etc. (Bonanno et al., 2010). After a disaster, some affected people choose to remain in their communities and rebuild, while others relocate to greener pastures (Lee et al., 2017). Prior work suggests that this choice may be influenced by the physical and social resources available to the community at large, as well as the individual's own resources (Karunarathne and Lee, 2019; Engel and Ibáñez, 2007; Lee et al., 2017). While prior work has extensively investigated the influence of community-level resources on both individual and aggregate measures of relocation, less attention has been paid to the influence of a person's own social capital derived from their social circle. Theoretical frameworks as well as a limited number of studies focusing on individual events and using simple measurements have suggested that a person's social capital can, depending on the specific form, either improve their ability to cope and recover in-situ or alternatively provide the resources necessary to make a smooth transition to a different locale (Aldrich, 2012a; Norris et al., 2008; Casagrande et al., 2015).

In Chapter 3, we study the relationships between two different forms of social capital, bonding, and bridging, of over 16,000 people located in communities corresponding to over a hundred disaster declarations by the Federal Emergency Management Agency (FEMA) to their relocation by using their tweets before and after a disaster. We use established geo-spatial analysis techniques to identify the specific pre- and post-disaster home locations for these individuals using geographical information in their geotagged tweets. Additionally, for each individual, we extracted their pre-disaster ego-network based on observed interactions and estimated multiple constructs, capturing different mechanisms, for each type of social capital. We analyze the association of these different social capital constructs with the likelihood of a person having relocated while controlling for a large number of community-level resilience and disaster-related attributes.

Our results show that different forms of social capital, measured from a person's social media, are indeed associated with their relocation choice. Measures for bonding capital, which correspond to close relationships that are embedded in the person's immediate social circle, show contradictory results regarding their influence. First, the level of overall embedding of the person's network decreased their odds of relocating after disaster. In contrast, a measure based on the retention of strong ties had a consistent positive effect on relocation. We suggest that this observation, in particular, reveals the complexity and nuance of the mechanisms linking social capital to relocation, including the competing mechanisms within the same broad class of capital (Engel and Ibáñez, 2007). Measures of bridging capital, which correspond to weaker and less embedded parts of a person's network, are positively associated with the odds of relocation, which aligns with the expectation that these connections are more geographically dispersed and represent information and support that can pull a person to other locales. Finally, our results show that these effects are largely consistent across different definitions of relocation (e.g., relocated vs. relocation distance), as well as over different post-disaster time periods (12, 24, and 36 weeks). As we discuss in Chapter 3, these findings of this study have implications for improving the characterization of individuals and communities from the perspective of resilience to disasters.

1.3.3 Open Source Developer Communities Centralize Coordination and Decentralize Work in Response to Increased Engagement After an Attention Shock

Over the last few decades, collaborative crowd communities that exist entirely online have become more common and play an important role within the modern information economy. Their online existence promotes wide accessibility and coupling with other aspects of the Social Web ecosystem, such as social media, search engines, and recommendation systems (Vasilescu, 2014;

Fang et al., 2022). As a consequence, social organizations operate in an environment that is more prone to certain types of external shocks, such as those driven by shifts driven by collective attention (Zhang et al., 2019; Lu et al., 2023; Fang et al., 2022). While prior work, particularly in the domain of organizational management, has considered both the effects of external shocks on organizations and their adaptive responses (Hambrick and Crozier, 1985; Pálovics et al., 2021; Meyer et al., 2005), the organizations that are the focus of these studies have largely been traditional organizations that exist primarily offline. Further, there are a number of other key differences that set online crowd collaborations apart from the traditional organizations. Crowds are made up of loosely knit geographically dispersed groups of people that organize their actions through informal guidelines and collectively developed protocols (Vasilescu, 2014; Vasilescu et al., 2015a). These organizations maintain relatively porous boundaries with their ecosystem, which allow people to join and leave them at their discretion. Due to this volunteer approach to work, crowd members are driven by nonmonetary motivations such as reputation and collective identity (Benkler, 2006; Benkler et al., 2015). In contrast, traditional organizations have clearly defined boundaries, structures of authority and coordination, and formal routines.

In Chapter 3, we investigate the effects of algorithmic attention amplification on external engagement with over a thousand public open-source software projects hosted on GitHub as well as the consequent adaptations of these organizations to them. We identify projects that experienced an attention shock through their appearance at the top of the GitHub Trending Page[2]. We use a potential outcome approach combined with difference-in-difference modeling to compare the adaptations of shocked projects to projects that experienced that similar organic growth and used the same programming language.

Our results show that projects receive substantially higher engagement from potential users and contributors after appearing on the trending page. In response to the sudden increase in questions, critiques, and contributions, experienced members of these project reconfigure its coordination and production processes. These core members serve as a central coordination backbone for an increased population of users and workforce while cutting back on their own development contribution. Instead, they direct the activities of the larger and more modular workforce comprised of many newcomers by delegating work and serving as gatekeepers for changes that are integrated into the project. Despite reconfigured coordination and production networks and the increased use of automation for quality control, the influx of users and contributors can overwhelm projects. We observe that the average responsiveness of the project to external engagements decline across all types by around 40%-50% irrespective of the level of effort required. These results have can in-

[2] It is an algorithmically curated list of software projects on GitHub that have experienced a notable surge in **organic attention**. While the algorithm is not public, we show that appearance on the trending page is correlated with simple measures of growth. The page is found at https://github.com/trending

form members of online crowd organizations of what to expect in the event of an attention shock as well the potential growth opportunities and pains that need to be managed. While we have stopped short of studying the effects of these adaptive behaviors on long-term growth, it is a natural next step to this line of work.

CHAPTER 2

Social Media Trajectories of Recovery in Disaster Stricken Communities

2.1 Introduction

Disasters have a multitude of effects on communities, including psychological and physical trauma, as well as disruption of social structures (Bonanno et al., 2010; Norris et al., 2002). Research has shown that communities can respond in a variety of different ways to these exogenous shocks (Tierney et al., 2001). Some prove to be resilient and are able to bounce back to their pre-disaster normal or even become even more well adjusted and tightly knit. For others, the effects of a disaster can be devastating, inducing negative effects along individual, social, and economic fronts (Bonanno et al., 2010). Additionally, the outcome for a community after a disaster is often mediated by the nature of the disaster itself; was there a warning? Do similar events occur frequently in the region? How severe was the initial impact on the community? Was it a prolonged hazard, such as a wildfire, or sudden but short-lived, such as a tornado? Answers to these questions can be the difference between a community that overcomes adversity to thrive or one that disintegrates (Tierney et al., 2001; Norris et al., 2002; Council, 2006).

In exploring these different questions, researchers have thoroughly analyzed how communities respond to disasters in terms of changes in affect, cognitive processing, social dynamics, and other dimensions (Bonanno et al., 2010; Norris et al., 2002; Perry, 2018). In addition, attempts have been made to develop theoretical frameworks that explain how the effects of a disaster and the subsequent response of the community develop over time (DeWolfe, 2000; Dynes, 1970; Alexander, 1993; Kates and Pijawka, 1977). For a single member of an affected community, the road to recovery can take many forms, including (i) chronic negative effects, psychological, physiological, and social, that last for years, (ii) negative effects that become more visible over time, or even (ii) a positive response despite adversity (Bonanno, 2004). Previous work has discussed how these individual socially embedded reactions to a disaster may shape the response of the larger community. Theories of disaster response, as well as empirical studies of disaster-stricken communities, propose that communities may go through a characteristic series of psychological and social stages as part of a "response trajectory" (DeWolfe, 2000; Dynes, 1970; Kates and Pijawka, 1977).

Community Response on Social Media. Over the past two decades, social media has become the dominant medium through which people, particularly in developed countries, express opinions, consume information, and engage socially. In the context of disasters, this has meant widespread use of social media as a tool and a medium for navigating and responding to the aftermath within and outside affected communities (Kogan et al., 2015; Houston et al., 2015; Lin et al., 2017; Brubaker et al., 2012). Social media has several advantages (and limitations) that are markedly different from the modes of communication and social interaction that preceded them and have influenced their use in the aftermath of disasters. Key among these, social media allow regular citizens to engage rapidly with large audiences and provide tools for parsing large volumes of messages to identify relevant information (e.g., hashtags, search functionality). As a consequence, social media such as Twitter and Facebook represent a primary, but not the only, channel through which those outside a disaster-stricken community, such as the general public, traditional news media, relief agencies, perceive its status and ongoing response to disaster. It is important to understand the extent to which the collective social media behavior of a community affected by disaster can provide a holistic view of the community affected in the aftermath of disaster. Does community response observed on social media contain regular patterns that generalize across different events? How is that behavior similar to or different from "response trajectories" that were identified from offline behavior within such communities before the widespread adoption of social media? In this regard, the existing empirical work falls short.

Studies that have explored the temporal evolution of community response on social media have limited themselves to a single disaster or a few of them (Lin and Margolin, 2014; Olteanu et al., 2015; Lin et al., 2017). We are unaware of any research that has explored the existence of patterns that are relevant in a wide range of disasters. Research along this avenue would have a number of implications. First, if it is established that the community response on social media follows regular patterns, it would allow us to better understand and track new disaster-stricken communities in a sign-posted timeline of their response in the aftermath. To the extent that the social media behaviors of affected communities can be mapped to their real experiences, the ability to track the community response can be very useful for organizations involved in disaster relief and recovery operations. Additionally, this would allow future research to investigate the relationship between contextual variables, such as community attributes, media attention, and support, and deviations in the social media response of individual communities from prototypical patterns. Finally, we expect that the extent to which communities talk on social media about different aspects of their experience will reveal those elements that are being highlighted in the consciousness of the outside public, as well as important gaps in that perception.

Research Questions. We begin with the existing understanding of the response of the community to disasters based on studies of offline behavior preceding social media. A survey of existing work reveals a number of commonalities (DeWolfe, 2000; Dynes, 1970; Kates and Pijawka, 1977). In the face of a disaster, people focus first on the safety of themselves and their families at the beginning of a disaster. Next, immediately after the initial impact, they rescue and help others in the community before any external assistance arrives. In turn, these altruistic behaviors lead to a short-lived subsequent period of enhanced community identity that suppresses some of the negative psychosocial effects. However, eventually communities face long-term ramifications of the disaster that lead to a resurgence of those effects and begin the long road to establish a new normal. Our first goal in this study is to establish whether community responses to disasters that is observed on social media follow some prototypical trajectory over time.

RQ1(a): Does community response reflected on social media exhibit a prototypical trajectory across a wide range of disasters?

If it does emerge that there are distinct patterns in community responses across a variety of disasters, those observations may or may not be explained by previous work, which brings us to a related question.

RQ1(b): Do the temporal patterns of the community response to disasters observed on social media align with existing stage models of disaster response?

Finally, given the numerous intrinsic differences between natural hazards and communities considered in our study, we expect that community responses will show some differences across disasters. Although each disaster is a unique event, previous work has explored the influence of coarse categorizations of disasters, such as type, severity, and duration, on how affected communities recover over time. This leads to our second research question.

RQ2: Are there distinguishable differences in community response trajectories across different disasters and are these correlated with broad disaster categorizations?

The Present Work. In this study, we begin by engaging with the existing literature to understand how individuals and communities respond to disasters considering multiple dimensions of their experiences. Based on this survey, we develop a comprehensive framework of five broad themes (psychological, social, sense-making, biological, and physical) that capture those experiences as the lens through which we observe community response. We apply this framework to Tweets from more than 200 US communities affected by disasters to quantify their short-term disaster response as multidimensional timeseries. We represent each community by Twitter users physically present in the affected area during the disaster, which amounts to more than 2 million users and more than 200 million tweets in all disasters. Finally, we cluster the response trajectories of different communities to establish if there are broadly applicable characteristic patterns of disaster response and if any variations upon these patterns are associated with properties of the specific hazard.

Contributions. This study makes three main contributions. First, we have developed an approach that characterizes community response to disasters over social media as a series of prototypical trajectories that are more temporally fine-grained and precise compared to theoretical models and empirical studies of individual events that precede it; we are able to unpack prototypical behavior over days instead of weeks or months. Furthermore, with this approach, we provide a more nuanced narrative of community response by simultaneously considering the many dimensions of that process, a necessary stepping stone to understand how different aspects of community experience, such as social and emotional dynamics, interact and contribute to eventual outcomes.

Second, our findings reveal that, despite wide variations in the properties of the hazards and affected communities, many aspects of community response reflected on social media (social, aspirational, biological, physical) follow broadly prototypical paths over the short term. These observations on social media broadly align with and validate prior theories on disaster response trajectories. Beyond confirming that community behavior over social media displays regular temporal patterns in the aftermath of disasters, our results reveal a number of previously unobserved themes in the community response. First, we observe a marked decline in conversation that implied productivity that coincides with the sense of risk and uncertainty that comes with disaster. Second, the onset of disaster coincides with people's focus contracting to their immediate social circle, a phenomenon that gradually abates over time, returning their social engagement to pre-disaster levels during our study period. Third, while we observe a sudden decline and gradual recovery of attention toward a number of aspects of life such as *health*, *finances*, *work*, and *food*, we note that references to *food* related dynamics take substantially longer to recover. As previous work has shown that the state of food systems is a reliable indicator of community health and sustainability, this observation hints at the potential utility of the trajectory of *food* related attention measured on Twitter to evaluate community recovery.

Finally, we show that not all behaviors considered in trajectory models are universally observed in disasters, a fact supported by observations from previous disaster research that certain behaviors, such as coping strategies, can vary depending on the particulars and nature of a disaster or the affected community (Bonanno et al., 2010). In our case, a minority of our disasters (25%) exhibit a substantially heightened level of emotions. These communities had very high levels of anger, incredulity, anxiety, and fear at the start of the disaster compared to others. Additionally, after initially trending downward in the aftermath of a disaster, within a few weeks fear and anxiety experience a resurgence within the social media feeds of these communities. This has implications for organizations providing relief and recovery assistance, as negative affects such as anger and anxiety play a mediating role in long-term recovery outcomes in communities (Forbes et al., 2015). This resurgence coincides with the continued signs of elevated cognitive processing associated with sensemaking within these same communities. These observations are consistent with an ongoing struggle to make sense of and adapt to a *new normal*. The possibility that these communities may be struggling in the aftermath of a disaster, more so than the majority, is reinforced by the fact that the amount of aid received by these communities, a proxy for the destructiveness of a disaster, is on average higher for these communities after accounting for the type of disaster and duration.

2.2 Background and Related Work

2.2.1 Dimensions of Disaster Response

Previous work has identified a wide range of different dimensions along which communities express their response to disasters (Bonanno et al., 2010; Norris et al., 2002; Grimm et al., 2012; Anderson et al., 2016). We take advantage of this literature to establish a thematic framework that allows us to build a comprehensive picture of disaster response in a tractable manner. Our framework is organized along five broad themes; (i) Psychological, (ii) Psychosocial, (iii) Sensemaking, (iv) Biological, and (v) Physical.

2.2.1.1 Psychological

Disasters impose substantial psychological stress on individuals and communities; affected individuals must deal with fear for their safety, loss of loved ones and homes, and disruption of their social network and support system. Previous work has theorized and, in some cases, empirically studied a range of different psychological manifestations in individuals as they navigate these challenging circumstances. They may manifest this distress in the form of posttraumatic stress disorder, depression, anxiety, and other related disorders (Bonanno et al., 2010; Yule et al., 2000; Cherry et al., 2021; Anderson et al., 2016; DeWolfe, 2000; Kar and Bastia, 2006). Although previous work has consistently documented the increased prevalence of these disorders in disaster-stricken communities, there is substantial disagreement on the magnitude of this effect, with estimates showing a wide variance (Bonanno et al., 2010). The level of psychological distress experienced within a specific community is likely to be influenced by a variety of variables, such as disaster characteristics, community demographics, and social resources.

Our interest lies in measuring the general psychological well-being of the community rather than identifying the prevalence of specific psychological disorders. As such, we limit our attention to the affective and cognitive dimensions of expression that are commonly associated with psychological well-being that are observed within disaster-stricken communities. For example, people diagnosed with one or more of a variety of clinically diagnosed psychological disorders, such as PTSD, depression, anxiety, and other mood disorders, are likely to exhibit elevated levels of negative affect, in the form of sadness, anger, or irritability, anxiety, and fear, among others (Association, 2013; Bonanno et al., 2010). Additionally, diagnostic definitions of these disorders also include cognitive and behavioral symptoms such as loss of energy and motivation, withdrawal, and helplessness.

2.2.1.2 Social

Social dynamics within a community before, during, and after a disaster are central to disaster response. In fact, increasingly, interpretations of disasters place emphasis on their socially constructed nature (Tierney et al., 2001; Drabek, 2018; Horlick-Jones, 1995). The social response to disaster, during and after it, can be considered in a number of related dimensions. Previous work has shown that the aftermath of a disaster can bring a community together. People are more likely to highlight their membership in the community and refer to that affiliation positively (DeWolfe, 2000; Kates and Pijawka, 1977). They may take advantage of that shared affiliation to exhort other community members to unite to withstand the disaster. In turn, this shared sense of community can increase social cohesion and lead to a more efficient and effective community response. Additionally, in the early stages of disaster recovery, community members are more likely to seek social support, be more prosocial, and develop solidarity. However, in the long term, the ability of communities to provide support may be exhausted, although many affected people are still in need (DeWolfe, 2000; Bonanno et al., 2010). A more narrow and personal aspect of the social dimension of community response is that individual concern for the safety and well-being of immediate family members, friends, and neighbors is a strong driving force behind behavior during the early stages of a disaster (Bourque et al., 2006; Bonanno et al., 2010; DeWolfe, 2000).

2.2.1.3 Sensemaking

At the beginning, during and after a disaster, members of the stricken community collectively make sense of their new, changed, and uncertain reality to adopt new survival norms (Tierney et al., 2001; Lindell, 2003). This includes evaluating the authenticity, source, and extent of the changes caused by the hazard, as well as possible actions that could be taken to mitigate the risks. As the disaster continues to unfold, sustained uncertainty regarding conditions makes the sense-making process challenging; the community may suffer from information scarcity about some aspects of the disaster, while in others they may have to sort through a vast amount of information from different sources of widely varying authenticity. Therefore, during and in the aftermath of a disaster, it is likely that members of the community collectively display an elevated level of cognitive processing. As time passes, more information naturally emerges, bringing with it increasing stability and certainty until eventually the community collectively achieves a coherent understanding of the circumstances.

2.2.1.4 Biological

Natural disasters are exogenous shocks that manifest in the physical environment of communities. Consequently, they can have direct and immediately visible effects on members of a community, including various degrees of discomfort and injury, including loss of life. Additionally, people are known to exhibit psychosomatic symptoms, such as headaches, dizziness, and soreness, due to disaster-induced psychological stress (van den Berg et al., 2005; Bonanno et al., 2010). The prevalence of such conditions, often medically unexplained, may indicate unresolved psychological and social issues among members of the community prompted by the disaster. Therefore, while a community may focus on immediate physical danger and harm in the short term after a disaster, there are other more long-term biological indicators that could characterize the trajectory of overall well-being.

2.2.1.5 Physical

Beyond the immediate risk of bodily harm, the most visible impact of a disaster is the destruction of infrastructure and the disruption of normal life. People in affected communities can have their homes and neighborhoods damaged or completely destroyed. Much of the attention of early disaster recovery research focused on this dimension, as it reflects the immediate goal of many large-scale relief efforts in disaster-stricken regions. Although more recent research increasingly considers the other aspects equally or more crucial in the recovery process, the intensity and duration for which physical reconstruction remains in community consciousness still have the potential to influence the overall recovery trajectory of the community (Lindell, 2013; Tierney et al., 2001). This is because many of the elements of the physical environment (Ex: homes, utilities, roads) correspond to the bottom tier in Maslow's hierarchy of needs (Maslow, 1954) and as such may be pre-requisites for addressing the more complex psychological and social needs faced by members of the community.

In the rest of the paper, this thematic framework serves as the overarching organizing principle. First, it is the lens through which we identify and catalog different dimensions of community response and the corresponding linguistic markers in tweets, and second, it is also the means by which we organize and interpret our observations of how community response changes over time.

2.2.2 Trajectories of Community Response and Recovery

In this section, we discuss theoretical models that shape the current understanding of how communities respond to disasters. These conceptual models often play an important role in the development of institutional practices to provide and support relief and relief to disaster-stricken communities (DeWolfe, 2000; Townshend et al., 2015). Here, they serve as a useful baseline for placing our own observations in the context of existing knowledge on community psychosocial dynamics during disaster recovery.

In general, conceptual models of disaster response and recovery characterize these collective processes as distaster-stricken communities moving through a series of steps or arcs for simplicity while acknowledging that some of these arcs often overlap in practice (Myers, 1994; Raphael, 1986; Dynes, 1970; Kates and Pijawka, 1977; Alexander, 1993; Williams, 1999). In these models, the community response begins immediately before the onset, when they may become aware of the imminent danger and take protective action (Myers, 1994; Williams, 1999; Raphael, 1986). This pre-disaster stage is usually known as the *warning phase*, as people actively disseminate warnings about the event within their communities in addition to taking other protective actions, a pattern of behavior that has been consistently observed both offline and online through social media (DeWolfe, 2000; Sutton et al., 2015). The degree of forewarning received by a community is likely to affect its response. Sudden and unexpected incidents are likely to evoke feelings of helplessness, fear, and vulnerability, while incidents that give substantial prior warning can cause feelings of guilt and blame within a community that is not adequately prepared (Tierney et al., 2001; DeWolfe, 2000). The length of this phase during a particular disaster depends on the level of warning the affected community received as well as whether the hazard, while active, continues to evolve.

The *warning phase* is followed by *impact* or *shock*, the period during which the hazard remains active within the affected community. Depending on the severity of the hazard, communities may experience a range of psychological effects ranging from anxiety and vigilance to collective panic. During this period, people focus on their personal safety and that of their family (Bonanno et al., 2010; Faas and Jones, 2017; Tierney et al., 2001). The impact phase can last for as long as a few minutes, as in the case of an earthquake, or hours or days, as in the case of hazards like storms and wildfires.

In theoretical models, the period following the onset of disaster is variously known as *heroic phase*, *honeymoon period*, *therapeutic community* etc. (Myers, 1994; Raphael, 1986; Alexander, 1993). During this stage, driven by a boosted prosocial orientation, people engage in relief efforts within their neighborhoods, which, in turn, evokes a strong sense of community spirit within the community that may last weeks. During the first few weeks after a disaster, this dynamic may suppress the negative psychological and psychosocial effects of the disaster.

Eventually, people exhaust their optimism and their capacity to support other community members. Concurrently, both external attention from the wider public and relief efforts begin to wane. This support vacuum marks the end of the *honeymoon* stage. At this point, affected communities take stock of the long-term repercussions of the disaster and the challenges to returning to normal. This period is known as the *disillusionment phase* (DeWolfe, 2000; Raphael, 1986) or the *Inner contradiction and crisis phases* (Williams, 1999) in reference to the cognitive dissonance that communities experience as they return to the harsh post-disaster reality after being buoyed temporarily through external support and social cohesion.

The disillusion phase can last for many months (DeWolfe, 2000). Communities experience a resurgence in negative affect, as well as maladaptive behaviors such as psychosomatic symptoms and substance abuse. Eventually, communities enter the*reconstruction* or *recovery* phase, where they come to terms with their new reality and return to stable behavioral patterns, which may or may not be similar to pre-disaster patterns.

2.2.3 Social Media Based Disaster Research

Social media is widely used during and after disasters by affected individuals, media, relief agencies, and volunteers to search for and disseminate information and coordinate relief work (Shklovski et al., 2008; Starbird and Palen, 2010; Lobb et al., 2012). Previous research based on Twitter has covered a range of topics, such as the extraction of disaster-relevant content, assessing situational awareness, as well as information seeking and diffusion behaviors of affected communities and their collective emotions (Imran et al., 2015).

Most of these previous studies have considered the information produced and shared on Twitter during disasters from the perspective of how it serves to improve awareness of the evolving circumstances among affected communities, the general public, and relief organizations. These analyses are interested in the types of information shared and their temporal variation. Information categories such as warnings and advice, donations, and volunteerism feature prominently during disasters and their aftermath (Olteanu et al., 2015; Imran et al., 2015, 2013). Although they provide useful contextual knowledge about Twitter use during disasters and how it can be used to improve relief efforts, they are less useful for understanding community well-being over time.

Less frequently, studies have explored the behavior of Twitter users affected by a disaster through comparative pre-disaster vs. post-disaster analysis (Olteanu et al., 2014; Murthy and Longwell, 2013) as well as the evolution of behavior over time (Olteanu et al., 2015; David et al., 2016). However, work that differentiates behavior during qualitatively different periods (e.g. before, during, immediate aftermath, and long-term) of the disaster's life cycle is rare. Similarly, work that draws a clear contrast between the community that is affected by an event and the geographically diffused population of other interested individuals is scarce. Kogan et al. (Kogan et al., 2015) is a useful case study of retweet behavior during Hurricane Sandy that incorporates both these elements. They compare the behavior of users who were present in the affected region with that of others who simply tweet about the disaster. Further, they consider the variation in behavior in discrete time intervals before, during, and after the disaster. In contrast, we focus only on the behavior of vulnerable users over time as a continuous evolving process.

Some studies have used social media to explore the psychological effects of disasters on affected populations. This literature has established that linguistic markers in tweets can reveal the psychological states, interests, and opinions of individual users and communities exposed to disasters (De Choudhury et al., 2014; Lin and Margolin, 2014; Olteanu et al., 2015). However, these studies either focus on a single emotion, such as fear or grief, or consider the general prevalence of different emotions during the entire observation period. On the contrary, we analyze the temporal evolution of the collective psychological state of communities along multiple affect and social dimensions. One of the first studies in this space, De Choudhury et al. (De Choudhury et al., 2014), examines the affective response and the resulting desensitization in the context of the "Drug War" related violence in Mexico. In a similar vein, Wen and Lin (Wen and Lin, 2016) estimate the prevalence of anxiety, anger, and sadness over time in tweets produced by a community affected by a terror attack. Furthermore, Garcia and Rimeé (Garcia and Rimé, 2019) analyzed collective emotions and language related to prosocial behavior and solidarity in French tweets after the Paris terrorist attacks of November 2015. However, their work is limited to the analysis of a single event and does not identify patterns of community response that generalize across a range of disasters. In this sense, the study by Saha and De Choudhury (Saha and De Choudhury, 2017) of the temporal dynamics of stress and other psychological markers in 12 cases of college gun violence using Reddit, although limited to a few events of a different type of crisis, is more closely related to our work.

2.3 Disaster Dataset

In this section, we describe the selected set of disasters, the procedure to define spatio-temporal boundaries for disasters, and Twitter data collection.

2.3.1 Spatiotemporal Boundaries

We use disaster declarations reported by the Federal Emergency Management Agency (FEMA)[1] between July 2011 and September 2016 to estimate spatiotemporal boundaries for reported disas-

[1] https://www.fema.gov/openfema-data-page/disaster-declarations-summaries-v2

ters. There are 530 declarations during this period. We extract six properties from each declaration; (i) *Disaster Type (e.g Flood)*, *Start Date*, *End Date*, *State*, *County*, and *Aid Amount*. Since FEMA makes regional declarations when the required criteria are met, we find that multiple declarations are frequently associated with a single event. Since we are interested in studying communities as a whole, we group FEMA declarations that correspond to the same *disaster event*. Our intuition for grouping a set of declarations into a disaster event is that they must: i) be of the same incident type, ii) be spatially close, and iii) be temporally close. Our grouping procedure consists of two steps:

1. **Group by type and time**: Declarations are split into ordered sequences of the same type so that the starting dates of consecutive declarations in a sequence are no more than 5 days apart. We chose the threshold of 5 days through trial and error, starting with one day and increasing one day at a time. For each threshold, we manually inspected a sample of the generated candidate groups and estimated how often and accurately their spatiotemporal boundaries aligned with actual events.

2. **Identify spatially connected sequences**: For each sequence of declarations, we construct a network from the corresponding states, where each state is a node, and there is a link between two states if they share a border. Each connected component in this network is assigned to a different disaster event. We assume that disaster events do not move directly between states that do not share a border. Note that this definition allows disaster events to span large geographic areas as long as they move across connected states.

For each identified disaster, we estimate the temporal and spatial boundaries as follows.

- Temporal boundaries: The period from the start date of the first declaration to the end date of the last declaration.

- Spatial boundaries: First, we identify the smallest rectangle that covers all counties recorded for a particular disaster. Next, we expand the boundary in all directions by 10% of the length of the rectangle to include potentially affected periphery regions not covered by the declarations.

Validation.

This procedure produces a set of 350 candidate disasters. To evaluate their validity, we verify the dates and location of disasters using Google Trends[2] and online news articles. Google Trends

[2] https://trends.google.com/trends/

provides a time series of the relative Google search volume for a query during a given time interval and geographic area. For each group, we search for the disaster type restricting to the temporal bounds of the disaster and the states involved. We verify that the search volume peaks around the time of the first declaration for most candidates. In cases where the search volume is unavailable, we search for news articles about the disaster to confirm the validity of our groups. When our best estimate for the start of a disaster based on this approach differs from the first declaration of the corresponding group, we use our estimate. The starting dates of 37% of the disasters are adjusted this way with mean and median adjustments of 5.6 and 3 days, respectively.

2.3.2 Social Media Data

We use the following procedure to identify Twitter users associated with each disaster and collect their tweets.

Disaster affected Twitter Users.

We use an archive of 10% of the Twitter stream to identify users who were in the affected area during the time of the disaster. We define affected users as those who had at least one geotagged tweet within the spatial boundaries of each disaster between one week before the start of the disaster and one week after the end of the disaster. Of the 350 disasters identified from the FEMA declarations, we select 203 disasters that had at least 100 affected users. This dataset contains 13 different disaster types; Severe Storm(s) (72), Fire (66), Flood (35), Severe Ice Storm (7), Hurricane (5), Snow (5), Tornado (3), Other (3), Coastal Storm (2), Earthquake (2), Mud / Landslide (1), Typhoon (1), and Volcano (1). The total number of users identified this way is 3,984,530.

Collecting Twitter User Timelines.

For each disaster, we focus on a period spanning 8 weeks, from 4 weeks before the start date to 4 weeks after it. We collect the 3,200 latest tweets (a limit imposed by Twitter), including retweets, for each user in our sample using the Twitter API. Approximately 8% of the users in all disasters had more tweets than this number at the time of collection. For some of these users, particularly very active ones, we are unable to collect all tweets during the disaster period. In addition, some users have made their tweets private or deleted their accounts. Finally, since our goal is to capture changes in the language of users before and after the disaster, we limit the analysis to users who posted at least once before and once after the onset of the disaster. This ensures that our baseline observations before the disaster are generated by the same population of users as those we observe after the disaster. We are able to obtain a total of 2,082,210 users who meet these criteria corresponding to approximately 227 million tweets. Figure 2.1 shows the distribution of the number of users across disasters after applying this final constraint. Figure 2.2 shows that disasters in our data set are spread over a wide range of geographic areas in the United States.

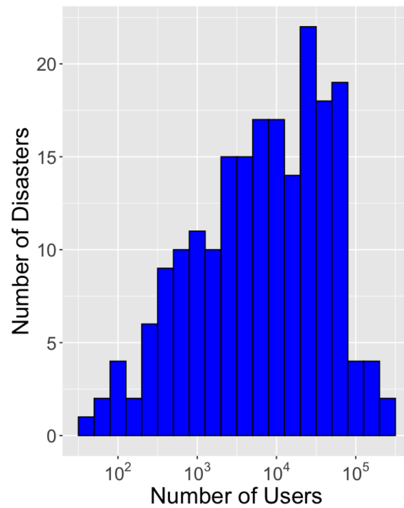

Figure 2.1: Histogram of the size of the disaster in terms of the number of associated Twitter users

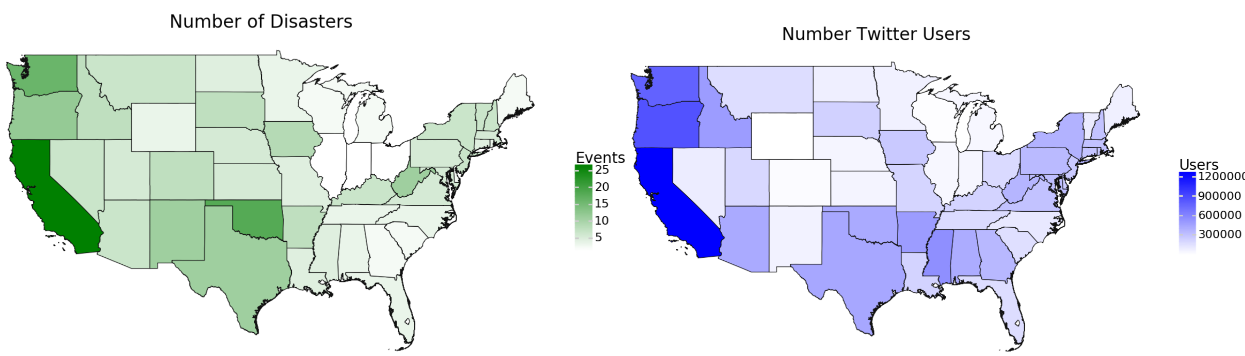

Figure 2.2: Data volume maps in the contiguous US. Left: Number of incidents per state in our dataset. Right: Aggregated number of users appearing in all disasters in each state

2.4 Methods

We begin this section with a discussion of the rationale for choosing a four-week period after the onset of disaster for our study. Next, we describe our methodology to characterize community responses. We use a lexicon-based approach to estimate different dimensions of community behavior based on tweets from affected users. We use these estimates to build a multidimensional time series of disaster response for each disaster that controls for pre-disaster behavior. Then we cluster these responses to identify groups of disasters that have similar trajectories along different dimensions. Finally, to discover any prototypical responses, we compare the contribution of dimensions to the differences between clusters.

2.4.1 Length of the Study Period

Two important methodological questions related to the study period must be answered when quantifying and comparing the community response in many disasters. The first is whether we should vary the duration of the study period for each disaster, since they may differ in terms of the duration of immediate impact. In this study, we choose to fix the length of the study period for all disasters, since having different study periods across disasters would make any form of temporal comparison challenging. Instead, we conduct a post-comparison statistical test to establish if our findings are affected by a relationship between the duration of a disaster and the community response.

The second question is what the length of the study period should be if it is to be fixed across all events. In this study, we choose a four-week period after the onset of the disaster for two reasons. First, due to normal attrition of Twitter users, the dataset has fewer users, and thus confidence in observed community behavior would be lower further away from the onset of disaster. Second, a longer study period increases the probability that a community will be exposed to new disasters or other shocks that would interact with their ongoing response to the original disaster.

2.4.2 Linguistic Dimensions of Disaster Response

We estimate the aggregate prevalence of linguistic markers associated with different themes of response for each disaster over a two-month period. A supervised learning approach to achieve this goal would have needed a large volume of human-annotated tweets. Instead, we use a lexicon-based approach. We use lists of words associated with the markers selected from validated existing lexicons to estimate a normalized index for each linguistic marker by counting the number of associated words in tweets. Previous work has shown that when using established lexicons, this approach is effective in estimating psychological and social markers in text, including tweets (De Choudhury and Counts, 2013; Cohn et al., 2004). Further, recent studies comparing the performance of lexicons against large language models (LLM) have shown that linguistic time series constructed from tweets using both methods correlate well with data from representative surveys and that, in this regard, lexicon performance is comparable to LLM (Pellert et al., 2022; Garcia et al., 2021).

We use LIWC (2015) [3] as the primary lexicon. One limitation of LIWC is that it contains relatively few categories that focus on emotion (*positive*, *negative sentiment*, textitanger, *anxiety*

[3]https://liwc.wpengine.com/

and *sadness*). One of the often mentioned affective dimensions in the literature on disaster response is *fear*. Since it is absent in LIWC, we use the corresponding word list from the NRC lexicon (Mohammad and Turney, 2013) to supplement our analyses with the dimension *fear*. Table 2.2 shows the list of dimensions used to represent the response to disasters (22 from LIWC, 1 from NRC).

2.4.3 Disaster Response Trajectories

For each disaster, we estimate an aggregate time series for each dimension. Let N_i be the number of tweets on day i of the disaster and $N_{p,i}$ be the number of tweets on day i that contain at least one word of dimension p. The measurement of category p on day i of the disaster is $T_{p,i} = \frac{N_{p,i}}{N_i}$. Next, we adjust the estimates to account for any differences in the content of people's tweets due to what day of the week it is. We achieve this by fitting an additive time-series model to the data and subtracting the estimated day-of-the-week seasonal estimates from our observations.

These numbers include retweets, replies, and quote tweets in addition to original tweets. We include retweets since the act of retweeting suggests that any affect in the tweet is shared between the original user and the retweeting user. This is particularly relevant in our case, as previous work has shown that when a disaster strikes, affected users preferentially retweet others from the same community (Kogan et al., 2015). In the case of replies and quote tweets, we include only the text of the replying/quoting user.

Each disaster in our dataset affected a different community defined by unique geographical boundaries. The typical behaviors of these communities can differ for a variety of social, cultural, demographic, and economic reasons. To compare the response of the community across disasters, we need to account for these differences in the typical pre-disaster behavior of communities. We do this by adjusting observed behavior in the temporal vicinity of the disasters by typical behavior uniquely estimated for each community from a corresponding pre-disaster baseline period. For each dimension p, we split the time into a baseline period of 3 weeks, beginning 4 weeks before the start of the disaster and ending 1 week before it, and a study period of 5 weeks from 1 week before the start to 4 weeks after it. Next, we construct a time series for the study period defined by the relative difference in intensity $T_{p,i}$ for each dimension p during the study and the baseline periods. More precisely, we estimate the baseline rate for dimension p, P_p, as the fraction of tweets that had at least one word from p from all tweets during the baseline period. The normalized estimate for p on day i^{th} of the study period is defined as $N_{p,i} = \frac{T_{p,i} - \mu}{\sigma}$, where μ and σ are the mean and standard deviation of the distribution of daily estimates during the baseline period.

Theme	Lexicon Dimension	References
Psychological	Anxiety	(Norris et al., 1994; DeWolfe, 2000)
	Anger	(DeWolfe, 2000; Forbes et al., 2015)
	Swear	(DeWolfe, 2000)
	Sadness	(Norris et al., 1994; DeWolfe, 2000)
	Fear	(Norris et al., 1994; DeWolfe, 2000; Grimm et al., 2014)
	Achieve	(Norris et al., 1994; DeWolfe, 2000; Grimm et al., 2014)
	Reward	(DeWolfe, 2000; Grimm et al., 2014)
Psychosocial	Social	(DeWolfe, 2000; Lindell, 2013)
	We	(DeWolfe, 2000; Gortner and Pennebaker, 2003; Richardson et al., 2014)
	Family	(DeWolfe, 2000; Tierney et al., 2001)
	Friend	(Tierney et al., 2001)
	Affiliation	(DeWolfe, 2000; Richardson et al., 2014)
Sensemaking	Risk	(DeWolfe, 2000; Lindell, 2013; Grimm et al., 2014)
	Cause	(Gortner and Pennebaker, 2003; Grimm et al., 2014; Coffelt et al., 2010)
	Tentative	(Grimm et al., 2014; Coffelt et al., 2010)
	Certain	(Grimm et al., 2014; Coffelt et al., 2010)
	Differ	(Grimm et al., 2014; Coffelt et al., 2010)
	Insight	(Gortner and Pennebaker, 2003; Coffelt et al., 2010)
Biological	Health	(Norris et al., 1994; DeWolfe, 2000; Gortner and Pennebaker, 2003; van den Berg et al., 2005)
	Ingest	(van den Berg et al., 2005; Grimm et al., 2014; Kuijer and Boyce, 2012)
	Body	(Norris et al., 1994; DeWolfe, 2000; van den Berg et al., 2005)
Physical	Home	(DeWolfe, 2000; Lindell, 2013; Tierney et al., 2001)
	Work	(Lindell, 2013; Tierney et al., 2001)
	Money	(DeWolfe, 2000; Tierney et al., 2001)

Table 2.2: Selected Lexicon Dimensions. All dimensions except Fear, for which used the NRC lexicon, are sourced from LIWC

2.4.4 Clustering Response Trajectories

Our first objective with this study is to establish the existence (or absence) of regular patterns in community response across different disasters. We approach this objective by using a clustering approach. Clustering can highlight aspects of community response over social media that differ across events substantially as these dimensions will contribute to the decision criteria that define the resulting clusters. Conversely, the dimensions that do not contribute to the clustering decision are those that have similar trajectories across the clusters (i.e., all disasters). Next, we describe how this rationale is extended to identify prototypical patterns of community response, as well as elements of response that show marked differences among groups of disasters.

The adjusted community response for each disaster is represented by a 5-week, 24-dimensional time series. Our next goal is to find clusters of disasters with similar response time series. The estimated time series suffers from noise to which most well-known time series clustering methods are sensitive, particularly under high dimensionality (Kotsakos et al., 2013). Therefore, we first smooth individual dimension time series by fitting locally weighted first-order regressions with a span of 20%. Next, we estimate the pairwise dissimilarity between response trajectories using the multidimensional dynamic time warping distance (DTW) (Kotsakos et al., 2013). DTW is generally superior to Euclidean distance in the case of time-series clustering, since it is purpose-built to account for temporal shifts; it finds the best compromise between alignment of the shapes of two time series and the cost of shifting (in time) necessary to achieve a given alignment. We use hierarchical clustering with Ward's criterion on the resulting trajectory-trajectory distance matrix to separate our dataset into groups of disasters that have similar community response trajectories. We evaluate the quality of the results for different numbers of clusters ranging from 2 to 10 using Silhouette Coefficient, Dunn Index, and Connectivity (C). We find that dividing the disasters into 2 clusters provides the best results based on all three metrics (Figure A.2). Disasters are unevenly divided among the two clusters, the smaller cluster (Cluster 1) containing 58 disasters and the larger one (Cluster 2) containing 145 disasters. Table 2.3 shows how different types of disasters are distributed between the two clusters. We generate a representative signature of each cluster using DTW barycenter average (DBA), a method that finds a representation that minimizes the sum of the squares of the DTW distance to the members of a cluster (Petitjean et al., 2011).

2.5 Results

Next, we discuss our findings in terms of disaster response trajectories and address how they relate to our research questions. Note that we base our observations and consequent interpretations solely on the trajectories relative to baseline period. This is true even for cases where we compare the

Type	Cluster 1	Cluster 2
Severe Storm	24	48
Fire	24	42
Flood	5	30
Severe Ice Storm	1	6
Snow	0	5
Hurricane	0	5
Other	1	2
Tornado	0	3
Earthquake	2	0
Coastal Storm	0	2
Mud/Landslide	0	1
Volcano	1	0
Typhoon	0	1

Table 2.3: Distribution of events by the type of hazard in cluster 1 and cluster 2

trajectories of different linguistic dimensions with each other. As an example, if, on a particular day, the prevalence of *Fear* in a community is one standard deviation above its mean during the baseline, while that of *Social* is only half a standard deviation above its corresponding mean, we would interpret this observation as *Fear* being more prominent within the community than *Social*, even if, in absolute terms, *Social* is much more prevalent in community tweets.

2.5.1 RQ1: Universal Patterns in Disaster Recovery Trajectories

While the two clusters correspond to distinct disaster response patterns, we find that the differences between them are not evenly distributed across the different dimensions. Further, the differences between clusters along a given dimension can stem from two sources; (i) scale of observations and (ii) shape of the time series. The scale of observations signal the intensity of the community experience, whereas the shape of the trajectory is more informative of how behavior evolved over time. To disentangle these two aspects of scale and shape, we estimate them separately as shown in Figure 2.3. For a given cluster, we estimate the characteristic scale of a dimension as the root mean square magnitude of its average trajectory. Consequently, the difference between the two clusters is the difference in the characteristic scales. To estimate the difference between clusters in terms of just shape for a given dimension, we normalize the average trajectory of each cluster by the largest magnitude observed in it. This normalization scales a time series to approximately the same scale with a maximum magnitude of one, while preserving its shape. Then, the difference in shape between the two clusters is estimated as the DTW distance between these normalized average time series. Figure 2.4 shows how the differences in dimensions between clusters are distributed between shape and scale. For dimensions above the 45ine, the difference across clusters is driven more by the scale of observations than the shape, while the reverse is true for those below it. We consider dimensions that fall within the lower left quadrant of the figure, defined by $(Scale < 0.5\ and\ Shape < 0.5)$, to be dimensions for which the clusters have similar trajectories. Thus, 14 dimensions with prototypical patterns: *achieve*, *reward*, *social*, *we*, *family*, *friend*, *affiliation*, *risk*, *certain*, *health*, *ingest*, *home*, *work*, and *money*. Figure 2.5 shows the average time series in all disasters for these dimensions. Next, we describe our findings based on these average time series and discuss how they relate to existing theory on disaster response.

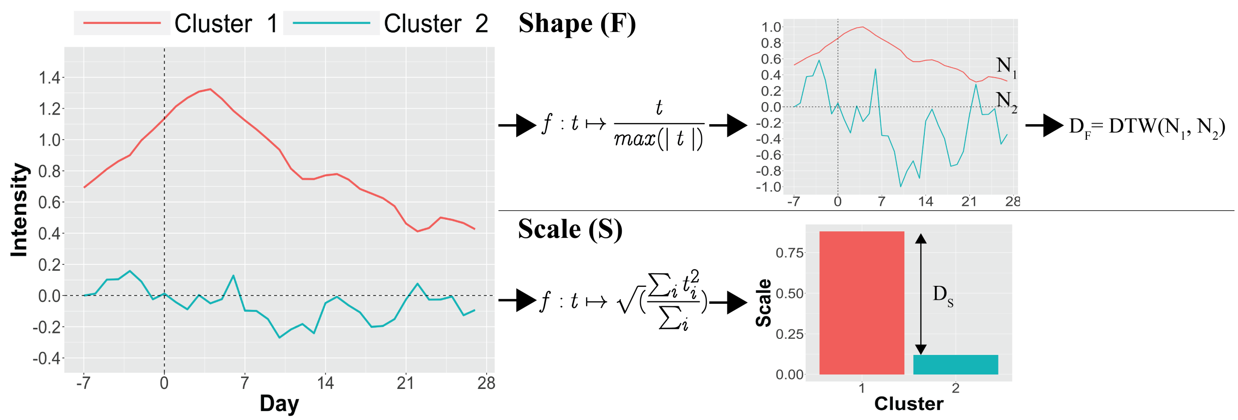

Figure 2.3: Scale and shape difference for *anger* between clusters 1 and 2. The difference in scale is measured as the absolute difference in the root mean sum of squares of magnitudes from the average trajectories of the two clusters. The shape contribution is measured as the DTW distance between magnitude-normalized average trajectories

2.5.1.1 Sense of Productivity is Lost

The normal functioning of a community depends on the stability of its physical and social environments. Disasters tend to suddenly and significantly disrupt these. In our results, the two prototypical dimensions of sensemaking, *Certain* (Fig. 2.5j) and *Risk* (2.5g), reflect how communities experience this disruption. We observe that communities experience a high level of uncertainty from before the onset of the disaster up to two weeks after it, before gradually recovering to normal levels. In contrast, communities perceive a significant threat to themselves only immediately after the event, that is, while the hazard is still active, which is reflected as a spike in *Risk*. A more detailed analysis showed that the increase in *Risk* is mainly driven by tweets associated with warnings (A.3)

In addition, we observe that the two psychological dimensions, *Achieve* (Fig. 2.5a) and *Reward* (Fig. 2.5d), have nearly identical prototypical patterns. They show a significant decline from baseline in the immediate aftermath of the disaster followed by a gradual return to normal by week 4.

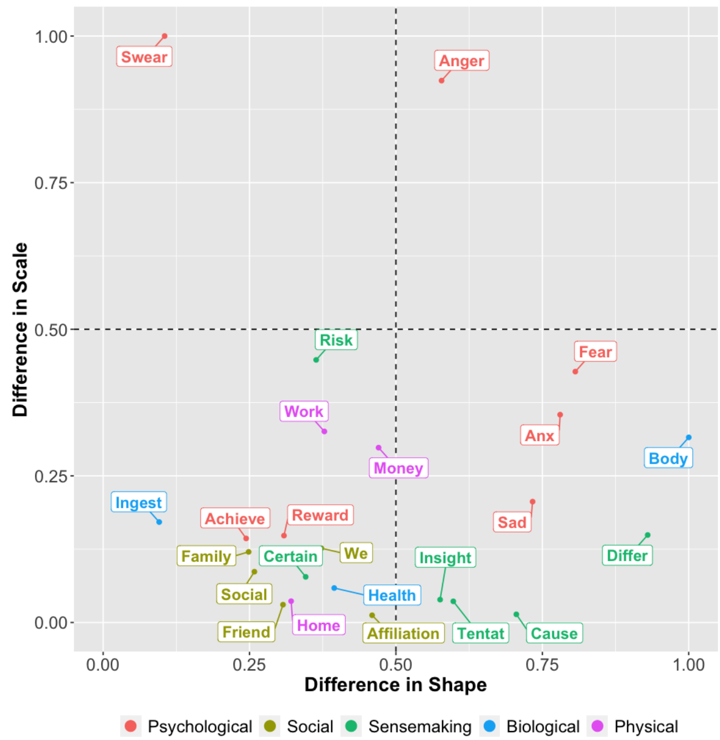

Figure 2.4: Scale vs. shape contributions to differences in dimensions. The values in both axes are scaled to the range of [0,1]

To unpack this observation, we examine the contribution of individual words to these categories (Figs. A.4 and A.5). In *Achieve*, the decline is driven by words that implied celebration (e.g. win, celebrat*, proud, award) or a positive outlook in general (e.g., improv*, challenge*, achieve, earned). The decline in these words is counterbalanced to some extent by an increase in comparatively neutral words (e.g. work, team*, lead). *Reward* follows a similar, if less clear, pattern with words that imply a positive outlook (e.g. good, goal, scor*, promot*, adventur*) declining, while more prosaic words (e.g. get, got, take, earn*) become more frequent. There are several possible behavioral changes within affected communities that, individually or in combination, explain these observations. In the aftermath of the disaster, when the devastation is easily apparent, people may feel less celebratory. Another possibility is that the disruption caused by the disaster reduces people's engagement with the spheres of life they associate with goals and a sense of achievement (E.g. Career, School, Sports). A third possibility is that people in the affected community are being more selective with their Twitter use, choosing to focus on disaster-relevant content and self-censoring any routine positive content that they deem irrelevant or inappropriate. We leave the disentanglement among these possibilities for future work.

2.5.1.2 Social Life Contracts and Gradually Relaxes to Pre-Disaster Levels

Given the disruption to local social life that can be caused by disasters, it is possible that communities would show a reduced level of social activity, at least for a while, after a disaster. However, previous work suggests that people would demonstrate greater community spirit and prosociality as they mobilize collectively to face the effects of a disaster.

Among the prototypical patterns in our results, *Social* (Fig. 2.5b) dimension, a measure of overall sociality, shows a substantial decline from baseline at the start of the disaster and gradually recovers to normal levels by week 4. Although this supports the first argument, a look at the contributions of individual words to this aggregate observation provides valuable nuance. In the aftermath of disasters, the observed decrease in *Social* is mainly due to a reduction in the relative rate of retweets (Fig. A.6) [4]. Communities tweet and retweet more often during and after a disaster, but are less likely to retweet compared to posting their own tweets. We suggest that the relative decline in retweets is due to affected communities becoming more discerning with the content they boost with their limited energy. Previous work has shown that vulnerable populations focus on locally produced content rather than the much larger volume of outside content during a disaster (Kogan et al., 2015).

2.5.1.3 Food Related Processes Rebound Slower than Other Aspects of Life

We observe that among the dimensions associated with biological and physical processes, disaster-induced perturbations have subsided quickly in all except one dimension, the *Ingest* dimension associated with food-related behavior. For dimensions *Health*, *Money*, and *Work*, the onset of disaster results in a decline in attention, but levels return to pre-disaster levels within a week. A word-level analysis of this short-lived decline shows that it corresponds to affected communities referring to routine preexisting behavior and conditions less often. For example, under *Health*, people were less likely to mention the words doctor*, fitness*, cancer*, and diet* among others (A.7) and under *Money*, words such as market*, shopp*, sale, and debt* occurred less often in their tweets (A.8)[5].

Compared to other dimensions that capture biological / physical processes, *Ingest* takes much longer (4 weeks), to return to pre-disaster levels from the perturbation (decline) observed around the onset of disaster. A word-level analysis shows that the decrease in *Ingest* is driven by words such as dinner, lunch, beer, cake, and coffee A.11. Other words such as water, eat, hungry, egg, rice, and thirsty actually increase in the aftermath of disaster [6]. Words that decrease in frequency appear

[4]The text of a retweet contains the prefix "RT" when it is recovered from the API

[5]These observations were consistent between the entire post-onset period as well as considering only the first week after onset

[6]We also observe a substantial increase in smoke*. However, we associate this increase with wildfire smoke as

to be associated with relatively high-level food needs and activities, while those that increase correspond to references to basic sustenance (Satter, 2007). We suggest that this observation reflects the reality that disasters inflict substantial damage to food infrastructure (restaurants, supermarkets, and grocery stores), limiting food diversity and accessibility within affected communities. It is more difficult to address why food-related activities take longer to rebound compared to those related to health, employment, and finances. One possibility is that food systems rely on relatively more types of other critical infrastructure (such as energy for equipment, water for growing, roads for transportation, communication technology for coordination), which are likely to be disrupted by disasters (Dobie et al., 2019). In fact, due to this dependence, the state of the food infrastructure has been suggested as a possible indicator of short- and long-term recovery from disasters.

2.5.2 RQ2: Variations in Community Disaster Response

We identify dimensions that have large differences between clusters as those that fall outside the lower left quadrant ($Scale > 0.5$ *or* $Shape > 0.5$) which includes all affective psychological dimensions (*anger*, *swear*, *sadness*, *anxiety*, and *fear*), four sensemaking dimensions (*insight*, *cause*, *differ*, and *tentative*) and the physical dimension of *body* (Fig. 2.6).

Overall, the two disaster clusters are differentiated by a consistent difference in the magnitudes of the time series. The smaller cluster of 58 disasters (**Cluster 1**) has markedly higher intensity at one or more points during the period across all dimensions that differ between clusters. In comparison, the larger cluster of 145 disasters (**Cluster 2**), with the exception of *swear*, remains at pre-disaster levels or drops below it. In relation to prior work, we hypothesize that the smaller cluster may correspond to more severely affected communities leading to a strong emotional response and a need to engage in sustained sensemaking (Kohn and Levav, 1990; DeWolfe, 2000). In the following sections, we first investigate if the differences between the two clusters may be explained by intrinsic properties of the disasters, and then, describe the main themes that emerge from the temporal patterns in Cluster 1.

2.5.2.1 Effect of Disaster Characteristics on Community Response

In discussing the findings, we refer to potential relationships between response trajectories and disaster characteristics such as the type, scale, duration, speed of onset, level of forewarning, and familiarity of the disaster drawn from existing disaster research (DeWolfe, 2000; Gaspar et al., 2016; Tierney et al., 2001; Saleem et al., 2014). Next, we evaluate the association 3 of these properties, type, scale, and duration, with the observed clusters of disaster response.

opposed to the culinary practice of smoking

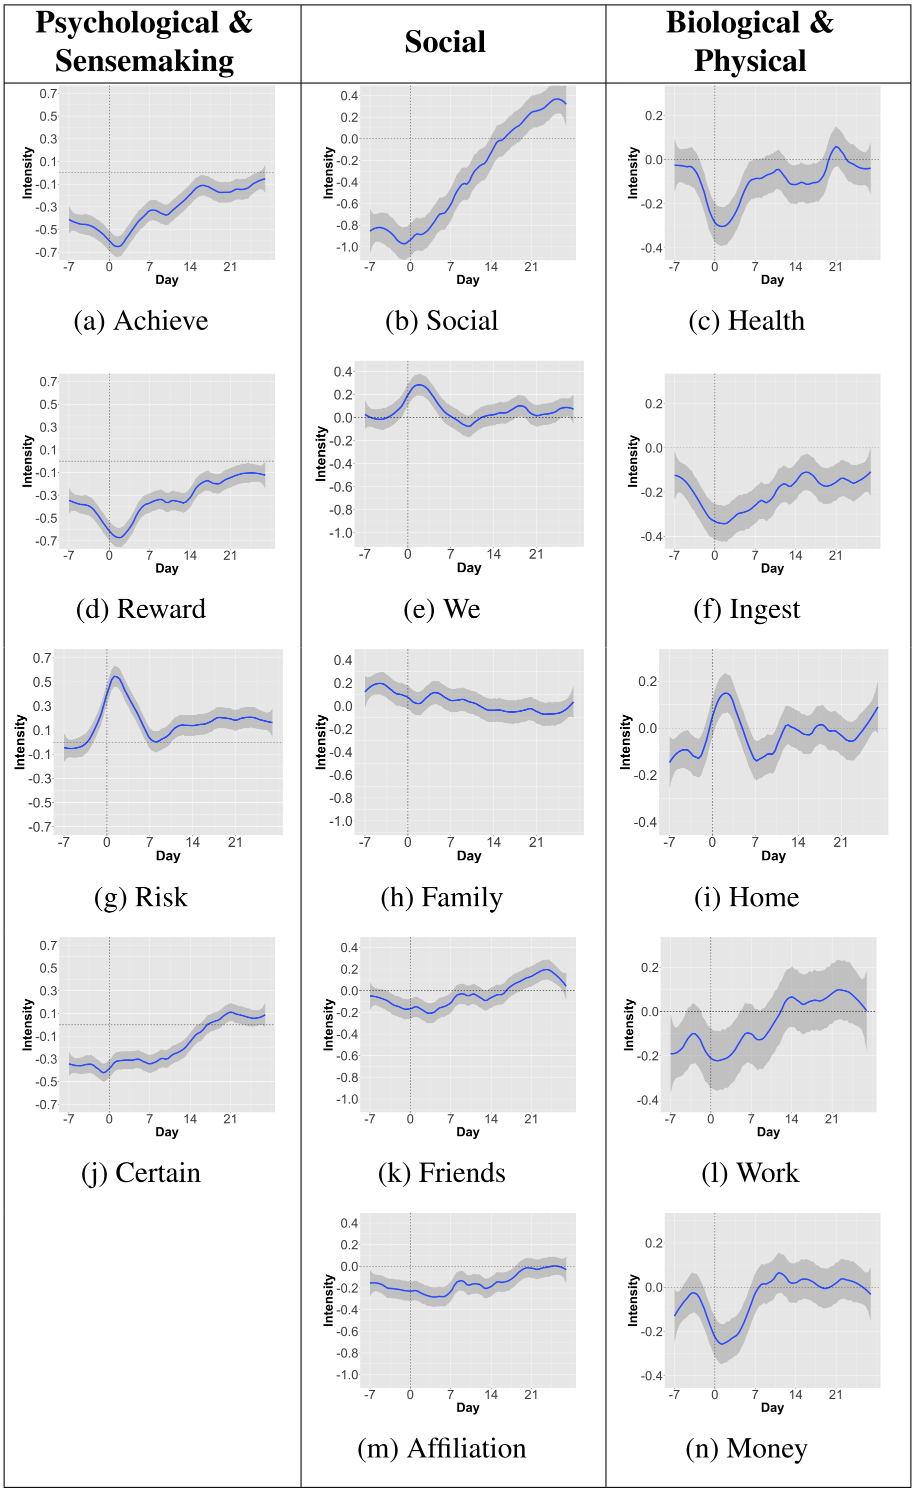

Figure 2.5: Dimensions with universal trajectories across disasters in the dataset. All figures in a column have the same intensity (Y) scale and the shaded regions are 95% confidence intervals

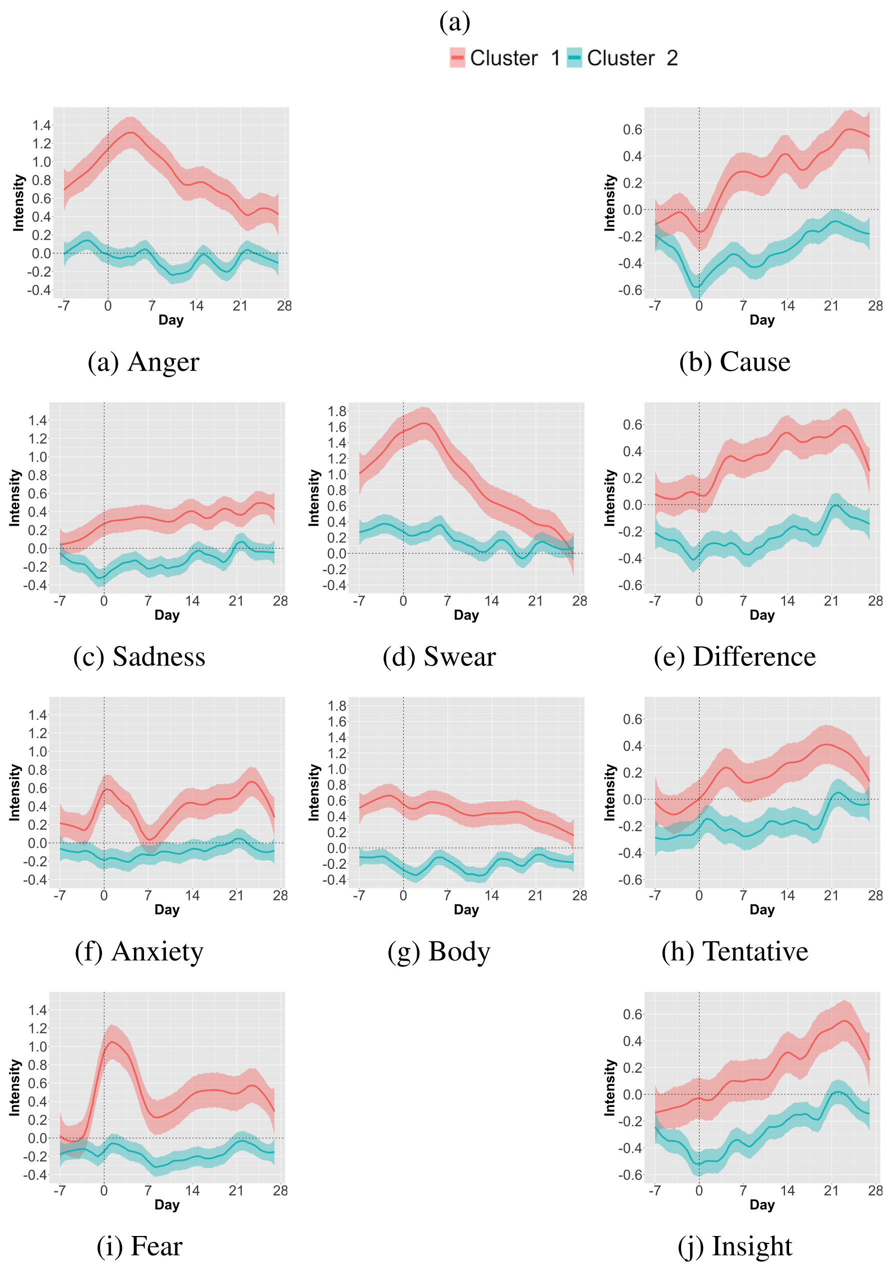

(a) Anger (b) Cause

(c) Sadness (d) Swear (e) Difference

(f) Anxiety (g) Body (h) Tentative

(i) Fear (j) Insight

Figure 2.6: Categories with noticeable differences across the two clusters. All figures in a column have the same intensity (Y) scale. The shaded regions are 95% confidence intervals

We derive the values for each property of a disaster in the following way. First, in terms of the type of disaster, severe storms, fires, and floods account for 85% of the 203 disasters used in this study. Therefore, in order to mitigate the effects of sampling bias, we limit the investigation to only disasters of these 3 types. Cluster 1 contains 24 storms (42.25%), 24 fires (42.25%), and 5 floods (9.5%), while cluster 2 includes 48 storms (40%), 42 fires (35%), and 30 floods (25%). Second, we consider the total amount of federal assistance received by communities through FEMA as a proxy for the scale of each disaster. Finally, we consider the number of days between the start date and the end date of a disaster according to the FEMA declaration as its duration.

We use a series of logistic regressions that consider all different combinations of properties to investigate the association of these properties with the membership of the disaster in cluster 1 or 2. For this purpose, we represent the type as three separate indicator variables (Storm, Fire, and Flood). As shown in Table 2.4, the only statistically significant relationship is between the scale of the disaster and the cluster in the complete model that controls both the duration and the type. This suggests that once the type and duration are taken into account, an increase of 1 standard deviation on the scale increases the odds of a disaster belonging to cluster 1 by 44%. This observation provides some support for our hypothesis based on the observation of trajectories that cluster 1 corresponds to disasters that have a greater impact.

Model	Scale(p)	Duration(p)	Storm(p)	Fire(p)	Flood(p)
Type	-	-	1.07 (1)	1.22 (1)	0.36 (1)
Duration	-	0.88 (0.465)	-	-	-
Scale	1.17 (0.313)	-	-	-	-
Type + Duration	-	0.87 (0.447)	1.05 (1)	1.25 (1)	0.36 (1)
Duration + Scale	1.22 (0.223)	0.83 (0.329)	-	-	-
Type + Duration + Scale	1.44* (0.046)	0.79 (0.218)	1.03 (1)	1.48 (1)	0.29 (1)

Table 2.4: Odds ratios by covariate from logistic regressions predicting the cluster associated with a disaster ($^*p < 0.05$)

2.5.2.2 Theme 1: Fear and uncertainty recur after a delay (Cluster 1)

In Cluster 1, we observe a zig-zag pattern of increase and return to baseline levels in *fear Fig. 2.6i*, *anxiety* (Fig. 2.6f), and *tentative* (Fig. 2.6h). These dimensions initially peak at the onset of disaster or in the immediate aftermath, but quickly return to pre-disaster levels within the first few days, only to increase again almost immediately. Communities maintain an elevated levels of fear, anxiety, and uncertainty in remaining two weeks within the study period. Interestingly, this subsequent peak in the trajectory is higher than the original for *anxiety* and *tentative*, but lower for *fear*. What may be driving this behavior within the communities of cluster 1?

Shortly after the onset of a disaster, members of communities come together in solidarity and experience greater social cohesion, a period sometimes known as the *honeymoon* phase in prior research (DeWolfe, 2000; Raphael, 1986). We suggest that this process temporarily mitigates some of the negative emotional effects of the disaster. The end of the *honeymoon* phase is marked by a decline in attention and support outside of the community, as a well as gradual exhaustion of social support within. During this period, communities are gradually shifting their attention to the long term and realizing that they have many challenges ahead on the road to recovery, which can lead to a resurgence of concern. The fact that expressions of *fear* within these communities don't reach the same levels of intensity as observed at the onset of disaster whereas *anxiety* and *tentative* exceed them, may be simply due to people being more likely to respond to the immediate dangers in the aftermath of disaster with fear. In contrast, we suggest, uncertainty associated with the long-term are more likely to elicit expressions of worry or concern.

2.5.2.3 Theme 2: Communities increasingly engage in sensemaking in the aftermath (Cluster 1)

Cluster 1 communities show a pattern of gradual increase in intensity in the cognitive processing dimensions *cause* (Fig. 2.6b), *difference* (Fig. 2.6e), *tentative* (Fig. 2.6h) and *insight* (Fig. 2.6j), starting from the immediate aftermath of disaster to nearly the end of the study period. For *difference*, *tentative*, and *insight* reverses in the last few days of the study period which may be an early sign that these communities are settling on a *new normal*. The fact that *insight*, associated with expressions of understanding, lags behind the other three dimensions and only peaks in the final week of the study period supports this argument.

We suggest that the presence of elevated levels of cognitive processing in cluster 1 communities, as well as the absence thereof in cluster 2 communities, that these communities experienced greater upheaval and had more difficulty facing the challenges brought on by disasters. This may be at least part due to Cluster 1 being associated with more destructive disasters. However, it also suggests that these communities are more likely to have lacked the institutional, financial, and social resources to weather the aftermath.

2.6 Limitations

We acknowledge some limitations in our work. First, the Twitter data for each disaster may not be representative in terms of demographic or socioeconomic groups within the communities. We begin with a 10% sample of all tweets and then use geotagged tweets to identify affected users. Geotagged tweets represent only about 1% - 2% of all tweets (Sloan and Morgan, 2015). Users who geotag tweets are more likely to be young, more wealthy, reside in urban areas, and belong to certain ethnic groups (Sloan and Morgan, 2015). While our study does not adjust for these biases, the fact that our many of our findings align with existing theoretical models of offline community experience suggests that deviations between different demographic and socioeconomic groups are more likely to be in scale rather than form.

Second, we have limited our analyses in this study to an observation period that extends only four weeks beyond the onset of the disaster. We opted for this short term analysis to manage computational complexity as well as the difficulty of disentangling responses associated within individual disasters when communities, or parts of them, experience multiple disasters or other relevant events over time. Prior work, both theoretical and empirical, suggest that recovery from disasters can take years. It is uncertain if our methodology of measuring Twitter trajectories and the community posts as a source of data are adequately sensitive and precise to measure community behavior over the longer time periods.

Third, in addition to the nature of the hazard, pre-disaster characteristics of communities, such as social and economic capital, inequality, and demographics, influence community resilience to the adversities of disaster (Bonanno et al., 2010; Tierney et al., 2001; Council, 2006). How these community properties interact with those of disasters to produce different recovery patterns is an interesting research question for future study.

Finally, in using aggregate community level estimates for social and affect categories of response, we implicitly assume that a community is monolithic. In reality, members of a community may demonstrate range of responses. There is a risk that the collective behavior observed in our study has been distorted due to the individual responses not being amenable to linear combination. This limitation can be resolved in future work in two ways; (i) conducting the analysis at a more spatially granular level, such as counties and census tract level, where demographic and socioeconomic data is readily available and (ii) employing statistical techniques such as latent growth mixture modeling to automatically extract different common recovery trajectories that play out within a single community experiencing a disaster.

2.7 Discussion

In this study, we demonstrate that community response to disasters over social media as measured by lexicon based time-series follow regular temporal patterns that generalize well to different disasters and communities.

This study demonstrates that lexicon based time-series derived from social media posts of people in disaster stricken communities can effectively capture many theoretically expected phenomena in the collective process of recovery. In particular, we find that there are broad similarities in community social media behavior in terms of some aspects of their lives (Ex: social, biological, aspirational, physical) despite significant variability in terms of the type, scale, duration, level of forewarning etc. However, the results also reveal that around a quarter of the communities demonstrate trajectories that diverge from the majority along some aspects of their experience. In particular, their social media conversations show a sustained levels of fear, anger, and anxiety as well as signs of an elevated cognitive load associated with sensemaking; both signs of ongoing challenges of adapting to post-disaster *normal*. In this section, we discuss the implications of this study and its findings for two domains; (i) disaster relief and recovery efforts and (ii) future disaster research.

2.7.1 Implications Disaster Relief and Recovery Assistance

Strategic Disaster Relief Planning. Relief organizations engaged in disaster recovery assistance face a difficult challenge. To effectively support communities in the aftermath of a disaster, they require two kinds of knowledge. First, at the strategic level, these organizations need to know how, in general, community response to disasters unfolds over time, in order to inform their policies for resource planning for assistance programs and training relief workers. The results of this study show that the behavior of disaster-stricken communities over social media follow a regular trajectory along many aspects of life and that in many cases this prototypical narrative aligns with prior knowledge of *real* (offline) community experience. Further, this comprehensive, but easily communicated, social media narrative is mapped to a more precise timeline than traditional models of community response. Therefore, we believe it can complement the existing suite of planning tools of relief organizations.

Community Response Monitoring and Forecasting. Even though an idealized narrative of disaster response is useful as an organizing principle for support organizations, each disaster is a unique event resulting from a specific hazard affecting a particular community. Thus, the community response to a new disaster follows its own unique trajectory that can deviate from the prototypical trajectory - For example, we show that communities differ in the intensity of their emotional and cognitive response to a hazard as well as the timing of their return to stable post-disaster norms. As a consequence, during the course of an assistance program, relief organizations are faced with the continuous need to reconcile facts on the ground with the narrative to align their activities with the needs of the affected community. Our work furthers support for this endeavor in three ways. First, our methodology provides a rapidly deployable, relatively low-cost, and near real-time approach to get a holistic picture of a community as it responds to a disaster - in the event of onset of a disaster, this would involve collecting easily accessible public social media posts from members of the affected community and generating a response trajectory using our estimation technique that requires only modest computing power. Moreover, since this approach accounts for idiosyncrasies of the community, these observations can be compared directly with the prototypical narrative. This leads to the second benefit - relief organizations can make early forecasts on the recovery process including the characteristic response cluster, areas that may need more or less attention compared to original expectations (e.g. material support, mental health services, information services and social programs), and more precise timing for crucial transitions (e.g. initial shock to the honeymoon period to the long-term perspective). In the supplementary section A.5, we demonstrate a proof of concept analysis to compare the forecasting capability of the response trajectories observed in this study. The results show that simply assigning an emerging community response to one of the clusters based on that community response until the present is able to provide a forecast of its future trajectory that outperforms time series forecasting methods using the same data around 75% of the time.

Evaluating Disaster Relief Interventions. Another challenge that relief organizations face is the difficulty of assessing the impact of their interventions in the aftermath of disasters. Our methodology and findings may be adapted for this purpose in the following manner. First, starting from the onset of disaster and as it continues to unfold, we construct the response trajectory for the affected community. We expect that even an early short-term fragment of this trajectory will contain quantitative markers for idiosyncrasies, such as the intensity of emotional response and the characteristic time to stable post-disaster behavior, stemming from the specific hazard and the specific community. Then, we are able to forecast future community behavior, in the absence of interventions, by adjusting the prototypical narrative to the observed eccentricities. Finally, we would be able measure the effect of assistance as the deviation of actual behavior from the expected after targeted interventions. In practice, this approach is limited to measuring the overall effect of assistance due to numerous overlapping relief programs that are normally active in a community after a disaster, rather than that of individual efforts. Further, work highlights an opportunity for the technology community to address these needs of relief organizations and disaster affected communities through the development of software tools that automate these functions.

Public Perception of Community Response. In addition to formal organizations dedicated to relief efforts, the public, either through community organizations or as individual Samaritans, invariably plays a role in supporting communities stricken by disaster. However, unlike the organizations, which are invested in providing relief and have established information channels regarding the needs of the affected community, the public relies on and is influenced by the freely available, and at times incidentally consumed, information in their every day sources of information. For many, the primary mode of such information is likely to social media (Walker and Matsa, 2021). One of the themes that emerged from the findings of this study is that , similar to what prior work on offline behavior has noted, community response to disasters over social media is a combination of multiple processes unfolding concurrently but not necessarily following the same shape or temporal order. Some processes such as emotional upheaval and community bonding operate on short time scales (days), while other dynamics such as sensemaking last much longer. Some processes, like emotional upheaval and sensemaking, are marked by an increase in their presence in community social media feeds while others, such as social engagement and nourishment, are marked by their absence. Only content that exists on social media can be amplified by it and brought to a wider audience. Therefore, the differences between different processes of community response is also likely to lead to a disparity in the attention of the public on the corresponding community needs.

2.7.2 Implications for Disaster Research

Recent decades have witnessed an increase in the worldwide frequency of natural disasters, to a large extent, driven by climate change (Hagon et al., 2020; Ritchie et al., 2022). Correspondingly, large-scale disasters that cause widespread economic and social damage and less observable effects such as internal displacements have become more common (NCEI, 2020; Coronese et al., 2019). These trends have added some urgency to our need to understand community response to disasters - in particular, around the questions of what makes a community resilient to disasters and how relief operations affect post-disaster recovery. In closing, we discuss the contributions of our study in this context as well as future research directions opened up by this work.

Measurement. It is challenging to study community response and resilience to disasters in a widely applicable and reproducible manner due to the many complex social phenomena that simultaneously unfolding along many interrelated dimensions. Traditional approaches, such as post-disaster surveys, interviews or even ad-hoc analysis based on observational data such as social media posts, are not suited for standardized measurement along these numerous dimensions for a wide range of disasters. Interview and survey data suffer from biases due being collected after the fact and are costly to collect. Further, whether using traditional or social media data, prior work has very often not accounted for idiosyncrasies of individual events - partially because many studies consider only a single event. Thus our most immediate contributions to disaster research are (i) a comprehensive thematic framework for measuring disaster response based on existing work and (ii) a scalable method, built upon that framework, for quantifying the trajectory of community response and that allows direct comparison between disasters.

Universality. In this study, we limit our focus to disasters in the United States over a five year period. A logical next step would be to investigate if our findings will generalize to across different cultures and times. i.e. Will community responses to disasters remain consistent with observations over time? Do we observe the same prototypical narrative of disaster response in communities outside the US in regions such as Asia and Africa?

Resilient Communities and Disaster Assistance. Resilience is an increasingly important topic in both disaster research and policy, as high impact disasters which impose large economic and social costs become more common. A resilient community re-emerges from the crises of disaster in a manner epitomized by the German philosopher, Friedrich Nietzsche's adage - "That which does not kill us makes us stronger" (Nietzsche and Levy, 1909). In the aftermath of disaster, they establish a *new normal* that is both better than the pre-disaster baseline (in terms of greater life satisfaction, social cohesion etc.) and are consequently better adapted to face future crises. Prior work is rich in proposed relationships between community properties, such as social capital and wealth inequality, and disaster resilience. Since our methodology captures the effects of disasters across different communities in a consistent manner, it can serve as the foundation for principled future investigations that relate properties of communities to their resilience to disasters. Further, we expect that it will have a similar utility for analyses that evaluate the effectiveness different assistance activities in improving disaster recovery outcomes of communities.

CHAPTER 3

Influence of Online Social Capital on Relocation Choice after Disasters

3.1 Introduction

Disasters have a multitude of effects on individuals and communities, including psychological and physical trauma, as well as the dismantling of social structures (Bonanno et al., 2010; Norris et al., 2002). In the aftermath of disasters, affected people face the choice of rebuilding in situ or relocating. The literature on migration, and especially forced relocation, offers a number of perspectives on the factors that contribute towards the decision to relocate (or not).

Early theories of migration frame the choice to relocate as an economic cost-benefit analysis. However, recent work paints a more nuanced picture with a combination of push factors that make people want to move away from a region and pull factors that attract them to new locales (Lee et al., 2017). Among factors that drive people away from a region, natural disasters feature prominently in addition to poor living standards, unemployment, conflicts/war, etc. while superior incomes, education, healthcare, urbanization, disaster safety, etc. are qualities that retain current residents and attract new residents.

Although natural disasters are exogenous shocks that exert a strong outward push from the affected region, this effect is moderated by several other factors related to the specific individual and the resources available to the broader community. Theoretical frameworks have suggested that an individual's social capital, i.e. the unique resources afforded by their relationships with a social circle, can influence their ability to cope and recover from disasters in place, or alternatively how successfully they may be able to relocate (Aldrich, 2012a; Casagrande et al., 2015; Norris et al., 2008). In this regard, most prior empirical work has employed coarse-grained measurements of community-level social capital, such as the presence of community organizations and religious affiliation (Karunarathne and Lee, 2019; Kyne and Aldrich, 2020), or simplistic measures derived from interviews and surveys conducted with small samples in the aftermath (Engel and Ibáñez, 2007; Lee et al., 2017). The primary limitation of these approaches is that they do not adequately characterize the specific nature, strength, or structure of interaction between individuals and their social circle. As a consequence, the rich literature on different forms of network social capital, such as bonding and bridging capital, distinguished by distinct structural mechanisms (Burt, 2000, 2007), remains underutilized in the work that characterizes the relocation choice and, more broadly, resilience, in the aftermath of disasters.

Social media provide an opportunity to address this limitation, as they capture detailed information about the extent and timing of how people interact with their online social circles. Not only would data from social media platforms such as Twitter and Facebook allow quantifying the pre-disaster levels of different types of social capital, they can be utilized to understand how disasters affect social capital. Recent work that has used social media data to characterize social capital and its effect on disaster resilience, coping, and evacuation choice (which is different from relocation) demonstrates the potential of this approach (Kaigo, 2012; Metaxa-Kakavouli et al., 2018; Page-Tan, 2021).

In our work, we extend this quantitative approach to assess the relationship between network social capital, as measured based on Twitter interactions of users affected by disasters, with relocation choice in the aftermath of disaster. While building on the existing body of theoretical and empirical work on social capital in the context of disasters, our work improves on this literature in the following ways. First, despite the potential relationship between evacuation and relocation, relocation is a distinctly long-term and more important phenomenon in terms of the effect it has on the trajectory of a person's life. Second, in contrast to prior quantitative work, we will use a large number of disasters of different types and scales that occurred in different regions of the United States during 2014 and 2015. This allows us to account for and interpret the confounding effects of the properties of the disaster, such as type, scale, and duration, as well as more generic and coarse-grained resilience characteristics of the affected communities, such as the push factors discussed earlier.

3.2 Background

In this section, we first introduce two different forms of social capital (**bonding** and **bridging**) that we consider in this study. We discuss the social mechanisms from which they arise, the types of social support they may afford, and how they are reflected in the structure of a social network. Then, we discuss network robustness, and it may be important in the context of leveraging social capital, particularly in the aftermath of an environmental jolt. We follow this with a discussion of current knowledge on how different forms of social capital may influence an affected person's decision to relocate after a disaster. Finally, we briefly discuss online forms of social capital, how they may differ from their offline counterparts, and their implications for how online social capital may be tied to relocation choice.

We also briefly introduce a new form of social capital, **linking** capital, that social scientists have recently proposed Aldrich (2012a); Magis (2010). As with the other forms of social capital, we discuss how this form of capital may influence relocation behavior in the following sections. Since the nature of this form of social capital makes it challenging to quantify using only structural features of a person's network, we leave the study of how this form of capital affects relocation choice after a disaster out of the core study. Instead, in an extension to it, we study the influence of this form of capital,as well as the other two forms, by measuring the prevalence of different types of relationship estimated using a machine-learning classifier.

3.2.1 Forms of Social Capital

3.2.1.1 Bonding Capital

Bonding capital arises from relationships between people that are "similar to each other in many important respects" (Putnam, 2000). Dimensions of shared identity for such ties include gender, ethnicity, age, and social class among others, and are commonly observed within family units and in strong friendships. From a structural perspective, these ties are associated with high levels of closure, in that if a person shares bonding ties with two others, those others are also more likely to share a bonding tie since they too are likely to be similar to each other. This means that a bonding tie is generally well embedded in a cluster of similar ties in a network that allows the effective use of social pressure and sanctions (Coleman, 1988; Putnam, 2000; Sabatini, 2009). As a result, bonding ties are associated with higher levels of trust, reciprocity, and collaboration. However, also as a consequence of closure, these ties are less useful to a person in terms of access to new information. As everyone in a cluster of embedded bonding ties is likely to share a common identity in many dimensions, they are likely to be in possession of the same information (Coleman, 1988; Granovetter, 1973). Given the structural characteristics associated with bonding ties (density and embedding), network measurements such as tie strength, tie density, and clustering are considered effective measurements of bonding social capital (Crossley et al., 2015; Metaxa-Kakavouli et al., 2018).

3.2.1.2 Bridging Capital

Compared to highly embedded bonding ties, *weak ties* between people who only interact occasionally and share only a few similarities, such as organizational affiliation, can bridge sparsely linked parts of the larger social network (Granovetter, 1973; Putnam, 2000; Burt, 2007). As a consequence of this sparse linkage, these *bridging* ties are not useful in enforcing social norms or in engendering high levels of cooperation. However, they are useful for accessing resources, in particular information, from disparate parts of the social fabric outside a person's local network (Granovetter, 1973). Since weak ties, which reflect only occasional interaction, are less effortful to maintain than strong ties, the majority of an average person's social circle is likely to comprise of weak ties.

Based on the properties of weak ties, the number of non-redundant ties that a person has to their social circle (Crossley et al., 2015; Burt, 2007) and the size of the network (Arnaboldi et al., 2016; Ellison et al., 2007) are natural measurements of bonding social capital.

3.2.1.3 Linking Capital

Linking capital refers to the resources that a regular individual can access through his or her connection to a person or entity with power or authority (Aldrich, 2012a; Kyne and Aldrich, 2020). Consequently, unlike the other two forms of social capital, which are horizontal in nature, linking capital corresponds to a vertical power dynamic (Kyne and Aldrich, 2020). In practice, it is challenging to discover the amount of linking capital embodied in a person's online relationships purely from structural measures of their network, because the key fact that differentiates this form of capital from the others is more related to the qualitative nature of the relationships. However, as seen in Metaxa-Kakavouli et al. (2018), the qualitative nature of a relationship (i.e., the type of relationship) can be used to reliably differentiate linking capital from other forms of social capital.

3.2.2 Network Robustness

The structural robustness of a network refers to the extent to which a network may be able to retain its connectivity and high-level structural features in the event of damage (loss of members and links between them) [1] (Iyer et al., 2013; Trajanovski et al., 2013; Dunne et al., 2002). In the context of social networks, structural robustness is related to, but is distinct from, social capital. Measures of structural robustness emphasize high network density and redundant linkages as strengths and uneven distribution of connectivity, particularly connectivity dependent on highly central nodes, as signs of fragility (Dunne et al., 2002; Sydney et al., 2010; Trajanovski et al., 2013). These definitions suggest a certain degree of positive (negative) correlation between robustness and bonding (bridging) capital.

An external shock such as a disaster can disrupt an affected individual's social circle, as well as the broader social fabric of the community. In effect, a disaster may result in, at least temporarily, the loss of some of its members and/or links between them. In this scenario, the structural robustness of the pre-disaster social network can provide an indication of how successful the person (or

[1] In fact, there are two interrelated types of structural robustness; (1) Robustness of connectivity to loss and (2) Robustness against cascading failures. In this study, we refer primarily to the former, because our context being social and economic networks. The latter is more relevant in technological and infrastructure networks, among others

community) may be in leveraging the social resources embedded in their regular social network in the aftermath.

3.2.3 Social Capital, Disaster Resilience, and Relocation

In the aftermath of a disaster, communities and their members utilize their personal capabilities as well as the resources available through social relationships to adapt to the new circumstances and begin the journey towards recovery. The availability of social resources, that is, social capital, from different sources has been shown to improve individual and collective outcomes after a disaster (Aldrich, 2012b; Murphy, 2007; Hsueh, 2019; Delilah Roque et al., 2020; Lee et al., 2022). Access to bonding capital provides individuals and communities with established and trusted pathways to coordinate the initial response to disaster, including rescue and relief operations. They also ensure that community members take economic risks that can lead to better long-term outcomes (Aldrich, 2012b; Chamlee-Wright and Storr, 2011). However, bridging ties tend to improve an individual's chances of obtaining relief supplies, access to information, and psychological support (Hsueh, 2019; Panday et al., 2021; Delilah Roque et al., 2020). Interestingly, in addition to providing access to different types of resources, the two types of capital show differences in their availability in the aftermath of disasters (Lee et al., 2022). Bonding capital arising from close ties, who may also be vulnerable to the effects of the disaster affecting a focal user, seems to decline or be exhausted in the short term. In contrast, bridging capital, which is more distributed within the network and potentially geographically as well, continues to be available much longer. Access to high-levels of bonding capital in the short term and high-levels of bridging capital in the longer term seems to result in the most resilient individual and community outcomes after a disaster.

In the aftermath of disasters, linking capital ties, particularly to government organizations responsible for disaster relief, can be crucial to ensure that communities receive adequate external support (Pfefferbaum et al., 2017; Aldrich, 2012a; Kawamoto and Kim, 2019). Therefore, we expect that access to linking capital within a community would lead to more resilient responses to the effects of disasters. On a different note, in the related scenario of evacuations due to hurricane warnings, Metaxa-Kakavouli et al. (2018) found that people with more linking capital, measured by the number of politicians followed on Facebook, are more likely to promptly evacuate than others. They concluded that people with linking capital were more likely to be exposed to hurricane warnings, which are usually disseminated by government sources, and that people would place more trust in warnings that they received from through their linking ties.

3.2.3.1 Influence of Social Capital on Individual Relocation Decision

The relationship between the social capital available to an individual and their relocation decision in the aftermath of a negative shock is more complex than its relationship to resilience. While the influence of social capital on relocation is mediated by resilience, at an individual level, which relocation outcome (i.e., staying or leaving) is resilient is a matter of perspective. When considering their relocation, people are likely to weigh options with respect to the comparative impact on their subjective well-being (Engel and Ibáñez, 2007). Therefore, an individual who chooses to relocate because their social capital, both inside their community and outside, ensures a convenient and profitable transition may be considered just as resilient as an individual who chooses to stay in the community because their social ties have provided them the resources to weather the shock and bounce back quickly. Prior work that has studied the impact of social capital on relocation due to an external shock is limited (Lee et al., 2017; Engel and Ibáñez, 2007; Airriess et al., 2008).

Taking into account the relationships between social capital and resilience, as well as the few studies that focused on relocation, some general observations can be made on how different types of social capital can influence relocation. These observations will rely on a key assumption regarding the correlation between bonding/bridging ties and geography. In short, given that maintaining ties takes effort and assuming that this is more convenient and efficient in person, we argue that bonding ties are more likely to be geographically local while bridging ties are less constrained by geography. Extending this argument, some prior work has reasoned that bonding capital, being more localized within the community, would be more likely to make an affected individual stay after a disaster both by supporting their recovery and also by representing an opportunity cost of relocation (Metaxa-Kakavouli et al., 2018). However, other studies on relocation have observed the opposite behavior, i.e. having more bond capital increases relocation in the aftermath of a shock (Engel and Ibáñez, 2007). They provide an alternative mechanism to explain this observation by focusing on a specific subtype of bonding ties, stable long-term relationships within the neighborhood such as neighbors and family. These relationships are likely to be more resilient to the effects of distance, while also being more invested in supporting their transition during relocation compared to more generic but still important ties. Compared to a more generic definition of *strong ties* that are characterized by high activity and position within the current structure of the ego-network, this mechanism highlights the cumulative value of relationship stability derived from contact over a long period or kinship.

Bridging capital, which is more likely to be distributed outside the community, would provide access to resources and information to make potential relocation more attractive. Similar reasoning has been used in recent prior work that considered the effects of social capital on evacuation behavior prior to a hurricane (Metaxa-Kakavouli et al., 2018). We formally state this reasoning as hypotheses in the next section.

3.3 Hypotheses

In this section, we describe the main hypotheses that are evaluated on the effects of social capital and network robustness on relocation choice after disaster.

H1 We expect that having more bonding capital will provide affected individuals with more local resources for recovery. However, these resources may influence in relocation choice through opposing mechanisms depending on the stability of the corresponding relationships

(a) We expect that bonding capital, without considering the stability of contributing relationships, would be a source of local support and an opportunity cost of relocation. **i.e., overall undifferentiated bonding capital will make relocation less likely**

(b) We expect that stable strong tie relationships characterized by continued engagement over time to be resistant to withering due to relocation and provide greater support during relocation. **i.e., bonding capital associated with stable long-term relationships will make relocation more likely**

We expect that having more bonding capital will provide resources for recovery within the community, encouraging people to stay. **i.e. bonding capital will make relocation less likely**

H2 We expect that having more geographically distributed bridging capital will give people access to information regarding opportunities outside the community as well as resources to make the transition smooth. **i.e. bridging capital will make relocation more likely**

H3 We expect that structural robustness of a person's network will improve their ability to leverage the resources of their network. **i.e. network robustness will enhance the effects of social capital on the likelihood of relocation**

3.4 Data

In this section, we describe (i) the set of disasters being used in this analysis, (ii) identification of Twitter users affected by a disaster, (iii) the collection of geotagged tweets and mentions for affected users, and (iv) the collection of tweets to establish the pre-disaster ego-network of these users.

3.4.1 Disasters

We rely on disaster declarations made by the US Federal Emergency Management Agency (FEMA) to identify the disasters used in this study. In particular, we focus on the disasters that occurred during 2014 and 2015. During this period, FEMA made a total of 153 declarations. We exclude two declarations from our analysis since they correspond to man-made disasters whose impact on communities was comparatively long-term and low intensity [2] leaving 151 declarations. Each FEMA declaration includes the following information about a disaster; (i) type of disaster (Flood, Fire, Storm, etc.), (ii) affected counties, (iii) disaster start date, (iv) disaster end date, and (v) amount of FEMA aid granted. Subsequent Twitter data collection, described in the following section, discovered potentially affected users for only 118 declarations. Table 3.1 shows statistics for these declarations aggregated by type of disaster.

Type	**# Events**	**FEMA Aid (Mean)** 10^6	**Duration (Mean)** Days
Severe Storm	49	24.76	13.37
Fire	32	0.74	9.97
Flood	16	21.12	16.00
Severe Ice Storm	9	39.72	6.44
Snow	3	12.69	7.00
Coastal Storm	1	0.43	1.00
Mud/Landslide	1	15.04	38.00
Earthquake	1	38.83	15.00
Tornado	1	10.10	2.00
Volcano	1	11.19	203.00
Other	1	22.46	3.00

Table 3.1: Disaster Statistics

3.4.2 Affected Twitter Users

Ideally, we would want to observe users directly exposed to the disaster and who are permanent residents of an affected area. In this study, we use a two-step approach to identify candidate users while minimizing unnecessary data collection. First, we use a 10% archive of Twitter data to identify individuals who posted at least one geotagged tweet from a county affected by the disaster during the two-week period centered on the date the disaster began. These individuals are more likely to have been directly exposed to the effects of the disaster. However, this first step doesn't

[2] The removed declarations correspond to the2014 Elk River Chemical Spill in West Virginia and the Flint Water Crisis in Michigan

account for the possibility that some of these individuals may not be permanent residents of the affected region. We address this limitation with the second step. We use the Twitter API to collect all geotagged tweets from candidate users for the 12 weeks immediately preceding the onset of the disaster (**baseline period**).

Next, we identify a home location for each candidate user as the county subdivision from which they produced geotagged tweets on the most days during the baseline period [3]. To limit the effects of sparsity of geotagged data on the reliability of these estimates, we only assign a home location to individuals if they satisfy two conditions. First, they should have geotagged tweets on at least five separate days during the period, and second, they must be observed at their most frequent county subdivision at least three separate days during the same period. Prior work has shown that variants of this approach can reliably estimate coarse-grained home locations for people using different types of geolocation data, including mobile phone Call Detail Records (CDR), GPS traces, and geotags from social media posts such as tweets (Calabrese et al., 2011; Kreindler and Miyauchi, 2021; Lin and Cromley, 2018). We were able to identify 138,233 Twitter users who resided in disaster-stricken communities in 118 FEMA declarations. Finally, we collected all geotagged tweets of these users in the 36 weeks after the onset of the disaster to identify their movements during that period.

3.4.3 Affected User Ego-networks

On Twitter, the ego-network, or social circle, of a user can be defined broadly as their friends, followers, and any other users they have interacted with through mentions. However, due to limitations with the Twitter API, it is not possible to uniquely identify friends and followers of a particular user at a given point in the past. Therefore, we limit our analysis to users with whom the affected user is observed to have engaged through reciprocal mentions. The parts of the ego networks left out correspond to passive connections that may be informational in nature. As we employ interaction-based measurements of social capital in this study, we argue that our observations are not biased by this choice.

Further, in order to make data collection more tractable, we limit the ego-network of an affected user to those other users (alters) that they mentioned and were, in turn, mentioned by them at least 5 times each over the pre-disaster period. We collect Twitter data to build the interaction network (i.e., ego-network of the focal user) of the alters selected in this manner as follows. First, for each affected user, we collect tweets that mention other users during the 24 weeks immediately

[3] County subdivisions are an intermediate level of administrative boundary in the United States. It is more granular than counties and is coarser than census tracts. For the purpose of this analysis, we used geospatial datasets acquired from the US Census Bureau for 2016, which, for comparison, contained 3,233 counties (or statistical equivalent entities), 36,639 county subdivisions, and 74,134 census tracts.

preceding the disaster. From among the set of users observed in this set, we select those who were mentioned by the focal user at least five times during the **baseline period**. Then, we collect tweets from these users that mention the focal user during the same period and remove users that did not mention the focal user at least five times. Finally, we collect all mentions between these selected alters during the same period.

3.5 Methods

In this section, we first describe how the Twitter ego-network is constructed using the collected data for each affected user we have identified. Second, we define our measurements to estimate bonding and bridging social capital, as well as robustness, using this network. Then we introduce two types of control variables; (i) intrinsic properties of disasters and (ii) community-level resilience variables. Finally, we describe the design of our core regression model to evaluate the relationship between social capital of individuals and the likelihood of relocation in the aftermath of a disaster, as well as a series of extensions to validate the robustness of the results.

3.5.1 Affected Ego Network

For each Twitter user associated with a disaster in our dataset, we construct an ego network for the 12 weeks preceding the onset of disaster. The edges of the network are directed and weighted by the number of mentions between users during this period. When using this network to estimate measurements of social capital that rely on relationships between alters, we exclude connections between the focal user and the alters as is standard for ego-network analysis.

3.5.2 Measures of Twitter Social Capital

In this section, we describe a series of social capital metrics, based on prior work in traditional and online settings, that are estimated for each affected Twitter user based on their interactions with other users, as well as the interaction network between those users.

Bonding Capital

1. **Density**: The density of a person's ego-network is an indicator of its connectivity. A higher density indicates a better connected ego-network which is associated with bonding capital. While network clustering and density are related, density captures network connectivity in a more general sense than clustering. A network with several tightly knit clusters of alters

that are not connected to each other can have moderate or high clustering, whereas density is more closely related to the overall connectedness of the network.

2. **Median Coreness**: The coreness of a network node is the value of **k** such that the node belongs to a **k-core** in the network where a k-core is a community of nodes such that each node is connected to at least k other nodes in the community. In this sense, the coreness of an alter within an ego-network reflects how embedded they are within it. Therefore, we consider the median coreness of the network to be an indicator of its cohesiveness.

3. **Strong Tie Retention**: Two issues make identifying significant relationships (i.e., strong ties) in a consistent manner across Twitter users purely using the number of interactions. First, users can have widely different levels of activity. Therefore, it would not be effective to use the same threshold for the number of interactions to distinguish between the *strong* and *weak* ties across all users. Instead, we define *strong* tie alters as individuals who are in the top X% in terms of the number of mentions by the focal user during a given period. Second, a person's activity level on Twitter, as well as the alters they engage with, is known to vary with time (Myers and Leskovec, 2014). This means that, by using our threshold-based criterion, the set of alters we would identify as being *strong* ties may also vary with the chosen time period. This outcome does not quite align with the sociological understanding of significant relationships which are perceived to be steady over long periods of time (Putnam, 2000; Casagrande et al., 2015). Therefore, we explicitly include the stability of interaction in our measurement of tie strength-based social capital. For a given user, let the set of strong ties from the 12 week period immediately before the pre-disaster baseline period be S and let $\{n_i \ where \ i \in N\}$ be their social social circle in the baseline period. Then, the retention of strong ties (r) is measured as follows.

$$r = \frac{\sum_i \delta(n_i \in S)}{|S|}$$

Bridging Capital

1. **Network size**: The alters in a person's ego-network include both strong and weak ties. However, since a person has only a finite amount of attention available and they focus disproportionately on important relationships, the network is made up of a small number of strong ties that correspond to a large fraction of their interactions and a much larger number of weak ties. Therefore, we consider the size of the ego-network as a proxy for bridging social capital.

2. **Structural Diversity**: A person's bridging capital is related to the extent to which they can

reach or leverage different social groups or communities that are not connected to each other. This argument relies on the idea that the resources and information available to tightly knit communities (bonding capital) are bounded and that an individual connected to several such groups can potentially receive information and support from all of them. We use the number of connected components in the **undirected** ego-network (i.e. weak components) of a focal user as a measure of the different sources of support available to them. Although the number of strong components (that is, the number of connected components in the directed network) is the more appropriate measurement for structural diversity (Ugander et al., 2012), we choose weak components for two reasons. First, due to the relatively sparse nature of Twitter mention networks, the number of strong components is strongly correlated with the size of the network ($r \approx 0.9$). Therefore, including both the size of the network and the number of strong components introduces a high level of colinearity regression models, making the effect estimates reliable. Second, we argue that observing mentions in one direction is a sufficient constraint to consider two Twitter users as connected when they both have reciprocal mentions with a third (focal) user.

3.5.2.1 Network Robustness

In addition to measuring different forms of online social capital reflected in the Twitter ego-networks of affected individuals, we also include multiple measures that capture how robust these networks would be to partial loss of communication due to the effects of disasters [4]. Structurally robustness networks would allow affected individuals to continue to channel the support of most of their ego-network in the aftermath of a disaster despite some loss of connectivity.

1. **Tie Strength Dispersion**: Individuals do not interact with their social circles (online or physically) in a uniform way. In general, most of their attention is taken up by a small fraction of their network. Therefore, it is likely that their ability to channel the resources of their network in the event of a crisis is more dependent on those ties. We measure the extent of this reliance (vulnerability) in terms of the smallest percentage of alters in their ego-network that would account for at least 90% of their mentions in the pre-disaster period. We consider this form of robustness to be relevant for social capital measures that are related to the size of a person's social circle. i.e., network size and structural diversity.

2. **Degree Inequality**: In a similar sense to the dispersion of tie strengths between the focal user and their social circle, the interactions between the people in that social circle can be

[4] Loss of communication, which in terms of the ego-network means the removal of nodes (alters) and/or edges (relationships), may be due to a range of reasons including the parts of the individual's social circle also being affected by the disaster, immediate and overwhelming demands on the time and attention the individuals, and failures in communication infrastructure

less than evenly distributed. If a small fraction of central alters in the network are involved in a majority of the linkages, then that network would be more vulnerable to disintegration if they are removed. Therefore, if those central members were to lose communication or not be available in the aftermath of disaster, it would limit the ability of the ego-network of an affected individual to effectively support them. We measure the concentration of connectivity in an ego-network as the Gini coefficient of the degrees of the alters. Assuming that a disaster disrupts the interaction between an affected person and each of their connections randomly, a social circle where connectivity is more concentrated around a few central individuals (i.e., high Gini) will remain more accessible and cohesive in the aftermath. We consider this type of network robustness to be relevant to social capital measures that are related to the level of connectivity within the ego-network. i.e., network density and median coreness.

3.5.3 Disaster Properties

The influence of a disaster on the choice to relocate among members of the affected community can vary depending on its intrinsic properties. For example, certain types of disasters (e.g. wildfires) may be more likely to cause damage to homes. In this study, we consider the type, duration, and scale (measured in terms of the amount of FEMA aid granted) as disaster properties that could influence relocation decisions.

3.5.4 Community Resilience Variables

The resources available to an affected person are not limited to their own and those of their social circle. They also have access to physical and social resources from their neighborhood and the broader community. Local organizations or emergent recovery groups may provide information, physical aid (e.g. food and shelter), and crucial emotional support in the aftermath of disasters. In addition, some communities may have access to better infrastructure (e.g. healthcare facilities and staff) or be more wealthy. Since these resources are likely to affect individual decisions to relocate or remain in the affected community and vary by community, we include a list of 20 *community resilience* indicators compiled by FEMA as part of the Resilience Analysis and Planning Tool (RAPT) program[5] as control variables in this study (Edgemon et al., 2020). The list includes 11 *population-focused* indicators and 9 *community-focused* indicators, as shown in Table 3.2. The population-focused indicators are point estimates for the distribution of certain individual attributes (e.g., age, income, employment) within communities, whereas community-focused indicators capture the qualities of the collective environment (e.g., civic organizations, healthcare) that

[5] The online tool and other resources are found at `https://rapt-fema.hub.arcgis.com/`

affect overall resilience of the community. While the RAPT data do not contain indicator measurements at the county subdivision level, 12 can be either precisely estimated from census tract data or are available directly. The remaining eight indicators, a majority of them community-focused indicators, are only available at the county level.

Indicator	Focus	Data Availability
% Population without Health Insurance	Population	Subdivision
% Population Unemployed	Population	Subdivision
% Population without a High School Education	Population	Subdivision
% Population with a Disability	Population	Subdivision
% Population without Access to a Vehicle	Population	Subdivision
% Population with Home Ownership	Population	Subdivision
% Population over 65	Population	Subdivision
% Population of Single-Parent Households	Population	Subdivision
% Population with Limited English Proficiency	Population	Subdivision
% Median Household Income	Population	Subdivision
% Income Inequality (Gini)	Population	Subdivision
# Civic/Social Organizations (per 10,000 pop.)	Community	County
# Hospitals (per 10,000 pop.)	Community	County
# Medical Professionals (per 1,000 pop.)	Community	County
% Population who are Religious	Community	County
% Mobile Homes in Housing	Community	Subdivision
# Public Schools (per 5,000 pop.)	Community	County
% Net Migration	Community	County
# Hotels/Motels (per 5,000 pop)	Community	County
% Vacant Rentals	Community	Subdivision

Table 3.2: Community resilience indicators from the FEMA Resilience Analysis and Planning Program (RAPT) used as controls

3.5.5 Identifying Relocation

In order to identify if a particular affected individual has relocated, we compare their home county subdivision during the baseline period to newly estimated home location in the 12 weeks immediately after the disaster. Home county subdivision after the disaster is estimated in the same manner as the baseline home location using the approach described in section 3.4.2. If, for a person p, their home location during the baseline period is $H_{p,A}$ and their home location after the event is $H_{p,B}$, then our relocation outcome variable (R_p) is defined as follows.

$$R_p = \begin{cases} 1, & \text{if } H_{p,A} \neq H_{p,B} \\ 0, & \text{otherwise} \end{cases}$$

3.5.6 Regression Models

We use a series of increasingly complex mixed-effects logistic regression model to establish how an affected individual's pre-disaster online social capital is related to their decision to relocate after a disaster. Let β be the overall intercept that reflects the base rate of relocation, **Social(p)**, the social capital features for person **p**, $r(p)$, the network robustness, **Comm(p)**, the set of community properties, **Dis(p)**, the set of disaster properties, **Dis.Type**, the type of disaster, and ϵ, the residual error. We include **Dis.Type** as a random effect in all models. In estimating the models, we excluded individuals for whom one or more social capital measures could not be calculated meaningfully. For example, if a person's Twitter ego-network, as defined by our activity constraints, contained only one person or no one at all, most social capital measures would be undefined for them. Slightly over 88% (121,884) of Twitter users identified as residents were excluded from the regressions in this way, leaving approximately 16,000 users.

The estimated series of models are as follows.

- M1 Base model with only the bonding and bridging capital covariates as predictors. i.e., $R_p = \beta + B_s\text{Social(p)} + 1|\text{Dis.Type(p)} + \epsilon$
- M2 Model with only interactions of social capital covariates with the robustness covariate. i.e., $R_p = \beta + B_s\text{Social(p)} + B_r\text{r(p)} + B_{s,r}\text{Social(p):r(p)} + 1|\text{Dis.Type(p)} + \epsilon$.
- M3 Model with social capital covariates and disaster properties as predictors. i.e., $R_p = \beta + B_s\text{Social(p)} + B_d\text{Dis(p)} + 1|\text{Dis.Type(p)} + \epsilon$
- M4 Model with interactions of social capital covariates with robustness and disaster properties as predictors. i.e., $R_p = \beta + B_s\text{Social(p)} + B_r\text{r(p)} + B_{s,r}\text{Social(p):r(p)} + B_d\text{Dis(p)} + 1|\text{Dis.Type(p)} + \epsilon$
- M5 Model with social capital covariates, disaster properties, and community properties as predictors. i.e., $R_p = \beta + B_s\text{Social(p)} + B_d\text{Dis(p)} + B_c\text{Comm(p)} + 1|\text{Dis.Type(p)} + \epsilon$
- M6 Model with interactions of social capital covariates with robustness, disaster properties, and community properties as predictors. i.e., $R_p = \beta + B_s\text{Social(p)} + B_r\text{r(p)} + B_{s,r}\text{Social(p):r(p)} + B_d\text{Dis(p)} + B_c\text{Comm(p)} + 1|\text{Dis.Type(p)} + \epsilon$

3.5.7 Robustness Checks

In order to validate the robustness of the relationships we may observe between online social capital of people, as well as other predictors, with their relocation choice, we also estimated two types of variants of the models introduced in section 3.5.6.

Outcome We used the following two alternative definitions for estimating relocation.

1. Distance (Log) between the (centroids of) pre-disaster and post-disaster home county subdivisions
2. Rank distance (Log) of the post-disaster home county subdivision from the pre-disaster one.

These alternatives are less sensitive to the uncertainty associated with the estimated home locations due to people's movements not conforming precisely to county subdivision boundaries, which our unit of spatial aggregation.

Post-disaster Observation Period There are two reasons why a post-disaster observation period of 12 weeks may not be adequate. First, due to the sparsity of geotagged tweets, we may be unable to reliably locate a home location for some users, and they would need to be discarded from the analysis. This may cause the downstream analysis to be less representative in important ways, such as in terms of socioeconomic attributes (Sloan and Morgan, 2015) or the severity of personal impact from the disaster. Second, affected individuals may not be able to relocate immediately in the aftermath of disaster. By limiting the post-disaster period to 12 weeks after the event, we may be either discarding or incorrectly identifying people who take longer to relocate. In order to account for these issues, we replicate the regression analysis using post-disaster durations of 24 and 36 weeks.

3.6 Results

In the following sections, we interpret the estimated models with respect to the main hypotheses of the study.

3.6.1 Bonding Capital (H1)

The results provide some support for our hypotheses regarding the different ways in which forms of bonding capital may influence a person's relocation choice in the aftermath of disaster. From our two measurements for undifferentiated overall bonding capital, **median coreness** supports

	M1	**M2**	**M3**	**M4**	**M5**	**M6**
Strong Tie Retention	7^{**}	7^{**}	7^{**}	7^{**}	7^{**}	7^{**}
Median Coreness	-9^{**}	-12^{**}	-8^{**}	-12^{**}	-8^{*}	-12^{**}
Density	5	9	6	11	5	-5
Network Size	9^{*}	9	7^{*}	8	8^{*}	8
Structural Diversity	4	7^{*}	4	8^{*}	5	8^{**}
Tie Strength Dispersion		-8^{***}		-8^{***}		-8^{**}
Degree Inequality		-2		-1		-3
Coreness: Degree Inequality		1		2		2
Density:Degree Inequality		6		7		5
Size:Tie Strength Dispersion		1		-1		-1
Diversity:Tie Strength Dispersion		1		1		2
Disaster Scale			4	5	18^{***}	18^{***}
Disaster Duration			21^{***}	21^{***}	23^{***}	23^{***}
Median Income					24^{***}	24^{***}
Income Inequality					1	0
% over 65					-3	-4
% Disability					4	3
% No High School					18^{**}	17^{**}
% Unemployed					2	2
% Uninsured					-21^{***}	-21^{***}
% Not English Proficient					-25^{***}	-26^{***}
% Home Ownership					-17^{***}	-17^{***}
% Single Parent					9^{*}	9^{*}
% No Vehicle					-8^{*}	-8^{*}
% Vacant Rentals					8^{**}	8^{**}
% Mobile Homes					-11^{**}	-10^{**}
# Civil Orgs.					10^{***}	9^{***}
# Hospitals					2	2
# Medical Professionals					-4	-4
% Religious					-3	-3
# Schools					12^{***}	13^{***}
# Hotels					-6^{*}	-6^{*}
% Net Migration					-9^{***}	-9^{**}
Intercept	14^{***}	14^{***}	14^{***}	14^{***}	13^{***}	13^{***}
Observations	16,372					
BIC	12,566	12,608	12,478	12,520	12,435	12,480

Table 3.3: Effects of Social Capital on Relocation Choice. Values for all covariates except for the intercept are percentage changes from the base rate of relocation (i.e., Intercept). The stars show statistical significance (*** p-val $<$ 0.001, ** p-val $<$ 0.01, * p-val $<$ 0.05).

the hypothesis (H1(a)) that overall bonding capital will make a person less likely to relocation. In contrast, **density** of the network does not show a statistically significant relationship with relocation choice. The median coreness of the network measures the extent to which a person's social circle is made up of one or more strongly embedded groups. A unit (standard deviation) increase in it decreases the odds that the affected individual will relocate by 9% to 12% across all models (Table 3.3). This suggests that the more strong ties (in terms of closures in the ego-network) a person has within their community, the more support they are likely to receive in the aftermath of a disaster. In addition, this tightly knit local social network represents a substantial opportunity cost to relocating.

Results also show that, as expected in H1(b) in comparison to overall bonding capital, the presence of stable core of strong ties (Strong Tie Retention) increases the likelihood of relocation. A unit increase in the individual's tendency to retain their strongest ties over time tends to increase the odds of relocation by 7%, which does not align with our hypothesis. This observation supports the mechanism proposed in (Engel and Ibáñez, 2007) that a person's closest relationships may be relatively robust to the strain that distance imposes should the focal individual chooses to relocate to the extent that they are not seen as an opportunity cost. Also, they may be qualitatively different from *strong* ties defined by closure in that they could mitigate the costs of relocation (e.g., look after property, assist with moving).

3.6.2 Bridging Capital (H2)

Our hypothesis that bridging capital makes relocation after a disaster more likely is supported by the results. The odds of a person relocating after a disaster increase with the size of their social circle at a rate of 7% to 9% across the models. The observed effect for structural diversity is not statistically significant ($p = 0.051$ in M5). However, the pattern of results across all models suggests that structural diversity has a similar, but smaller, effect to network size on relocation choice.

3.6.3 Network Robustness and the Effect of Social Capital (H3)

We find no evidence to support our hypothesis that structural robustness of the network enhances the effects of social capital on a person's decision to relocate after a disaster. This may be due to two reasons. First, our measures of robustness may not be sufficiently sensitive to capture any relationships on the sparse Twitter ego-networks after controlling for the pre-disaster level of social capital. Second, the effects of the disaster on the ego-network may not be randomly distributed as we have assumed. For example, if the closeness of the relationship is correlated with geography, an affected person's strong ties are also more likely to be affected by the disaster.

3.6.4 Effects of Disaster Attributes and Community Resilience

Disaster Properties

Both the scale and duration of a disaster are positively associated with the likelihood of an affected person relocating afterward. This is expected since both variables capture the magnitude of the shock on affected community and the larger the disruption, the more likely people would feel the need to relocate. While disaster duration has a relatively constant effect of increasing odds of relocation by approximately 20% across all models (**M3**, **M4**, **M5**, **M6**), the scale of destruction only exerts a significant effect on relocation after accounting for community resilience attributes. This suggests that community resources can counteract the effects of disasters of different scale in the short term, but sustained hazards are more likely to force people to move regardless.

Community Resilience

In general, the observed effects of community resilience attributes on the choice of relocation of its members can be encompassed within the following simple narrative. If an individual is more likely to have access to personal and community resources that improve their resilience (e.g, income, schools, community organizations), they are more likely to relocate. If, instead, they lack certain resources (e.g., insurance, home ownership, low English proficiency, not owning a vehicle) or own property within the community, they are less likely to relocate.

Additionally, the perceived attractiveness of the affected community (% Migration, % Vacant Rentals) also influences relocation choice. A community that is seen as more attractive is more likely to retain people after a disaster. Conversely, a disaster may accelerate migration out of a community that is already experiencing a decrease in population.

3.6.5 Geography of Social Capital and Relocation

In building our hypotheses regarding the relationships of different forms of social capital and post-disaster relocation choice, we relied on the assumption that the type of social capital associated with a relationship is also related to the geographical proximity of the parties involved in it. We test the validity of this assumption by estimating the fraction of alters from a person's Twitter ego-network that reside within their community (i.e. home county subdivision) and including it in the complete regression model except for terms for interactions between social capital and structural robustness (**M5**). Due to the rarity of geotagged tweets, we were only able to establish the home locations of at least one alter for only 2,724 affected people out of 16,372 included in the original regressions. For each such individual, we are able to identify, on average, the pre-disaster home locations of approximately one-third of the alters. We estimate the proportion of ego-network that is in each focal person's community using this sample of alters. Results show that a unit increase in the fraction of a person's social circle that live in the same community decreases the odds of relocation by 16% ($p = 10^{-3}$) which aligns with the expectation from our assumption that bonding capital is more likely to be derived from geographically co-located individuals and bridging capital is more geographically distributed.

3.6.6 Prevalance of Relationship Types and Relocation

People are likely to reach out and receive engagement from different parts of their network during different stressful events. Such differences in support and the nature of support provided are correlated with the nature of relationships between a person and different parts of the network (Aldrich, 2012b; Burt, 2007; Putnam, 2000). Prior work has shown that there are a variety of relationships that people maintain on Twitter, and these align quite well with categories of offline relationships (Choi et al., 2021). Furthermore, different types of ties become more active and engage more with a focal Twitter user when they experience different life events, such as becoming unemployed or a romantic breakup (Choi et al., 2023).

In our main analyses, we rely entirely on structural features of a person's ego-network to measure their access to two types of social capital (i.e., bonding and bridging). Here, we use a complementary approach to evaluate the effects of these two types of social capital, in addition to the effects of linking capital, which is challenging to measure using structural measures.

In this complementary analysis, we measure the different forms of social capital in terms of the prevalence of different types of relationships that the literature on social capital associates with them Aldrich (2012a). To identify the type of relationship between each person and each of their alters, we use the pre-trained deep learning classification model described in (Choi et al., 2021). This model assigns one of five different types of relationships (social, family, romantic, organizational, and parasocial) to pairs of Twitter users based on a number of features, including linguistic cues from tweets sent between them and the features of their profiles on the platform. Based on the literature on social capital, we propose the following hypotheses that relate the prevalence of different types of relationships and the types of social capital they embody to the relocation choice.

H4 **Family** and **Romantic** ties generally correspond to strong ties. Therefore, we consider the prevalence of these types of relationships to be forms of bonding capital. **Similar to H1(a), we expect that having more ties related to family and romantic categories will make relocation after a disaster less likely.**

H5 We expect that **Social** ties on Twitter as a whole is indicative of the broader community that a user engages with. Therefore, we consider the prevalence of social ties to be a measure of bridging capital. Additionally, **Organizational** ties (i.e., co-worker or hierarchical relationships) also represent a weak form of bridging capital that is less likely to provide the same level of support as family, romantic, or even social relationships Choi et al. (2021). However, these ties may be able to provide information that could be useful in the aftermath of a disaster. **Similar to H2, we expect that a higher prevalence of social and organizational ties will make relocation after a disaster more likely.**

H6 **Parasocial** relationships are defined as asymmetric relationships where one party places substantially more value in the relationship than the others Choi et al. (2021). This type of relationship is broadly the most consistent with linking capital, even though it has certain limitations. On social media platforms such as Twitter, parasocial relationships often fall into the celebrity-fan category than other categories such as institution-citizen or politician-citizen that are more clearly associated with linking capital. However, on social media, politicians often have celebrity status and both them and celebrities serve as opinion leaders that influence both formal institutional action and informal collective action Hunt and Gruszczynski (2023). We expect that having more linking social capital will mean that a person is able to find or acquire access to resources that would be useful in rebuilding in the aftermath of a disaster. **In other words, we expect a higher prevalence of linking social capital will make relocation less likely.**

Method. We summarize the ego-network of each affected person in terms of the number of relationships of each type present[6] and fit a new series of regression models for relocation choice with these five measures replacing the social capital measures of the original analysis (**M1**, **M3**, and **M5**).

Results. Table 3.4 shows the estimated effects of different types of relationships on the choice of relocation of an affected person during the first twelve weeks after the onset of a disaster. We find no evidence to support our hypothesis that family and romantic relationships would make relocation less likely in the aftermath of the disaster (H4). This is surprising given the consistent reference to close ties that these types of relationship generally represent in the literature on bonding capital. One explanation may be that family and romantic relationships observed on Twitter may not be representative of offline relationships in volume, nature, or both.

In the case of the hypothesis that social and organizational ties representing different forms of bridging ties will make relocation more likely (H5), we observe mixed results. First, the prevalence of **social** category, which is the most frequent in the observed ego-networks, makes a person more

[6]The mean ego-network consists of 8.3 relationships, of which, 4.3 are social ties, 1.1 are family ties, 0.7 are romantic, 1.3 are organizational, and 0.9 are parasocial

likely to relocate in the aftermath of disaster in the first two models. We consider the observed relationship as further confirmation that bridging capital increases the likelihood of relocating after a disaster. We note that the effect is not statistically significant when community-level resilience variables are included in the model (**M5**) which requires further investigation. Additionally, we did not observe a statistically significant effect of the prevalence of organizational ties on relocation, suggesting that these ties, as measured on Twitter, do not provide any support or incentive to either relocate or stay in the aftermath of disaster.

Finally, the prevalence of parasocial relationships significantly decreases the odds that a person would relocate in the aftermath of a disaster. This supports our hypothesis (H6) that linking capital would help an affected person access resources from institutional sources or other individuals or entities with authority or power, which in turn would make it more likely that they would rebuild within the neighborhood.

Relationship Type	**M1'**	**M3'**	**M5'**
Social	9.1^{**}	8.4^{*}	5
Family	-2.4	-2.1	-1.8
Romantic	2.3	1.2	0
Organizational	-0.4	-0.4	2.5
Parasocial	-15.2^{***}	-16^{***}	-11^{**}

Table 3.4: Effect the prevalence of different relationship types on relocation choice during the first 12 weeks after the onset of disaster. Values are the changes in likelihood of relocation for a unit (SD) increase in the prevalence of each relationship type. The stars show statistical significance (*** p-val < 0.001, ** p-val < 0.01, * p-val < 0.05).

3.6.7 Robustness Checks

The results for models estimated considering alternative definitions of relocation (boolean, log distance, log rank distance) over different post-disaster periods (12, 24, and 36 weeks) are generally consistent with the observations from the core set of models discussed in the main text except for the case of ego-network density which shows a statistically significant positive effect on the rank distance of post-disaster home location from the pre-disaster one (Figure B.1). This observation contradicts our hypothesis and requires further study.

3.7 Limitations

Although this study considers a relatively large number of people (more than 16,000) in a range of different communities (76) and different types (10) of disasters, these individuals are drawn from among Twitter users. Even among Twitter users, our dataset is sampled from individuals who use geotagged tweets and are reasonably active on the platform. Users who geotag tweets are more likely to be young, more wealthy, reside in urban areas, and belong to certain ethnic groups (Sloan and Morgan, 2015). We have mitigated, but probably not completely eliminated, the effects of these biases on our results by including a variety of community-level resilience characteristics.

Due to limitations of the Twitter API, it is infeasible to collect data for the entire social circle of each affected individual. As a consequence, the ego-networks that are studied here are samples that correspond to individuals with whom our focal users had a relatively strong reciprocal pattern of communication (5 mentions in direction over 24 weeks). Therefore, only around eighth of the originally identified individuals could be used in the regression analyses that evaluated the relationships between relocation choice and different forms of social capital. Since ego-network collection is predicated to an extent on the strength of the ties, our observations are likely to be based on only a fraction of the full range of variation in social capital among affected users. It is possible that the magnitude of the estimated effect of social capital on relocation is different within this population than within the average population.

Finally, in this study, we assume that a person's Twitter ego-network provides a fair reflection of social capital afforded by their full, online and offline, social network. While there is some evidence to support that online and offline ego-networks are similar at least structurally (Dunbar et al., 2015), there are likely to be differences in the motivations and mechanisms for maintaining relationships over social media as opposed to in the real world. Although our observations on the effects of social capital on relocation choice broadly align with the hypotheses derived from traditional social capital literature, further investigation is necessary to confirm the mechanisms through which social media social capital may influence relocation.

3.8 Discussion

This study investigated the relationship between a person's social capital measured on social media and their relocation choice in the aftermath of a disaster. In line with prior work that has argued for the importance of social capital in post-disaster adaptation and resilience (Cong et al., 2018; Adeola and Picou, 2014; Aldrich and Meyer, 2015; Norris et al., 2008), our results show that a person's social capital is likely to influence their decision to relocate. In particular, we have demonstrated that the resources afforded by the immediate network of a person affect relocation choice even after accounting for social capital embodied in the larger community, socio-economic and infrastructural community attributes, as well as the nature of the specific disaster.

The observed relationships between different forms of social capital and relocation after disaster, largely align with our expectations based the literature of social capital and migration/relocation applied to the post-disaster context. This includes support for a rarely discussed mechanism related to the subtype of bonding ties defined by their stability (Engel and Ibáñez, 2007) that seems to oppose the more well-known effect of overall bonding capital in encouraging people to rebuild in situ. In fact, for this meachanism that a person who maintains a core of strong ties over time is more likely to relocate in the aftermath of a disaster is the most consistent observation related to social capital across all model variants. We believe that further study is needed to better understand this mechanism.

We also find that a person's bridging capital, derived from distant and sparsely connected parts of their network, increases the likelihood that they relocate after disaster. The size of a person's network is more relevant in the short term, with the effect on both relocation and the distance declining over longer periods. On the contrary, having links to more distinct social groups seems to increasingly influence the likelihood of relocation in the long term (Figure B.1). This suggests that, like bonding capital, there may be multiple mechanisms that link bridging capital to relocation. For example, the size of a person's social circle may indicate the ease with which each person is able to make accommodations to temporarily move out of the community. However, to relocate permanently, they would need access to more rare and substantial resources such as permanent housing and employment. A person who interacts with many disparate social groups is likely to be more successful in this regard (Granovetter, 1973; Burt, 2007).

Finally, we find that the prevalence of parasocial relationships, which are indicative of linking social capital, a new form that signifies a person's awareness of and trust in the institutional and power structures significantly increases their odds of continuing to remain in the neighborhood after a disaster. This observation implies that this knowledge is likely very useful for acquiring additional resources to rebuild their lives within the community (Kawamoto and Kim, 2019; Aldrich, 2017). From a methodological perspective, we note that this observation was made possible by utilizing attributes of affected users that went beyond purely structural properties of their network such as elements of their virtual identities and linguistic features of their interactions. It suggests the potential for developing more nuanced social capital measurements that combine structural and qualitative attributes of relationships observed on social media.

Implications Disasters, both natural and man-made, affect more than 300 million people worldwide, each disrupting their physical and social environments each year (Ritchie et al., 2022). Increasingly, governments, disaster relief organizations, and social workers recognize the importance of resilience in understanding the post-disaster outcomes for people and communities (Aldrich, 2012a; Edgemon et al., 2020). Social capital is known to be a key contributor to community resilience by providing individuals with access to resources and support beyond their own, as well as making communities more cohesive as a whole. However, in characterizing disaster-stricken communities, governments and relief agencies are generally forced to rely on coarse-grained collective measures of social capital such as the presence of community institutions. This study has shown that while these resilience indicators at the community level do indeed explain post-disaster behavior to an extent, the personal networks of affected individuals and the resources embedded therein also play an important role. Furthermore, we have shown that proxies of a person's social capital measured over social media are likely to be somewhat reliable indicators of their overall social capital. This is suggestive of the potential of social media-based measurements to augment the existing estimates of individual and community-level vulnerability to disasters, which would enable relief organizations, community managers, and social workers, among others, to calibrate and plan their supportive activities. Finally, the results of this study, including our observations on community-level resilience, highlight the tension between individual and collective resilience in the context of relocation. We observe that access to resources, both in social and other forms, often increases the likelihood that an individual would leave the community after a disaster. Conversely, individuals who are underprivileged and have more localized social capital are less likely to relocate. Therefore, at least for a fraction of the population, relocating in the aftermath of a disaster represents a resilient adaptation. In contrast, for a community, its members choosing to relocate after a disaster is clearly an adverse adaptation, as it further depletes collective resources and disrupts the already strained social fabric. Therefore, from the perspective of governments and community leaders, the ideal scenario would be one in which people find that they can rebuild in situ instead of relocating. Future work should investigate the interaction between individual social capital and community dynamics to understand how individual incentives and community outcomes can be aligned.

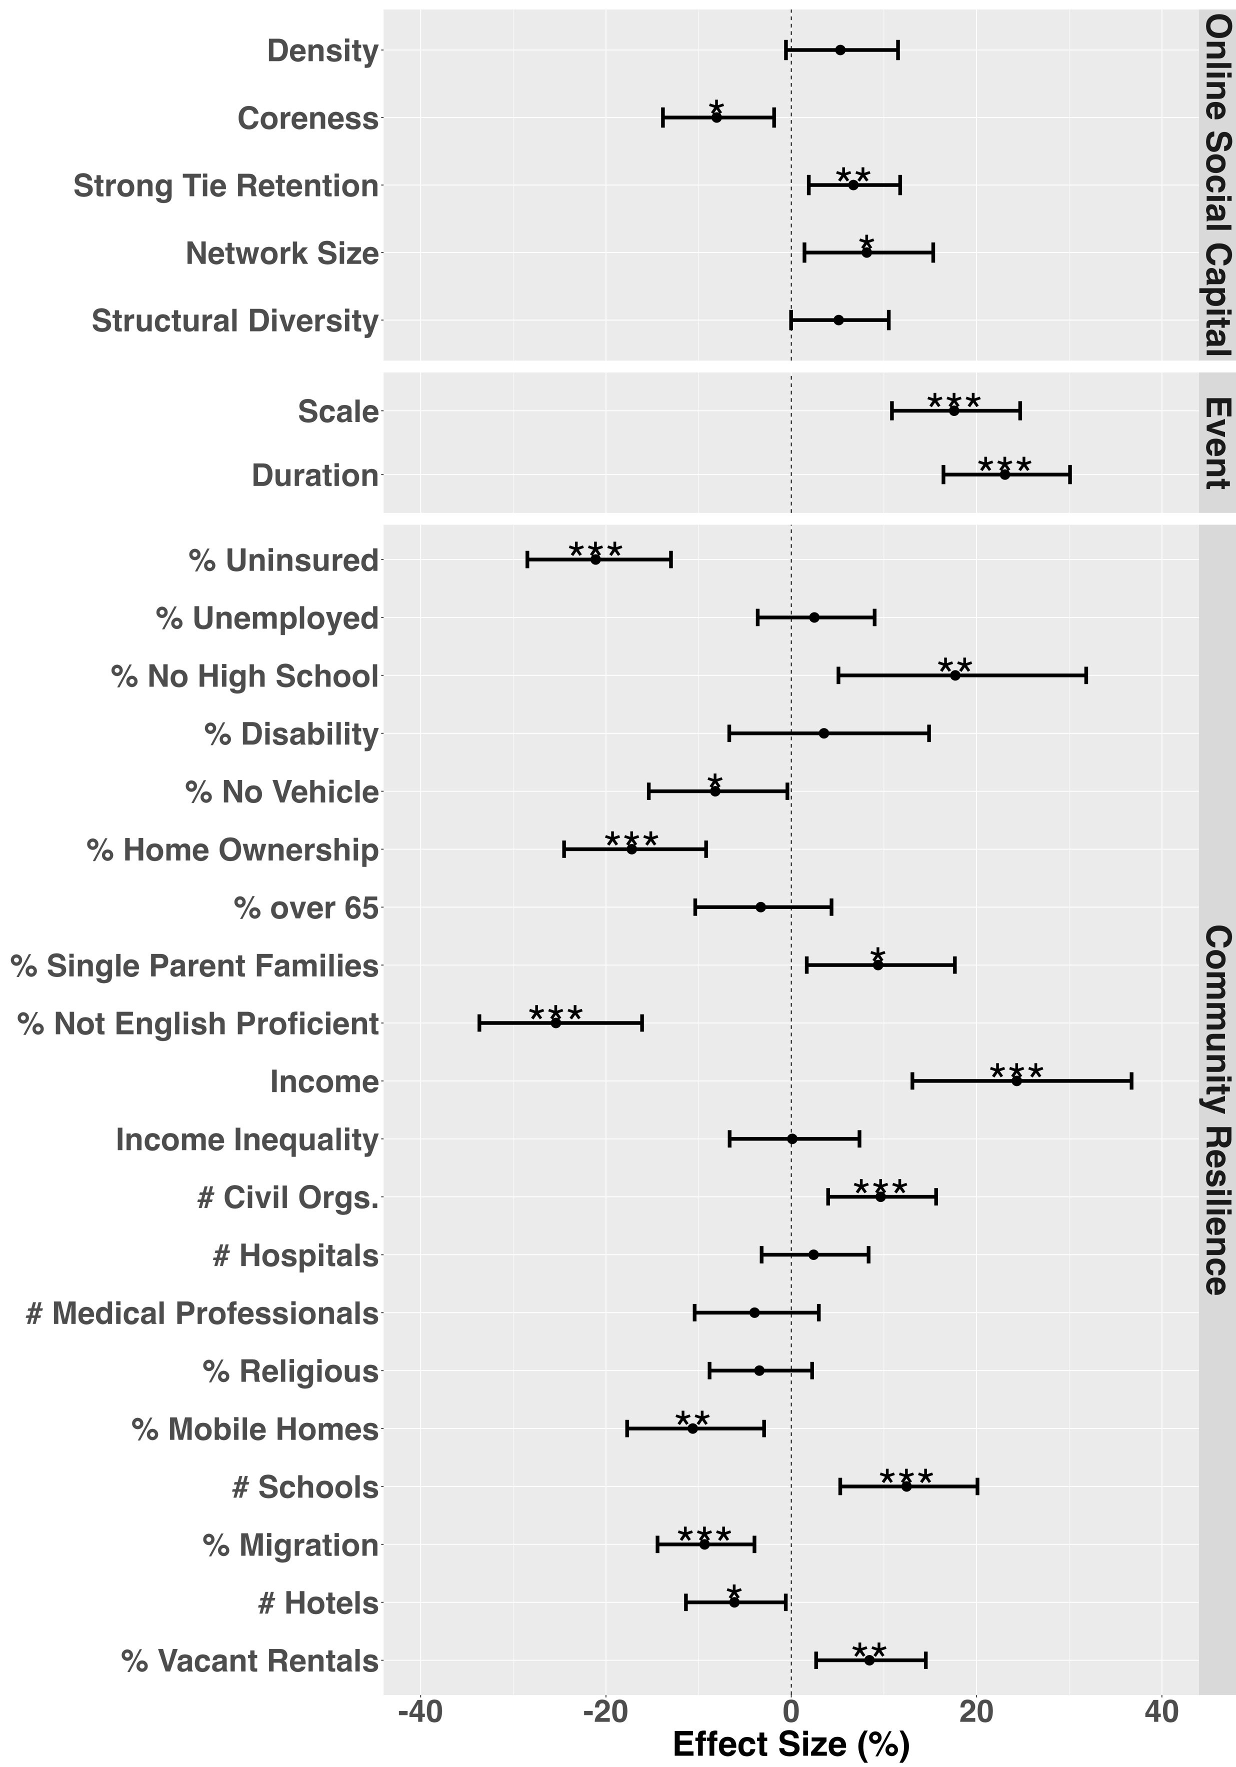

Figure 3.1: Co-variate Effect Sizes on the Likelihood of Relocation after a Disaster (**M5**). The error bars show 95% confidence intervals. The stars show significance (*** p-val < 0.001, ** p-val < 0.01, * p-val < 0.05).

CHAPTER 4

How GitHub Projects Respond to Increased Attention

4.1 Introduction

Collaborative crowd communities have recently become commonplace within the information economy. They represent a significant departure from traditional organizations in terms of worker motivation, organizational structure, and operating dynamics. Community members are driven by non-monetary motivations such as reputation and collective identity (Benkler et al., 2015). Further, these communities consist of loosely-knit and geographically dispersed teams with limited formal structure, roles, or routines. This is in contrast to organizations, which have clearly defined boundaries, authority structures, and formal routines (Benkler, 2006). Crowd communities such as open source software (OSS) projects exhibit growth, maturation, and decline driven by both endogenous and exogenous processes similar to more-traditional organizations. However, due to their differences from traditional organizations, our understanding of these crowds is limited. This is especially true for teams that are growing rapidly due to exogenous shocks that draw increased attention (Keegan et al., 2013; Zhang et al., 2019, 2017a).

A common mechanism by which projects receive such shocks on GitHub is by appearing on the trending page. The default view of the trending page displays 25 projects that have demonstrated a recent growth in productivity and attention (Blog, 2013). Projects on the trending page have the potential to attract significant attention, from users and future contributors. This represents both an opportunity and a challenge for their members. They may be able to expand their workforce by attracting committed and capable new contributors. Yet, their capacity to coordinate and remain responsive to their audience might be strained by a deluge of prospective users and contributors (Keegan et al., 2013). Furthermore, the vision of the original team may be at odds with the aspiring newcomers, leading to friction and negative interactions that may undermine further growth.

Motivated to characterize these challenges, opportunities, and their consequences, we study the behavior of new contributors and the adaptive response of existing members from the perspective of organizational change (Child and Kieser, 1981). We first characterize the scale of growth and the interests and commitment of newcomers. Then we inspect how core members adjust their work routines in response. Finally, we employ a theoretical framework of distributed leadership to contextualize the evolution of the coordination structure.

Our results show that trending results in a veritable explosion in attention and engagement. Types of engagement include starring the project to indicate interest, using the project software and reporting issues with it, suggesting additional features and contributing code and other content to the project. However, the vast majority of this external attention is restricted to shallow short term interactions with the project—more so than external engagement under *normal* conditions. Despite this, the scale of the attention shock means that there is still substantial growth in team size and the work being done. Subsequently, core members shift toward coordinating the growing workforce and move away from direct production. Even with this new administrative orientation, members struggle to keep up with increased coordination requirements leading to backlogs, especially with regards to outsiders, such as taking longer to respond their engagements. Finally, we find that as members adopt a more comprehensive approach to managing the burgeoning workforce, they allow outsiders to unofficially fill in moderately important positions. In addition, members introduce technological solutions to improve work quality and efficiency—all in line with a more distributed coordination structure.

In sum, we (i.) characterize the scale and nature of shocks resulting from a project being thrust into the limelight and (ii.) shed a light on the resulting collaboration dynamics between the original team and interested outsiders from the community. This understanding will provide the necessary background for better equipping such teams to manage and take advantage of the deluge of good Samaritans drawn to projects during times of rapid growth.

4.2 Related Work

Organizational Growth and Change. Organizational scholars have studied the impact of rapid growth in organization size. Early research focused on the benefits of such rapid growth such as an increase in organizational resources (Whetten, 1987; Baker and Cullen, 1993; Kiviluoto, 2013). Scholars reasoned that these increases can allow organizations to seek out more opportunities for profits (Habersang et al., 2019) and that they can reduce internal conflict as there is less need for internal competition for resources (Baker and Cullen, 1993).

At the same time, organizational scholars have become aware of the dysfunctional consequences of growth. First, increases in size can lead to greater work complexity (Blau, 1970; Whetten, 1987) and, thus, difficulties in coordination (Baker and Cullen, 1993; Liozu et al., 2014; Nicholls-Nixon, 2005). Second, as organizations grow, relationships between employees become less personal—

decreasing organizational cohesion (Liozu et al., 2014). Organizations often struggle with ensuring that these new employees learn their shared norms and work practices (Whetten, 1987). This is especially problematic when growth occurs rapidly (Liozu et al., 2014). In some cases reduction in cohesion is accompanied by increases in conflict (Nicholls-Nixon, 2005). Finally, increases in complexity and decreases in cohesion contribute to ineffectiveness and inefficiency (Nicholls-Nixon, 2005). This can lead to organizational processes being strained, stopping the organization from meeting its goals (Whetten, 1987).

Recently, organizational scholars have shifted toward identifying approaches to mitigate the dysfunctions that result from growth. According to this stream of research, organizations must change their management style and structure to adapt to growth (Liozu et al., 2014; Nicholls-Nixon, 2005). To accomplish this, organizations must decentralize decision making and introduce more formal rules and procedures (Baker and Cullen, 1993; Whetten, 1987). Both decentralizing and introducing rules and procedures allow organizations to rapidly respond to changes in the environment ,while maintaining control to facilitate the coordination of work. Organizations should also restructure to ensure that individuals or sub-units do not get overwhelmed (Nicholls-Nixon, 2005). This often involves creating new sub-units and assigning them new work and/or redirecting work to them from existing sub-units. Finally, organizations must build and maintain relationships among existing and new members (Nicholls-Nixon, 2005).

Drawing from these findings, we develop hypotheses for how GitHub teams respond to rapid growth due to going trending. However, we also consider the important differences between crowd collaborations and traditional organizations. For instance, traditional organizations have well-defined expectations in terms of employee productivity. In contrast, crowds are characterized by volunteers, with different levels of commitment, contributing due to non-monetary motivations (Benkler et al., 2015). As a result, participation in these communities is often characterized by a long-tailed distribution where most contributors do very small chunks of shallow work while a small core group performs the bulk of the work including overall coordination. (Ortega et al., 2008; Laniado and Tasso, 2011). This workload inequality makes coordination of volunteers more manageable (Raymond, 1999; Arazy and Nov, 2010; Kittur and Kraut, 2008; Romero et al., 2015).

GitHub. Prior work on GitHub has focused on understanding the collaborative dynamics on the platform. Researchers have studied the social and collaborative interactions that take place on the site and how such interactions impact the software. Qualitative studies, mainly based on interviews, have found that users infer each other's goals, expertise, and interests based on their logged actions and profile information. This information is then used to make decisions about which projects to contribute to and how to organize collaborations (Dabbish et al., 2012; Marlow et al., 2013; Vasilescu et al., 2015a). Other work focused on measuring large-scale statistical properties such as the power-law-like distribution of the number of contributors and watchers per project, the low levels of reciprocity in the followers network, the geographical distribution users (Lima et al., 2014), and how coordination efforts by developers scale with the size of a project (Romero et al., 2015). Researchers have also characterized the properties of GitHub repositories at large scale including the fraction of personal and inactive repositories and the fraction of repositories that use pull requests (Kalliamvakou et al., 2014).

Another line of research focused on identifying the properties and routines of successful GitHub collaborations. Characteristics indicative of success include diversity in gender and tenure (Vasilescu et al., 2015b), the use of automated processes such as continuous integration (Vasilescu et al., 2015c), teams with a high fraction of members with a history of collaboration (Casalnuovo et al., 2015), and low levels of multi-tasking (Vasilescu et al., 2016). Rather than characterizing collaboration patterns on GitHub or identifying successful collaboration dynamics, our goal in this paper is to understand how such dynamics change when exposed to attention shocks.

Shocks on Crowd Collaboration. Past research has investigated the effects of external shocks on crowd collaborations in other crowdsourcing platforms such as Wikipedia (Keegan et al., 2013; Zhang et al., 2017a, 2019, 2017b). This work is unique in two important ways. First, OSS development involves substantially more complex and diverse work than Wikipedia.
Projects on GitHub incorporate the hallmarks of crowdsourcing collaboration in openness and promoting external engagement, much like Wikipedia. However, they also demonstrate a clear insider vs. outsider dichotomy in terms of authority enabled by access privileges. Further, these projects face growth and competitive pressures from other similar projects that also rely on the same audience and volunteers for success. Second, Github provides tools for setting access privileges, categorizing, tracking and reviewing work as well as project management. That allows us to explore more nuanced questions regarding response of crowd collaborations to shocks that parallel what is known about traditional organizations.

4.3 Hypotheses Development

In this section, we build a research framework informed by prior work in organizational change management as well as participation and coordination dynamics of crowd collaborations.

4.3.1 Community Engagement

The GitHub trending project page is the primary mechanism by which the platform introduces high quality repositories to the community (Begel et al., 2013). Therefore, it serves as a source for developers and users to find repositories to contribute to and use (Jiang et al., 2017). Furthermore, the trending page provides a star button for each trending repository. Therefore, we posit that being featured in the GitHub trending list will lead to a substantial increase in community engagement with the repository. Thus, this is our first hypothesis.

Hypothesis 1: Trending will increase community engagement with a repository.

(a) Trending will increase observed community interest, measured through starring and forking behavior.

(b) Trending will lead to an increase in external contributions.

4.3.2 Growing Pains

Organizational theory suggests that traditional organizations with rapidly growing sales or market share are chronically short of resources and labor. This commonly leads to employee hires to compensate for growth. This, in turn, causes a number of growing pains as the existing organizational structure is strained by the increased number of new and untrained employees and new tasks associated with growth (Hambrick and Crozier, 1985; Fombrun and Wally, 1989). We expect a similar trend for the repositories that are featured on the GitHub trending page. These repositories are likely experience a sudden increase in the number of external contributors (Hypothesis 1.b). We posit that this sudden increase will similarly strain the capacity of the core repository team. Therefore, we expect that increased workload for the members will translate into delays and backlogs in tasks.

Hypothesis 2: The shock will strain the capacity of the project and create backlogs of tasks.

4.3.3 Adaptation

We draw on organizational theory to build hypotheses as to how the members of a trending GitHub repository will respond to the increased attention and contributions from the community,

4.3.3.1 Work Routines

In firms where the labor force is rapidly growing, employees from the pre-growth period—being more experienced and knowledgeable about firm culture—are compelled to take on more management responsibility (Hambrick and Crozier, 1985; Child and Kieser, 1981). Similarly, we posit that GitHub repository members would have to spend more time responding to queries and coordinating the work of the increased number of outside, and potentially inexperienced, contributors. Correspondingly, we expect them to also scale back their own development work.

Hypothesis 3: Members will take on a more administrative role

(a) Members will do more organizational work, such as responding to and directing external contributors.

(b) Members will do less development work.

4.3.3.2 Coordination

Growing organizations manage the increasing complexity of coordination by adopting a more decentralized and modular coordination style. They can achieve this through empowering lower-level employees, specialization of roles or the division of work among overlapping working groups (Whetten, 1987; Hambrick and Crozier, 1985; Child and Kieser, 1981; Miller, 1994). However, this flexible structure can lead to a decline in its core values and culture, which play an important role as a cohesive force. To compensate, management may choose to make its values explicit through formalized procedures and highly visible symbols and slogans (Hambrick and Crozier, 1985; Fombrun and Wally, 1989; Whetten, 1987; Platt and Romero, 2018). Here, we employ Gronn's Distributed Leadership framework as the lens for viewing these phenomena in GitHub (Gronn, 2002; McDonald et al., 2014). Gronn proposes that organizations that employ a distributed approach to coordination will demonstrate three broad properties; (1.) They encourage workers to dynamically form temporary teams for spontaneous collaborations, (2.) Employees establish relationships and responsibilities through collaborations and receive recognition for their contributions, and (3.) Working practices are institutionalized as part of the organizational governance. We operationalize these observations as follows.
Hypothesis 4: After the shock, repository coordination will become more open and decentralized.

(a) Members will increase their collaborative engagement with outside contributors.

(b) Outside contributors will take on more central roles within the work routines.

(c) Collaboration will take on a modular structure.

(d) Members will reinforce core values through automated enforcement.

4.4 Background on GitHub and Data

GitHub is a popular social coding platform that provides tools for collaborative software development, social networking, and reputation, project, and access privilege management. The affordances of the platform have resulted in a large community of open source projects and developers.

Due to the complexity of software development, GitHub provides several features for scaffolding of common types of work. The availability of these affordances on GitHub provides us with convenient measures of collaboration dynamics.

GitHub users can engage with a public repository in various ways. Users can make changes to a repository by submitting a *pull request*. Requests need to be approved by a member of the repository in order to be merged. Non-members can also submit and make comments on *issues*, which are used to report bugs, request changes, or to make general comments. Those interested in creating an independent copy of the repository for their use can do so by *forking*. Finally, users can signal approval of a repository by *starring*.

Some actions are restricted to only repository *members*. Repository members can do anything non-members can do in addition to adding new members, directly *pushing* changes to the code without a pull request, approving or disapproving pull requests, categorizing issues and contributions, and closing issues, among others. This allows the members to effectively steer the project but also introduces an explicit coordination bottleneck.

We rely on two different data sources: (i.) GitHub trending data, (ii.) GitHub project trace data. We use (i.) to identify the attention shocks and use (ii.) to examine how the GitHub community and the trending project team respond to the shock.

4.4.1 GitHub Trending Data

GitHub Trending page lists a set of repositories—updated 8 times a day—to characterize "what the GitHub community is most excited about today"[1]. One can filter the trends by language—all languages are shown by default. These languages include *C++, HTML, Java, Javascript, PhP, Python*, and *Ruby*, which are the top 7 languages that appear on the language selection drop-down list[2]. For each trending repository, website visitors see the owner/repository name, primary language name, repository description, a star button, and a list of the top five contributors for the project. According to GitHub, in order to choose the trending repositories they "look at a variety of data points including stars, forks, commits, follows, and pageviews, weighting them appropriately. It's not just about total numbers, but also how recently the events happened." (Blog, 2013).

In this study, we scrape the GitHub Trending page every three hours—to make sure each trending change is captured—for 7 months (6/27/18 - 1/31/19). We collect this data for all languages combined, as well as individual languages. We identify triplets (r, t, l) of repository r and time t such that r was trending at time t in language l and was not trending any time before during our scraping period [3]. To examine the projects with strong trending effects, we further limit our anal-

[1] https://github.com/trending

[2] Top languages list is personalized for logged in GitHub users. We use the non-personalized ordering here.

[3] In order to account for left censoring, we only consider trending repositories 14 days after the beginning of the

ysis to the projects that appear among the top 5 trending results irrespective of whether a project appeared directly or rose through the ranks. We identify 1,107 such triples [4].

4.4.2 GitHub Projects Trace Data

We use GitHub event stream data to examine how a repository team and the broader GitHub community respond to a repository being featured in the trending repositories list.

4.4.2.1 GHArchive Data

We use GHArchive archives (https://www.gharchive.org/), which provides events information about GitHub repositories, to retrieve event data between January 2018 and February 2019 (inclusive) for all repositories. This archive includes timestamped data on a large number of events including stars, pushes, pulled requests, issues, comments, member additions, among many others. There are approximately 350K repositories with at least one star or fork in the timeframe of our shocks (6/27/18 - 1/31/19).

4.4.2.2 REST API

While the GHArchive data is rich, it still does not include some important events that are central to our measures. As such, we complement that data by using GitHub REST API in three ways. First, we retrieve the language of each repository with at least one event within the study window. This important feature is used in our propensity score matching step (details below). Second, we collect data about 36 additional event types such as *subscribe*, *renamed*, and *labeled*, related to actions on issues and pull requests to study group dynamics. Third, we collect data on commits including author ids and status updates to identify automated tasks. Unlike GHArchive, data collection through the REST API is rather slow. Therefore, we collect this data for our shocked repositories and the set of control repositories identified using propensity score matching. In total, there are roughly 7K such repositories.

4.4.3 Preprocessing

We perform two preprocessing tasks that are important for ensuring the reliability of the our outcome measurements: (1) member identification and (2) removal of bot activity.

---

scraping period

[4] We originally identified 1,297 triplets. However, 190 were created at most 7 days prior to when they were trending or had a very small amount of activity during that 7 day period (i.e. no stars or forks). We are unable to compute various measures for these repositories and thus we removed them from our study, resulting in 1,107 triples.

4.4.3.1 Member Identification

The evaluation of many of our hypotheses (2(a), 3(a), 3(b), 4(a), and 4(d)) rely on our ability to distinguish repository members from external contributors. To identify members, we use all events from GHArchive between 01/01/2018-01/31/2019. A user is labeled as a member if they perform any event that is restricted to members only. One limitation of this approach is that there could potentially be users who have been members of repositories since before 01/01/2018 but did not use any membership privileges during the period. However, these users are not exercising any member-level authority and are no different from outsiders in their observed behavior. Therefore we expect such users to have little or no effect on our outcome measurements.

4.4.3.2 Removal of Bot Activity

GitHub allows teams to use task automation or *bots* for tasks such as spam prevention and code quality assessment and deployment. Bots generate tremendous amounts of activity/events that are not useful to understand people's collaborative dynamics. We use the following two fold method to identify and remove bot accounts from our study. First, we filter out GitHub accounts with usernames that end with either "bot" or "[bot]". Next, to identify other bots that do not follow this standard, we perform a simple classification. We find an activity level, in terms of the average daily number of pushes when an account was active, above which an account is considered a bot. To evaluate different thresholds, we bucket accounts by number of pushes between ∞, 1000, 500, 200, 100, 50, 25, 13, 7, 4, 2, and 0. We manually label 10 randomly selected accounts from each bucket as bots or human by observing their activity. Next, we find the threshold that maximizes the macro F1 score weighted by the distribution of accounts in the complete data set. We find that the best threshold is 13 average daily pushes ($F1 = 0.78$). Based on this classification, we remove 38 bot accounts in 54 treated repositories and 45 bot accounts in 67 control repositories.

4.5 Measurement

In this section, we describe the construction of measures to test our hypotheses. Table 4.1 provides a summary of all the measures.

4.5.1 External Engagement

For hypothesis 1, we need to measure the engagement of external users (non-members) with the repository. We measure general interest and contributions separately. To measure external interest (H1a), we track the number of stars and forks received during a period. We use both the raw number as well as a log normalized version. For a measure M_t at time period t, the normalized version of the measure is $N_t = \log\left(\frac{M_t}{M_{t-1}}\right)$. We use this normalization for any metric that is a count or a duration. We also measure external contributions (H1b) using three measures that capture contribution at different levels of commitment; (1) Number of external contributors — users who perform any action in the repository other than starring or forking, (2) number of issues opened by external users, and (3) number of pull requests submitted by external users.

4.5.2 Backlogs and Productivity

For hypothesis 2, we need to measure the productivity of members to assess whether the flow of external contributors expected after the shock creates backlogs. We measure productivity in two ways.

First, we define *response delay* as the average amount of time it takes a member to respond to an issue or pull request. This measures the overall responsiveness of the members. Second, we measure *closure efficiency* as the effectiveness of the members in closing open tasks. To do so, we first define the concept of a *stale* issue or pull request. We consider a task that has remained open for 60 days as stale. We assume that stale items are probably outdated and not will not lead to further action. We then consider the amount of time available to close a task during a period as the minimum of the time remaining during the period and the time remaining before it becomes stale. Finally, we measure the repository issue (pull request) closure efficiency, as 1 minus the average fraction of time taken by members to close active issues (pull requests) and the time available to them.

4.5.3 Member Administrative Orientation

For hypothesis 3, we need to evaluate whether members in shocked repositories are becoming more administration oriented. We separately measure volume of administrative work (H3a) and development work (H3b). For development work, we measure (1) the total number of pushes and pull requests by members and (2) the total number of lines of code or text members added using pushes and pull requests. For administrative work, we measure the total number of issues (pull requests) closed.

4.5.4 Coordination Processes

For hypotheses 4, we need to measure various aspects of a project's coordination patterns and openness.

Member Engagement with Outsiders. Members engage outsiders primarily through pull requests and issues. Therefore, we measure member engagement with outsiders (H4a) using (1) the average number of issues (pull requests) edited by a member and the (2) the total number of outsider pull requests accepted during a period.

Table 4.1: Measurement of GitHub project behavior used to test our hypotheses described in section 4.5

	Group	**Measure**
External Engagement (H1)	External Interest	# Stars # Forks
	External Contribution	# External Contributors # Issues Opened # PRs opened Avg. # Files Edited by Outsiders
Backlogs (H2)	Productivity	Issue Response Delay PR Response Delay Issue Closure Efficiency PR Closure Efficiency
Member Orientation (H3)	Development Work	# Pushes and PRs # Lines of code
	Admin. Work	# Issues Closed # PRs Closed
Coordination Processes and Decentralization (H4)	Member Engagement with Outsiders	Avg. # Issues Edited by Members Avg. # PRs Edited by Members Total # PRs Approved
	Decentralization	Outsiders in Top File Editors Outsider Centrality-File Net. Outsider Centrality-Issue Net. Outsider Centrality- PR Net.
	Specialization	Modularity of File Network Modularity of Issue Network Modularity of PR Network
	Automation	Binary use of Automation # Automated Tasks

Decentralization. We quantify decentralization within projects by examining the distribution of collaborative effort between members and outsiders. We use two types of metrics to evaluate the changing role of outsiders[5]. First, in order to assess the extent of the impact that outsiders have on

[5] Another potential way to examine decentralization is to quantify the rate with which outsiders transition to project members. However, recruitment of members is, in general, a slow process and we expect our study period may not be long enough to capture changes.

project's output, we measure the percentage of outsiders among the top x% users in terms of files edited.

Second, we measure the role of the outsiders in the collaborative process by assessing their degree centrality in three types of collaboration networks: collaboration in (1) files, (2) issues, and (3) pull requests. We construct each of the collaboration networks as follows. First, we build a bipartite network that indicates which users edited which items. In the file-user network, we add an edge between a user and a file if the user has edited the file. In issue-user and pull request-user networks, we add an edge between a user and an issue (pull request) if the user has taken any action in relation to it (e.g. open, close, comment, label, review, pin). Next, we construct the collaboration network by getting the projection of the bipartite network on the users such that the weight contributed by an item acted on by two users is inversely proportional to the total number of users that have acted on it (Newman, 2001). Rather than measuring the change in the mean centrality of outsiders, which would be problematic since the size of the networks can change, we measure the fraction of outsiders that are among the top $K\%$ most central nodes. We test several values of K between 10% and 50%.

Specialization. Modularity of a network measures the degree to which it is organized into groups of nodes that interact more often with each other than with outsiders. In the file, issue, and pull request collaboration networks of a repository, we employ it is an indicator of the extent to which contributors specialize on different aspects of the project, such as contributing code for a particular feature or addressing certain types of issues.

Automation. GitHub project teams can use automation to perform a variety of routine tasks, such as running tests against submitted pull requests, edit-locking stale issues and pull requests for spam prevention, and build and deploy code. In general, these tasks reflect the values and standards members wish to uphold in addition to being a mechanism for increasing productivity. We consider two measures in this context; (1) A binary value indicating if a project uses any task automation and (2) the unique number of such tasks.

4.6 Methods

In the following sections, we detail our approach for constructing a comparison group of repositories using propensity score matching, operationalizing our measures, and testing our hypotheses using difference-in-difference.

4.6.1 Propensity Score Matching

We use propensity score matching (PSM) to select a comparable set of matched repositories (Rosenbaum and Rubin, 1985; Imbens and Rubin, 2015). Broadly, this involves four steps; (1.) We estimate the likelihood (i.e. propensity) of a repository receiving the treatment (i.e. going trending on a given day) through a regression with an appropriate set of covariates described below. (2.) For each treated repository, we find the nearest neighbors w.r.t. a combination of the propensity and covariate similarity, (3.) We stratify the resulting set of shocks and matches by their propensity scores, and (4.) Within each strata, we estimate the overall similarity of the matches to the shocks w.r.t. covariates to ensure we have adequately accounted for selection effects. In the following sections, we describe this process in detail.

4.6.1.1 Covariates

While we are unaware of the exact design of the algorithm that GitHub uses to select trending repositories, we know that it is largely based on the size and the growth of external attention measured by stars and forks (TechCrunch, 2013; Blog, 2013). In fact, according to the former head of GitHub's marketing team, "stars are the primary consideration for whether a repository is trending or not" (Begel et al., 2013). Based on this we use the number of daily stars and forks for estimating the likelihood of each repository-date combination having experienced a shock. Our goal is to match each shocked repository to a control set of non-trending repository-date combinations that had a similar propensity to be shocked based on the scale and trend in attention growth in the preceding week. To do so, we condition on the following covariates. First, to account for volume, we use daily number of stars and forks a repository received over the week preceding the day of the shock. Second, to account for daily changes, we also use Log ratios of daily numbers of stars and forks of consecutive days, with the count of the more recent day as the numerator, over the week preceding the day of the shock. To be a candidate for the matched set, a repository-date combination must meet the following conditions: (i) the repository must not appear anywhere on the GitHub's trending page during our study period and (ii) the repository must receive at least one star or fork during the week prior to the date.

4.6.1.2 Propensity Scores

We model the likelihood of a repository trending as a logistic function of our covariates to estimate propensity scores for each repository-date in our dataset. Given the imbalanced nature of our data (1,107 positive and 27 million negative instances), we use weighted logistic regression and apply a weight to each class proportional to the class size. The resulting model clearly differentiates between shocks and non-shocks as seen by the propensity score distributions in Figure 4.1.

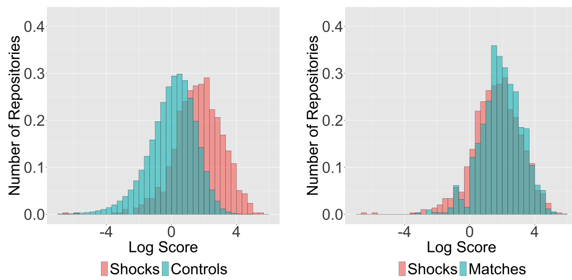

(a) (b)

Figure 4.1: Propensity Score Distributions. (a). Shocks vs. All Non Shocked. (b). Shocks vs. Matches

4.6.1.3 Matching

We now use the estimated propensity scores to define our control set. Rather than simply using the repositories with the closest propensity score, we use a hybrid method that, for each shock, first defines a sufficiently similar neighborhood using the linear (i.e. log-odds) version of the propensity score (10% of the pooled standard deviation $\sqrt{\frac{1}{2}\left(SD^2_{shocks} + SD^2_{not\_shocks}\right)}$) and then selects the nearest neighbor repositories of the *same programming language* in terms of the Mahalanobis distance based on the covariates (Rosenbaum and Rubin, 1985). The version of the Mahalanobis distance we employ weighs its different components by their effect sizes from the regression. This way, we prioritize those covariates that are highly predictive of the propensity score (Greevy Jr et al., 2012). If we do not find an adequate number of matches within propensity score neighborhood, we resort to selecting the remaining matches from the next nearest neighbors in terms of the propensity score itself (Rosenbaum and Rubin, 1985). Following this procedure, we select 5 matches for each of the 1,107 shocks that remain in our dataset.

Figure 4.2c shows time series of the number of daily stars and forks for shocked and control repositories vs. the number of days from the day of the shock. We observe that both groups of repositories exhibit an increase in attention prior to the shock – which is expected since GitHub selects repositories with increasing visibility for their trending page, and control repositories should exhibit a similar pattern by the design of the matching procedure. However, after the shock, control repositories return to their baseline much faster than shocked ones. This is also to be expected since control repositories were not listed on the trending page and thus do not receive an additional boost in attention.

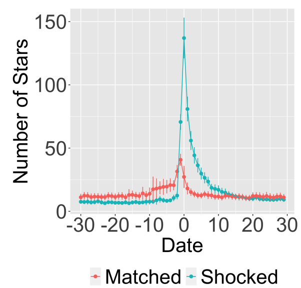

(a)

(b)

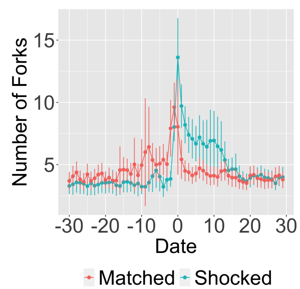

(c) Number of stars (a) and forks (b) vs. number of days from the shock for control and shocked repositories.

4.6.1.4 Stratification

Next, in order to account for any selection effects that may remain after matching, we stratify the combined shocks and matches dataset using the linear propensity score such that the behavior of shocked repositories is compared only against that of matched repositories in the same strata (Imbens and Rubin, 2015). We apply the following procedure:

(1) Remove all matches that have a linear propensity score (PS) less than the minimum PS among the shocks (PS_{min}) (11 matches removed).

(2) Remove all shocks that have a PS greater than the maximum PS among the matches (PS_{max}).

(3) The initial block is defined by [PS_{min}, PS_{max}]

(4) Repeat the following steps until there are no blocks left to split

    (a) For each block, calculate the t-statistic (**t**) for the mean difference in PS between shocked and matched repositories

    (b) If $\mid t \mid \leq 1$, treatment is sufficiently uncorrelated with PS. There is no need to further split the block

(c) If $| \ t \ |> 1$, Split into two blocks along the median PS of the block if each new block has at least $max(3, K + 2)$ observations, where K is the number of covariates.

This approach splits our data into 6 consecutive blocks with different numbers of shocks and matches (Table 4.2). All except the smallest block retain approximately the 1 to 5 ratio between shocks and matches. In the smallest block with 44 shocks corresponding to the highest PS scores there are only 3.7 times as many matches.

Table 4.2: Propensity Score Block Statistics

Block	**Lower Bound**	**Upper Bound**	**Shocks**	**Matches**	**t**
1	-2.56	4.88	548	2768	-0.33
2	4.88	12.58	278	1379	-0.21
3	12.58	24.60	140	688	0.11
4	24.60	39.54	65	350	0.28
5	39.54	57.64	32	176	-0.02
6	57.64	375.5	44	163	0.96

4.6.1.5 Balance

We verify that matches are sufficiently comparable to the shocks by evaluating the balance of the covariates using the standardized mean difference (SMD). First, we estimate the maximum value that the SMD of a covariate can have if group membership (shocked or matched) accounts for less than 1% of its variance using the formula $r = \frac{d}{\sqrt{(d^2+1/(pq))}}$, where r is the square root of the variance, d is the SMD, p is the fraction of shocks and $q = 1 - p$ (Cohen, 1988). Considering the 1:5 ratio between shocks and matches overall, we arrive at an upper bound for the SMD of 0.269. For each covariate, we estimate an overall SMD across all the blocks by using the weighting scheme from (Imbens and Rubin, 2015) that takes into account the number of shocked and matched repositories as well as the distribution of the covariate values in each block. For all our covariates, SMD calculated in this manner satisfy the estimated upper bound constraint. The maximum covariate SMD is 0.22 and the mean, when weighted by their corresponding effect sizes in the regression, is 0.13. This indicates that match quality is substantially better for covariates strongly associated with the propensity.

4.6.2 Hypothesis Evaluation

4.6.2.1 Difference-in-Shocked Analysis

We evaluate the first order behavioral changes in the treated repositories before and after the shock. For a behavior, y, we measure the change in a given treated repository as $\tilde{y}_t - \tilde{y}_{t-1}$, where $\tilde{y}_{t-1}$ and $\tilde{y}_t$ are the median measurements of y before and after the treatment respectively. We evaluate the statistical significance of this difference using Kruskal-Wallis one-way variance test (Kruskal and Wallis, 1952). Given that the control repositories also exhibit increased outside engagement, the observed response may be driven by a combination of organic growth and the attention shock. Thus, in order to the disentangle these effects, we also employ a diff-in-diff model.

4.6.2.2 Difference-in-Difference Model

We use a difference-in-difference (DD) regression approach (Abadie, 2005) to compare the changes in shocked and matched repository behaviors to test our hypotheses described in Section 4.3. Difference-in-Difference is a quasi experimental approach that estimates the effect of a treatment in observational data by comparing against the outcome in the treatment set to a counterfactual derived from a suitable control set. We use this measure to characterize the change in each of the behaviors of interest (described in Section 4.5) 30 days before and after the shock. In order to account for propensity score based stratification in our model, we use a fixed effect for the block to which an instance belongs. This mixed effects model is summarized by the following:

$$y_{it} = \beta_0 + \beta_1 t + \beta_2 I + \beta_3 tI + \beta_4 B + \beta_x C + \epsilon$$

where y_{it} is the outcome for repository i for behavior y, t is an indicator variable for time period (0 during the period before the reference date and 1 after it), I is an indicator variable for treatment (1 if i is in the treated set or 0 otherwise), B is the PS block to which i belongs, $\beta_x C$ is a placeholder for controls and ϵ is residual error. The coefficient β_3 represents the difference in changes in outcome y_{it} over time between the treatment and the control group.

4.7 Results

In this section, we describe the results for our hypotheses based on the *difference-in-shocked* and *difference-in-difference* estimates.

4.7.1 H1: External Engagement and Contribution

4.7.1.1 Engagement

Figure 4.3 shows that the number of stars and forks received by a repository increases substantially after being shocked. The difference estimates show that a shocked repository receives approximately 107% more stars and 67% forks than before on average. Also, this increase is even larger when compared with the control (237% for stars and 131% forks). This suggests that attention on shocked repositories is growing rapidly, while the attention on the control repositories is rapidly returning to baseline.

4.7.1.2 Contribution

We also observe $\approx 50\%$ increase in the number of users who performed at least one action other than starring and forking (Figure 4.3). Further, this observation remains statistically significant for higher thresholds of 2 (33%), 3 (28%), and 5 (7.5%) actions. We observe similar but smaller effects for the number of pull requests (5%) and issues opened (25%). Similar to our observations with low-effort engagement, the increase in external contributions in shocked repositories is even larger when compared with the controls which seem to be slowing down. However, when considering the number of files edited by external contributors we find that the external contributions are not scaling at the same rate. In fact, we see that on average an outside contributor edits 37% fewer files and that this effect persists, to a lesser extent (10%), when compared with the control set. This provides a more nuanced perspective of increasing outsider contribution, where most contributions from the expanded outsider community are relatively shallow. Taken together, our results support hypotheses 1(a) and partially support 1(b) – there is a large increase in engagement, in terms of software interest, but also in willingness to contribute, albeit shallow ones.

4.7.2 H2: Backlogs

With external engagement increasing, we expect members of shocked repositories to be overwhelmed by the growing number of tasks that need their attention. While, some evidence of this emerges, our observations depict a more complex portrait of the situation (Figure 4.4). After the shock, the amount of time it takes a member to respond to a new issue or a pull request increases substantially, by 30% and 42% respectively. This suggests that members are indeed feeling the strain of an increased workload. This decrease in responsiveness is even larger compared against control repositories where external attention is declining when compared to the reference date.

However, we do not see conclusive evidence suggesting that members in shocked repositories are becoming less efficient in actually closing pull requests or issues. While the efficiency of dealing with pull requests seems to decrease slightly ($\approx 4\%$), this slow down is no different from the one observed in control repositories, suggesting that this slowdown is a consequence of increased organic attention, but not of the shock itself. In addition, issue closure efficiency does not show any change in shocked repositories either with or without comparison to the controls.

Overall, members of shocked repositories are slower with regards to their initial response towards external engagement. However, they are still managing to keep up with the workload by following through on tasks. In summary, our results partially support hypothesis 2.

4.7.3 H3: Member Administrative Orientation

How do the members of shocked repositories compensate for the increasing demands on their time from outsiders? Figure 4.5 shows that, as we hypothesized, members of shocked repositories become more administration oriented.

4.7.3.1 Admin Work

We observe that members are trying to keep up with the increased administrative overhead of engaging with outsiders. The number of issues and pull requests they close increase substantially after the shock, with the number of issue closures increasing by 33% and pull request closures increasing by 11% after controlling for the number active items (issues or pull requests) during each period.

4.7.3.2 Development Work

Members of shocked repositories are doing substantially less development work in the aftermath of the shock. On average, they submit 38% fewer pushes and pull requests in the period after the shock. They fall behind members of control repositories even more (42%). We also inspect the lines of code (or text) changed by members as this metric can be more illuminating in this context. Again, we observe that members in shocked repositories reduce the number of lines they change—albeit by a smaller percentage of 18%. This decline is much more substantial (57%) when compared to the control set.

Overall, we find that members in shocked repositories forego development work to coordinate external engagement and contributions, providing support for hypotheses 3(a) and 3(b).

4.7.4 H4: Coordination Processes and Decentralization

Based on our hypotheses, we expect that GitHub projects that experience an attention shock will become more open and decentralized. In general, our results largely align with the expectation and provide partial support for hypothesis 4 (see figure 4.6). However, there are several important caveats and robustness tests discussed below.

4.7.4.1 Member Engagement with Outsiders

The results show that members are trying to accommodate the greater volume of incoming contributions. The number of pull requests approved increases by a 12% for shocked repositories. However, this increase is no different from that of the control set. Therefore, this increasing openness appears to be a function of the organic growth of a repository rather than the effect of an attention shock. We also do not see a significant change in the number of issues edited by members among shocked repositories. However, in comparison to the control set, the story changes — there is indeed a large increase (30%) in the number of edited issues by members in shocked repositories relative to control. This suggests that shocked repositories are maintaining their engagement at higher levels than the control set. In terms of engagement on external pull requests, we observe that individual members in shocked projects edit moderately fewer pull requests (17%) than before, but, in fact, they are editing more, almost 44%, relative to the control set. We find support for hypothesis 4(a), but the effect is due to both natural growth and the shock; and the specific type of engagement (issues or pull requests).

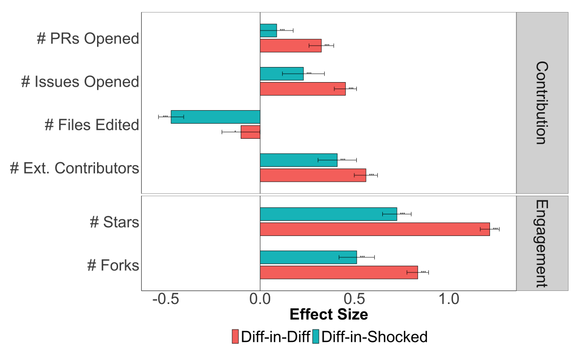

Figure 4.3: External Engagement (H1). The bars show the effect sizes for the Diff-in-Diff regression, which represents the effect of the shock on each measure relative to the control group as a percentage, as well as the difference in medians before and after the shock for shocked repositories only. The error bars show 95% confidence intervals. The stars show significance (*** p-val $<$ 0.001, ** p-val $<$ 0.01, * p-val $<$ 0.05). Other bar plots in this section follow the same format.

4.7.4.2 Decentralization

We find evidence for hypothesis 4(b). Members of shocked repositories do indeed allow outsiders to take on more central positions in work, thereby adopting a more decentralized brand of coordination. However, there is one important caveat. First, we observe that, among shocked repos-

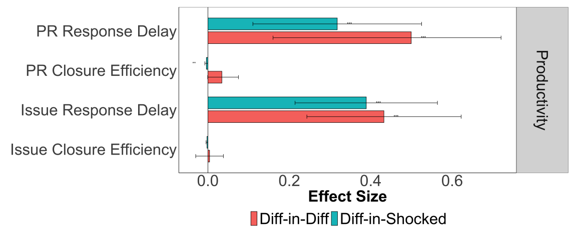

Figure 4.4: Backlogs (H2)

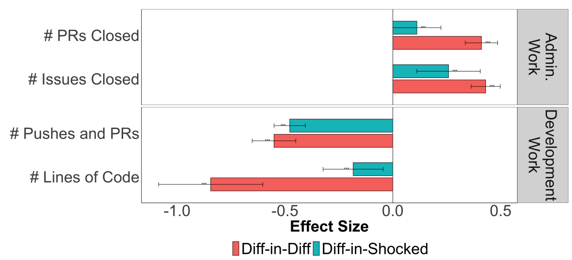

Figure 4.5: Member Admin. Orientation (H3)

itories, the fraction of outsiders among the top 50% contributors with most files edited increases by approximately 8%. This increase remains roughly the same in magnitude when compared with control repositories (7.2%), indicating that, in the aftermath of the shock, outsiders are allowed to make more substantive contributions than they would have in other times of increased attention. Robustness tests show that this phenomenon persists at a higher threshold of 40% but not beyond. However, in contrast to this observation, we observe that the fraction of outsiders among the top 50% most central contributors with highest degree centrality in the file collaboration network does not change in the shocked repositories with or without comparing to the controls. Therefore, it appears that, while outsiders are being allowed to make greater contributions, these contributions may be relatively isolated from core parts of the project, where development is likely to be more collaborative.

Less experience and effort is required to contribute to discussions surrounding issues and pull requests than to contribute code. Therefore, we expect that it would be easier for outsiders to take on more central roles in the collaboration related to issue and pull request discussions. Our results for the fraction of outsiders among the top 50% contributors with highest degree centrality for issue and pull request collaboration networks show that this is indeed the case. We observe that the number of outsiders among the most central contributors increases by 3% in the issue collaboration network and by 5% in the pull request network. In both cases, this effect is amplified when compared against the controls (7% and 8% among issue and pull requests networks, respectively), indicating the increased external attention drives decentralization. Additional robustness tests show that the increasing importance of outsiders in the issues collaboration networks is con-

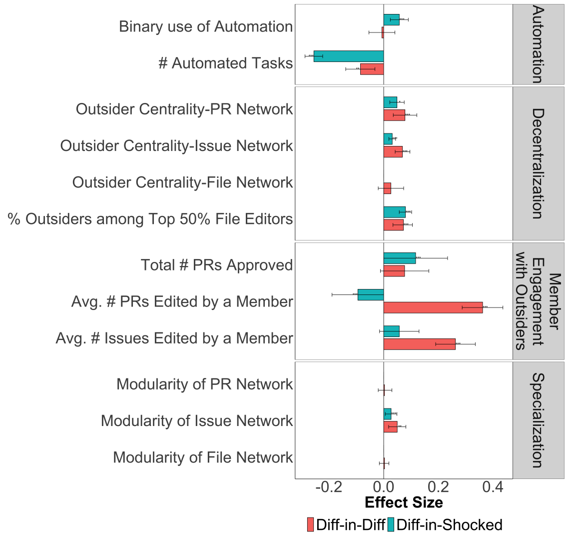

Figure 4.6: Coordination Processes and Decentralization (H4). Missing bars indicate that all values for the corresponding measurement are zero.

sistent even at higher thresholds of top 40%, 30%, 20%, and 10%, while it extends only up to the top 40% in the pull request network. Thus the centrality of outsiders is more persistent with issue discussions which require less effort and expertise than pull request discussions. This serves as further evidence that outsiders work within a hierarchy of commitment and experience as they become more central.

4.7.4.3 Modularity of Collaboration

Our results provide partial support for hypothesis 4(c), which predicts increase in specialization by members and outsiders to deal with coordination costs. First, in the file and pull request networks, which are associated with contributors with relatively high levels of commitment, we observe no effect on the modularity of collaboration in the shock set, with or without a comparison to the control set. However, in the issue collaboration network, which is associated with relatively low effort, we do see a small increase in modularity of 7.5%, which remains true when compared with the controls (5%), indicating that some level of specialization does occur. Why does file and pull request collaboration, which stands to benefit perhaps even more than issues, not show any effect? We argue that this may be due to the relatively high complexity of these contributions. Code contributions and discussions about pull requests require more effort and expertise, which take more time to build up compared to issue discussions, which often serve as landing site of general queries.

4.7.4.4 Automation

Regarding automation, our results provide only partial support for hypothesis 4(d). Shocked projects become slightly more likely to adopt automation if they haven't already (6%). However, this increased likelihood is no different in the control set. This suggests that the shocked projects that adopted automation, would have probably adopted automation anyway as part of their natural growth. Further, contrary our expectation, the number of automated tasks in shocked repositories decreases by 20% in the aftermath and this remains true to a lesser extent even when compared to the controls (8%). What explains this pattern? Organic growth that leads to a shock is sometimes associated with a concerted drive in development to produce a software release. The number of automated tasks right before a release (and the shock) could be substantially higher than at other times. This would inflate automation in the *before* period. Future work should determine if this potential mechanism is indeed at play.

4.8 Limitations

Our paper has several limitations. First, we only consider the short-term impact of the shock. Whether the shock has a long lasting impact remains as an open question. Second, our analysis is based on logged user actions such as committing and submitting pull requests, but ignores the text in comments or measures of code quality in pull requests. A more detailed analysis of this content could uncover new dynamics that we do not capture here. Third, we test our hypotheses independently of each other. While this approach is sufficient to characterize the effects of the attention shock on the collaboration dynamics, it is possible that there are dependencies between our collaboration measures. For example, the level of modularity in the issue collaboration network may impact the time to address issues. Future work can examine interactions between our measures through the use of structural equation models. Fourth, we do not consider the impact of repeated shocks on a repository. Additionally, since we do not have access to trending data prior to our 7 month period, we cannot be sure that the repositories that were trending during our period were not trending prior to it. Fifth, we ignore aspects such as the application and quality of repositories when selecting our controls. This could affect the quality of matching. However, because we selected controls based on GitHub's statements about trending repositories and covariate balance was within acceptable limits based on standard metrics, we deem our control dataset to be adequately similar to the trending projects. Finally, we characterize the response by repositories to attention shocks but do not address whether the observed responses are beneficial or harmful. Future work should uncover which types of responses lead to desirable outcomes such as higher quality and increase in adoption.

4.9 Discussion

Our research expands the current understanding of open source software crowds by studying how they respond to attention shocks. Trending GitHub crowds actually respond or adapt similarly to successful organizations despite their differences. Based on the analysis using the distributed leadership framework, core members attempt to adapt their coordination style by distributing work responsibility. For, example, they allow outsiders to take more central yet unofficial positions. They also increasingly adopt automation as a means of enforcing community norms and work practices. The project also self-organizes into a modular structure while core members restructure their work. Despite this apparent positive response by core members there are some signals that they are being overwhelmed, at least in the short term. Due to the increase in contributions from outsiders, core members strain to coordinate work and begin to accumulate backlogs, which lead to delays in responding outsiders. This study provides implications for GitHub crowds responding to attention shocks. First, most of the engagement from arriving outside contributors are shallow. This has the effect of increasing the work of core members, making it potentially difficult to leverage the efforts of these arriving outside contributors. This begs the question: How can core members make the most of such increased resources? One approach is to make easily doable small tasks highly visible through public to-do lists that outsiders can work on. This can help outsiders easily identify urgent needs of the project and make it more likely that their contributions will be accepted.

Second, despite the best efforts of core members, a backlog of pull requests and issues pile up, and while they eventually close them, there are long delays in responding to them. This low responsiveness has the potential of dissuading the contributions of outsiders, including those who could have become productive contributors. How can crowds being impacted by attention shocks prevent low responsiveness from running off arriving outsiders? One possibility is to automate responses to pull requests and issues using boilerplate responses (e.g. Thanks! Will get back to you) (McDonald et al., 2014). An automated response might give outsider the sense that members are aware of their contributions and will eventually review them.

Finally, new work groups emerge and arriving outsiders become increasingly central. This may explain why members are able to keep up with closing issues and pull requests. However, this is more likely to be a spontaneous self-organized response rather than a planned strategy. An open question is how GitHub should manage the expectations of members of repositories that are listed on the trending page. Our findings suggest that this can create both opportunities for more attention and contributors, but also a potential disruption to their routines. Providing members with a warning of what is likely to come their way could help them be more prepared to herd the deluge of good Samaritans.

CHAPTER 5

Conclusions

5.1 Summary of Contributions

5.1.1 Overview

In this dissertation, we have utilized quasi-experimental study design in combination with large-scale digital trace data from the Social Web to investigate how people respond to the challenges brought on by sudden external events in two different settings: (i) natural disasters and (ii) online crowds experiencing a sudden burst in popularity. As a whole, these studies demonstrate the potential for the combination of quasi-experimental study design and "Socia Web" data to advance the causal understanding of a domain that is generally infeasible to study experimentally. These studies, each using a different quasi-experimental setup, also highlight the validity issues that this framework needs to account for, especially when generalizing causal claims beyond the online settings they are based on to offline behavior Diaz et al. (2016); Oktay et al. (2010).

5.1.2 Adaptaton to Exogenous Shocks

Human society is increasingly vulnerable to frequent cascading shocks due to coupled social, political, technological, and ecological networks. As a consequence, top-down reactive interventions, a mainstay of government policy toolboxes, have become less effective and may even induce undesirable adaptations (Rabe, 2004; Taleb, 2007; Trump et al., 2017). In this context, it is important to understand emergent human adaptations to adverse external events,in different individual and collective settings, the desirability of the corresponding outcomes, and the attributes and processes that lead to desired outcomes. This knowledge can inform proactive policies that focus on building social resilience in different settings.

In this dissertation, we find that people demonstrate a range of emergent adaptive behaviors in the face of external shocks that disrupt their normal functioning and that these adaptations take place largely through the leveraging and reconfiguration of social linkages.

In chapter 2, we find that communities experiencing disasters rely on a normally diffuse sense of collective identity that comes into sharp focus in the event of a disaster as a focal point to respond cohesively in the immediate aftermath. Most communities appear to weather the effects of disasters well at least in the short term, though a minority seem to experience much more substantial negative effect in terms of long-lasting and recurring patterns of negative affect as well as an increasing engagement in sensemaking overtime, indicating that a continuing struggle to come to terms with their new reality. Beyond these differences, community responses follow prototypical pathways along social, aspirational, biological, and physical dimensions. Both these prototypical patterns and the differences that suggest variations in the level of successful adaptation can inform strategic community resilience planning at the state and regional levels. Furthermore, our findings demonstrate the potential to use, with careful calibration, social media traces to track and forecast community response trajectories and, by extension, to evaluate the efficacy of relief and recovery activities.

In Chapter 3, we find that a person's decision to relocate after a disaster is associated with the social capital reflected in their social media. We find that bonding capital, operationalized as derived from relationships that are highly embedded in the overall network, increases the chances of a person rebuilding after a disaster. This aligns with the expectation that a person's immediate social circle is likely to provide them with material, emotional, and other support necessary to rebuild. In contrast, another different bonding capital construct, that of durability of strong ties, increased the odds that a person would relocate. This may be due to a different mechanism, one in which strong ties, such as family members, are more resilient to the effects of distance, helping to make a smooth transition to a different region (Engel and Ibáñez, 2007). This argument is further supported by the observation that the presence of community institutions, a coarse-grained measure of localized social capital, also increases the likelihood of relocation.

Finally, in Chapter 4, we find that collaborative crowd communities demonstrate adaptive behaviors that incorporate elements observed in both individual and much larger communities that interestingly somewhat align with adaptations of traditional organizations. In the face of a large influx of users and potential contributors resulting from algorithmic amplification, members of these online social organizations seem to adopt a distributed leadership approach that is characterized by (i) centralized coordination and gatekeeping and (ii) delegation of production to a newly recruited decentralized modular workforce. The former strategy parallels collective identity-based coordination in communities in the aftermath of a disaster, while the latter strategy is similar to individuals taking advantage of their social capital. Despite these adaptations, these communities struggle to cope with the deluge of engagement, as evidenced by the decrease in short-term responsiveness.

This dissertation has not directly addressed the interactions between different social levels that simultaneously produce emergent behaviors across them. However, our findings indicate that similar social processes play an important role in determining adaptations at different social levels. Importantly, we find that the same mechanisms that lead to desirable outcomes may produce undesirable outcomes at one end. We observe this phenomenon in Chapter 3, where individuals with access to bridging capital, as well as other socioeconomic resources, find it more desirable to relocate, while individuals who lack those material resources and only have geographically localized social capital are more likely to remain. From the perspective of the first group, relocation is a resilient adaptation. But from the perspective of the community, losing some of its members after a disaster has disrupted its social fabric and strained its material resources, it is an adverse adaptation. Even at scales beyond the focal community, this leads to geographical sorting that exacerbates issues of inequality. Considering these implications, future work should to build on our findings to investigating multiscale interactions to discover regimes where adaptive behavior leads to resilience across all scales.

5.2 Future Directions

The goal of the studies in this dissertation has been to advance the causal understanding of how people respond to external shocks in different settings by utilizing online data for large numbers of events within a quasi-experimental study design. In addition to making substantive contributions to the relevant domains and to the field of computational social science in general, this type of work needs to enrich the policy discussions on building resilient physical and online communities. Our studies fall short of this yardstick of policy impact to various extents. Our investigations of how geographically co-located populations (communities) and individuals respond to natural disasters contribute temporally fine-grained and nuanced narratives to the domain. However, the findings and implications of these studies are tempered by issues of selection bias when we try to generalize them to the offline behavior of people affected by disasters Garcia and Rimé (2019). While the studies utilize some methdological strategies and arguments based on prior validation work in the domain, they do not quite meet the high standard necessary for directly informing policy without further improvement. In contrast, our study of how online crowds adapt to the sudden popularity brought on by a platform announcement does not suffer from this internal validity concern to the same degree. This study focuses on behaviors that are native to the online environment that has no direct equivalent in the offline setting and as a consequence, within that narrow context, quasi-experiments with online trace data can extract causal narratives that are immediately meaingful in platform-level policy discussions.

This dissertation tackled the first two of the three substantive questions associated with human adaptation to external shocks that were introduced in Chapter 1 – (i) what are the different ways in which people respond? and (ii) what intrinsic and environmental factors drive the differences in responses. However, we did not directly address the third question regarding which responses constitute resilient adaptations beyond the ones discussed that leveraged prior work. It is this question that truly highlights the policy relevance of this domain. Those responsible for developing policy for supporting resilient communities, whether they are state policy makers allocating resources and developing strategies to shield physical communities from natural or manmande disasters or owners and designers of online social or work platforms, need to know which types of behavior from affected people lead to “desirable” outcomes and how types of interventions could turn different causal knobs to ensure people adapt successfully.

In the following sections, we consider how future work can build on the studies discussed here and incorporate new methodological advances from different domains to address these key limitations that limit policy impact.

5.2.1 Fidelity of ”Social Web” observations as reflections of offline behavior

One of the key limitations of the work discussed in this dissertation is that beyond the relevance of people’s behavior on the Web in and of itself as an important dimension of behavior, any applicability of findings made by utilizing digital trace data to offline behavior is limited to the extent that those behaviors are an accurate and unbiased reflection of it. Despite a host of recent work that has utilized social median and other digital trace data to investigate behavioral phenomena that have implications for both online and offline behavior and the often implicit assumption that online behavior is adequately representive of offline behavior, there remains a clear gap in knowledge regarding the mapping between online and offline behavior in many important behavioral settings.

Although research can be performed to build a comprehensive understanding of this mapping between the on-line and the off-line, across any and all dimensions of interest — a valuable, if monumental task — in our context of understanding human responses to unexpected challenges, a more focused approach is appropriate. Several previous studies that used trace data such as call detail records have resolved or at least mitigated this mapping issue using mixed methods strategies (Blumenstock et al., 2015; Aiken et al., 2022). These studies employ surveys, with stratified sampling along the dimensions where representativeness is important, to collect data regarding offline attributes and behaviors, including those that correspond to digital trace data that are used in the studies. The survey data are then linked to trace data either directly through unique identities or through common measurements that can be estimated from both the survey and the trace data. This mapping can now be used to calibrate trace data through techniques such as transfer learning. Similar approaches need to be used to strengthen the case for the applicability of our findings more broadly beyond the virtual sphere.

5.2.2 Resilience of adaptive responses

In this dissertation, we investigate how people, individually and collectively, adapt to unexpected challenges and to a more limited extent how different intrinsic and environmental attributes determine which responses materialize. The next question that naturally follows is which adaptations result in desirable outcomes for focal individuals and collectives, as well as the externalities associated with them. As briefly discussed in Chapter 3, individually resilient behavior such as a person easily relocating to greener pastures can be maladaptive from the perspective of the larger community (loss of resources, geographic segregation, etc.). There is substantial theoretical and qualitative work, particularly at the individual level, that has associated intrinsic attributes of people and their adaptive responses to unexpected challenges to resilient or negative outcomes (Bonanno et al., 2012; Norris et al., 2002). However, there are relatively few studies that have employed empirical strategies based on measurable constructs to develop links between different responses and the desirablility of outcomes, particularly when observational behavioral data are used (Garcia and Rimé, 2019). A key reason for the dearth of such studies is that it is often challenging to determine the subjective desirability of response behaviors using the same data being used to identify those behaviors. For example, in Chapter 3, we were able to distinguish between individuals who chose to relocate from an area affected by disaster and those who did not, but it was not possible to determine whether those choices were resilient or maladaptive in each case. There are several strategies that can be pursued in future work to address these limitations. One potential approach that has been used in prior work (Garcia and Rimé, 2019) is to use the same behavioral dimension that is used to measure the adaptive response or a related measure that can be derived from the same data to evaluate the desirability at a later time after the response has been measured. An example is the measurement of the time taken for the elevated negative affect expressed on social media in the aftermath of a disaster to return to pre-shock conditions given the initial intensity of the negative response.

APPENDIX A

Supplemental Materials for Chapter 2

A.1 Normality of Pre-disaster Variation in Lexicon Dimensions

We standardize the intensity time series for each linguistic dimension during the 35 day (5 week) study period with the mean and standard deviation in intensity estimated from the 21 day (3 week) period immediately prior to the study period. This normalization is meant to make time series for each dimension comparable across all disasters. However, the extent which this objective is achieved relies on the underlying assumption that the pre-disaster variation in each dimension for each disaster conforms to a normal distribution. We conduct Shapiro-Wilk tests for the pre-disaster variation of each dimension across all disaster (24 dimensions x 203 disasters) to verify how broadly this assumption holds the dataset. As seen in Figure A.1 We find that at a significance level of 5%, We fail to reject the null-hypothesis of normality for at least 75% of the events for 23 of the dimensions. In the case of the remaining dimension, *family*, pre-disaster variation we fail to reject the normality hypothesis for only 64% of the events. We also visually inspected Q-Q plots for a number of dimension-disaster pairs to ensure that our intuition aligned with the p-value of the statistical test.

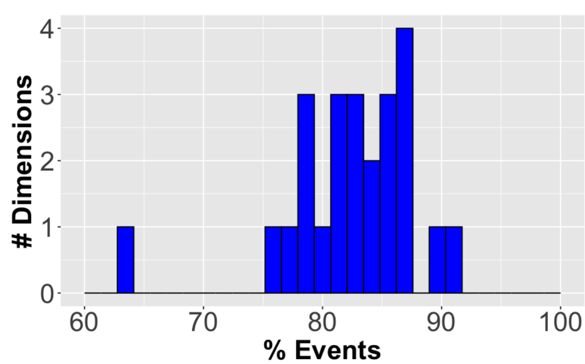

Figure A.1: Histogram of the number of lexicon dimensions in terms of the percentage of disasters for which we cannot reject the hypothesis of normality ($\alpha = 0.05$)

A.2 Clustering Quality Measurements for Different Numbers of Clusters(N)

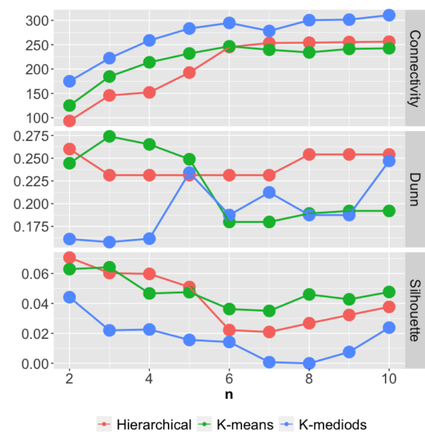

Figure A.2: Silhuette coefficient, Dunn Index, and Connectivity for $N \in 1 : 10$. Higher Silhuette coefficient and Dunn index values, as well as lower Connectivity, indicate higher quality clustering

A.3 Word-Shift Analysis of Prototypical Dimensions Before and After the Onset of Disaster

A.3.1 Risk

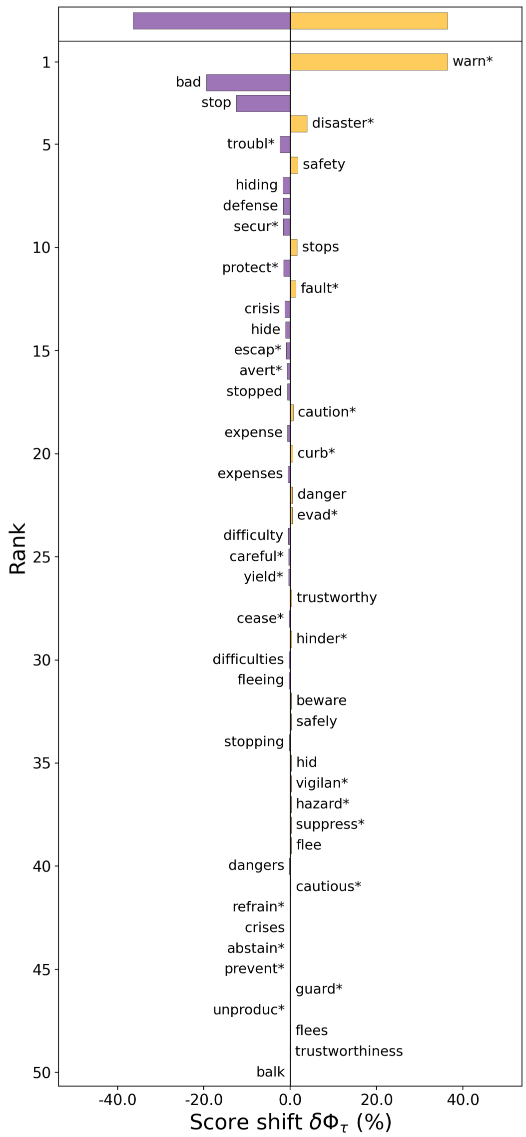

Figure A.3: Top 50 overall proportional wordshifts for **Risk**

A.3.2 Achieve

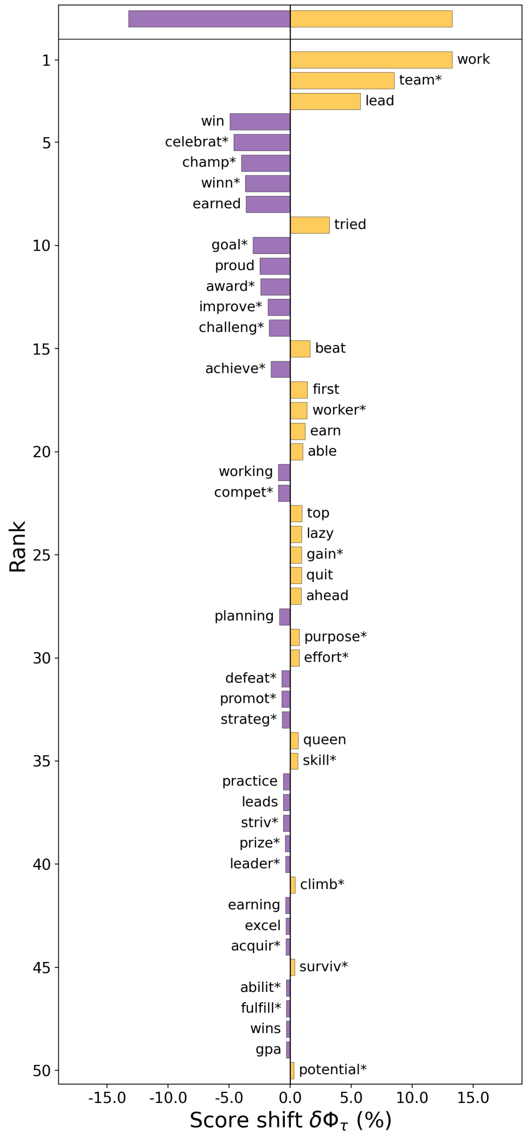

Figure A.4: Top 50 overall proportional wordshifts for **Achieve**

A.3.3 Reward

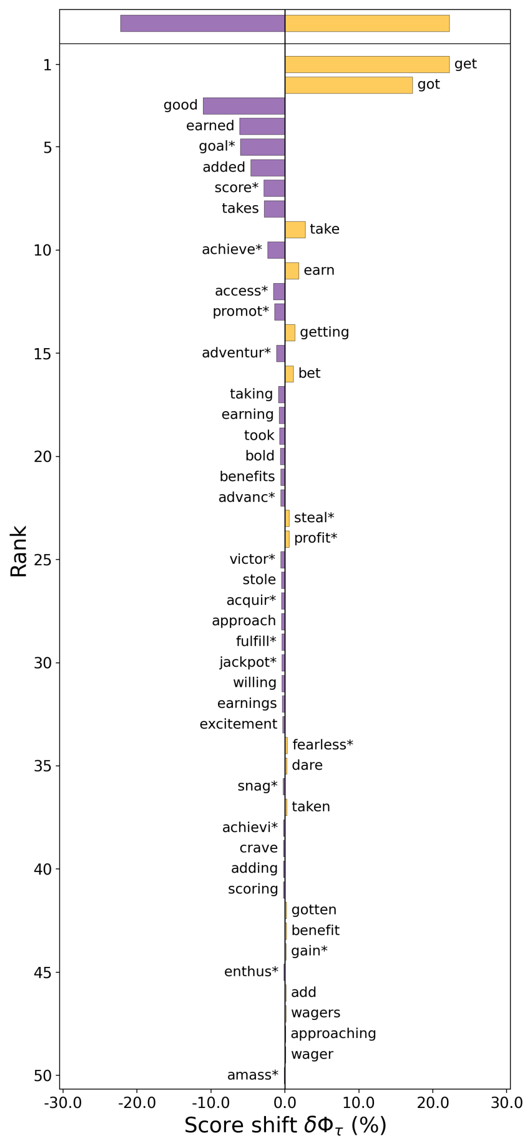

Figure A.5: Top 50 overall proportional wordshifts for **Reward**

A.3.4 Social

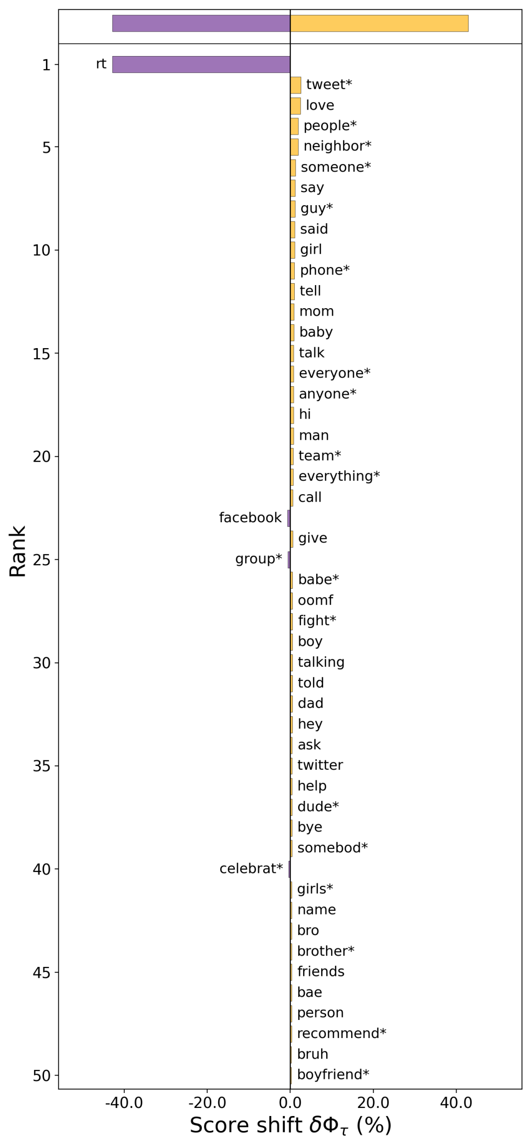

Figure A.6: Top 50 overall proportional wordshifts for **Social**

A.3.5 Health

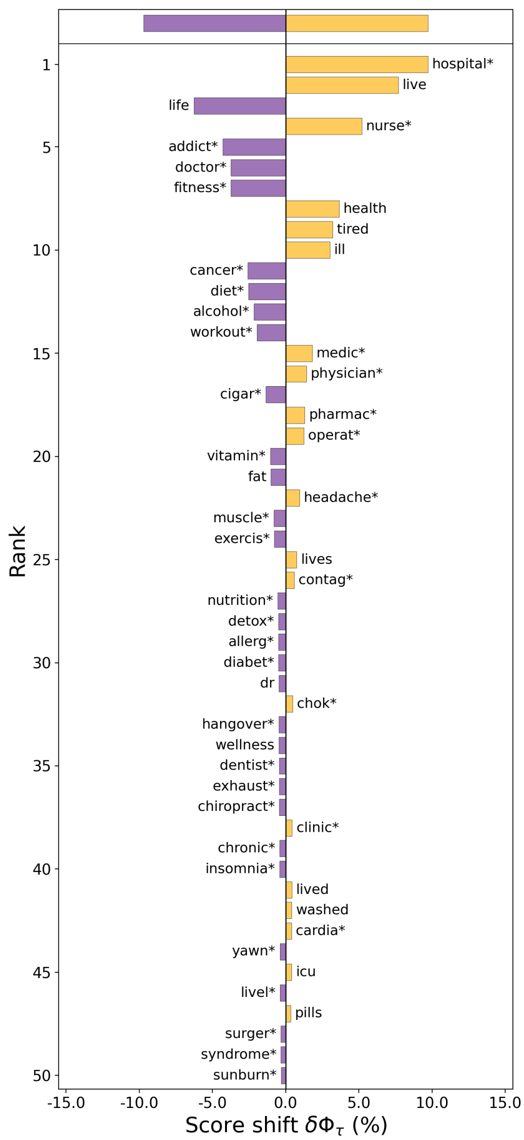

Figure A.7: Top 50 overall proportional wordshifts for **Health**

A.3.6 Money

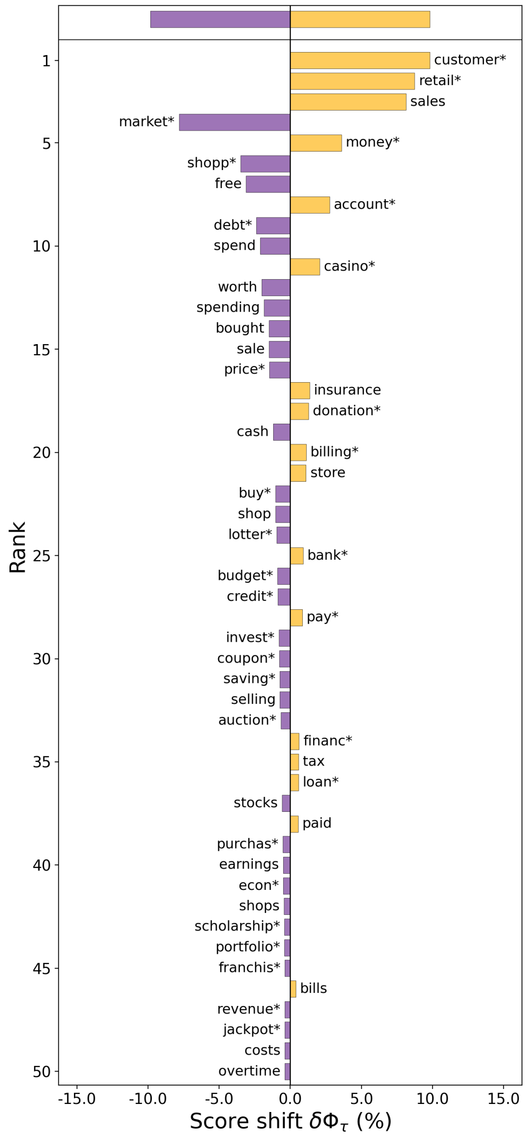

Figure A.8: Top 50 overall proportional wordshifts for **Money**

A.3.7 Work

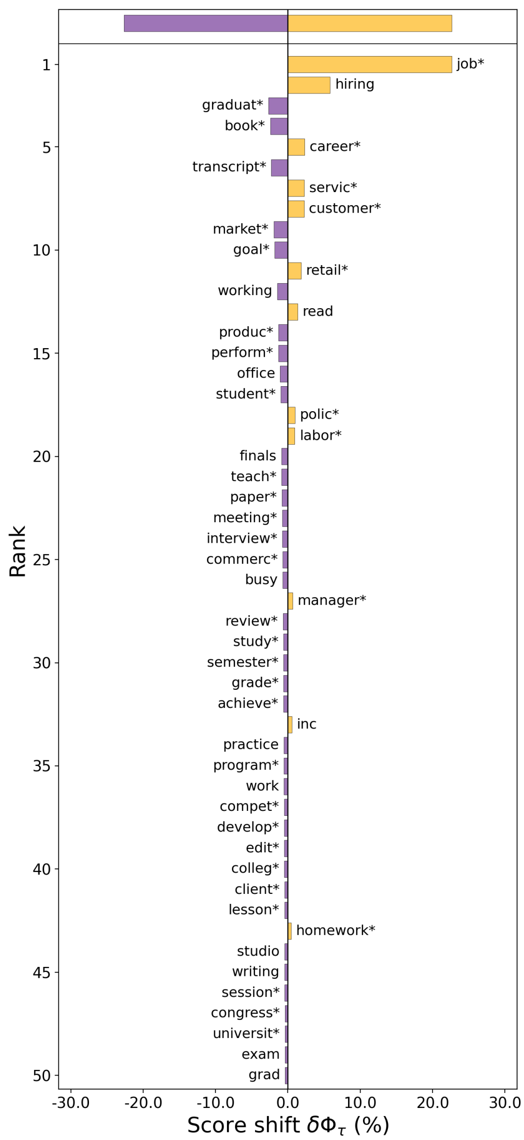

Figure A.9: Top 50 overall proportional wordshifts for **Work**

A.3.8 Home

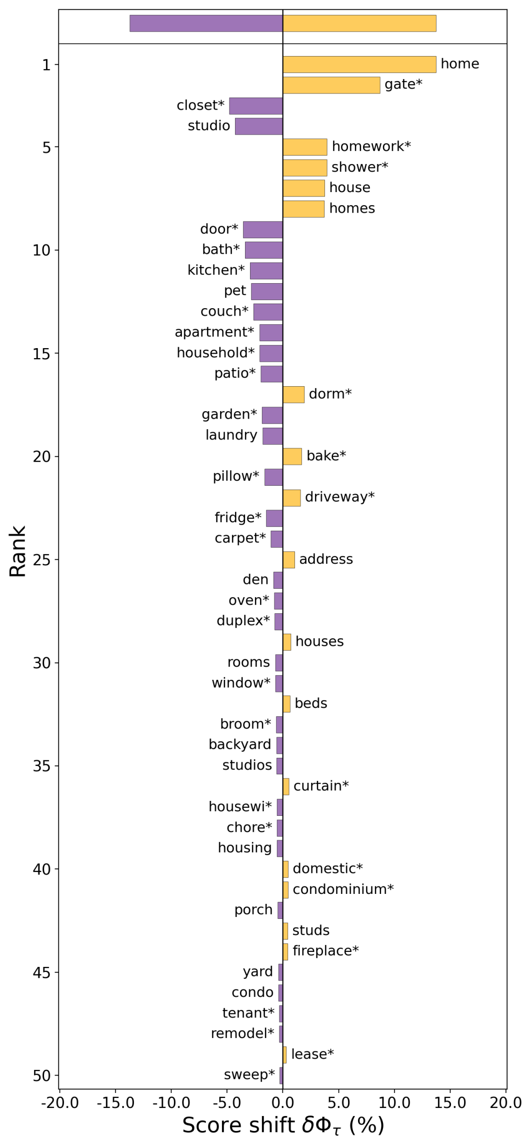

Figure A.10: Top 50 overall proportional wordshifts for **Home**

A.3.9 Ingest

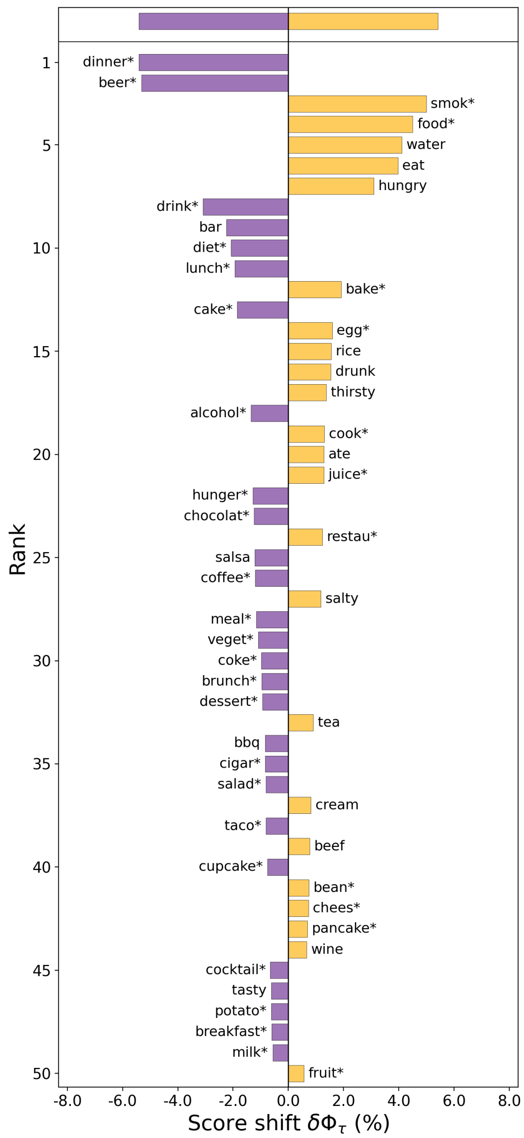

Figure A.11: Top 50 overall proportional wordshifts for **Ingest**

A.3.10 Word-Shift Analysis for Dimensions that Differ Between Clusters

A.3.11 Sadness

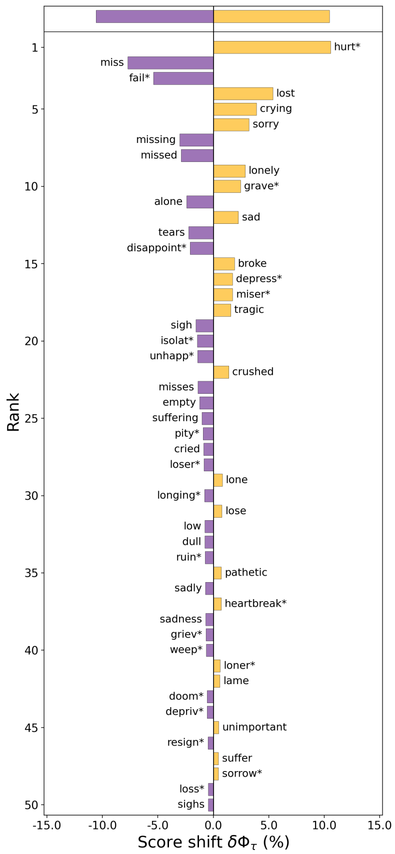

(a) Cluster 1

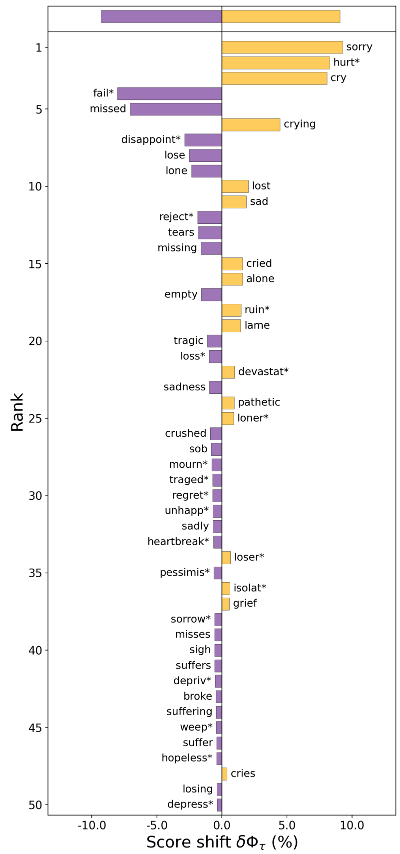

(b) Cluster 2

Figure A.12: Proportional wordshifts for **Sad** for the two clusters of disasters

A.3.12 Anger

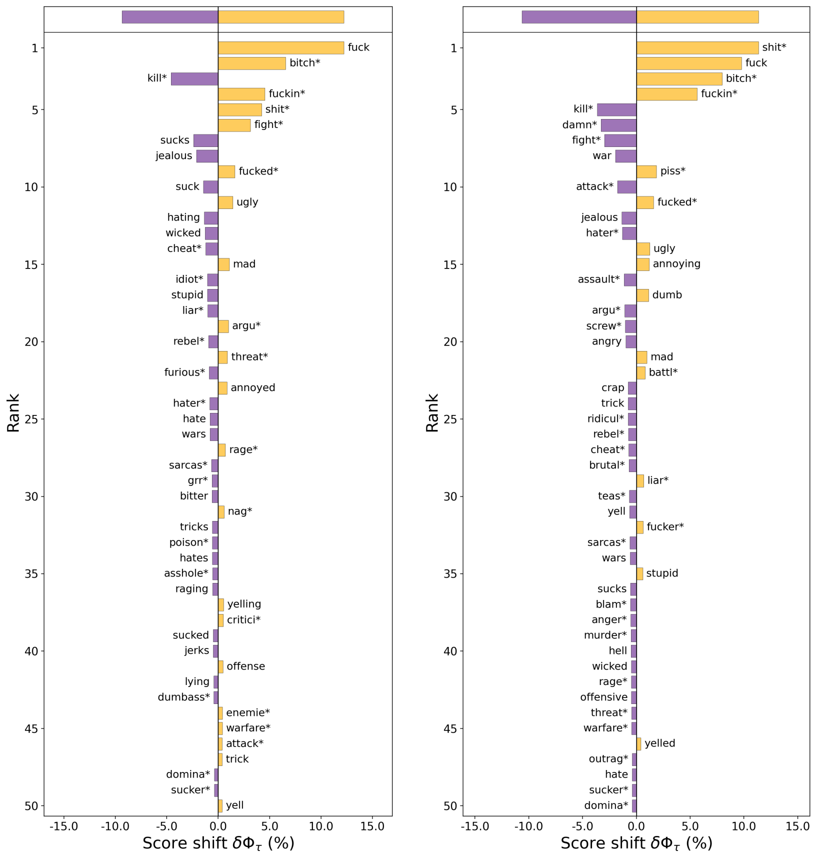

(a) Cluster 1 (b) Cluster 2

Figure A.13: Proportional wordshifts for **Anger** for the two clusters of disasters

A.3.13 Swear

We have excluded the wordshifts for *swear* as they contain words that are offensive, particularly when read out of context. The observations from these wordshifts bear a close similarity to those of *anger*

A.3.14 Fear

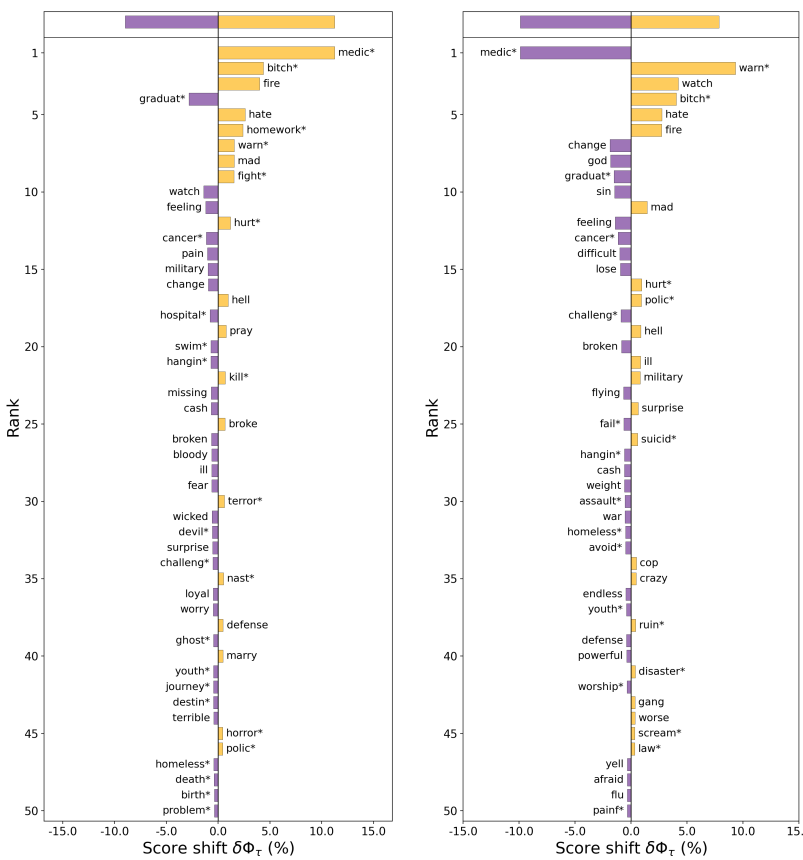

(a) Cluster 1 (b) Cluster 2

Figure A.14: Proportional wordshifts for **Fear** for the two clusters of disasters

A.4 Exemplar Tweets for Selected Lexicon Dimensions

The text of the tweets shown below have been annonymized by removing usernames and identifying hashtags as well as by paraphrasing where necessary.

A.4.1 We

(1) RT <username>: What are some inspiring scriptures about HOPE?? We all need a little bit of hope today

(2) Tornado sirens. Here we go I guess

(3) <username> I posted this once we got back to the dorms, crazy

(4) <username> Hold on! We need you all to give the rest of us hope and bring some sense to the state

(5) RT <username>: We can decide to be hopeful - don't have to wait until we feel like it

A.4.2 Ingest

(1) He is just feeling upset because he didn't get free *food* from red cross

(2) First food in three days. (at <username>)

(3) I can't drive in this how tf am I supposed to eat!!!!?

(4) Dinner is served! #paleo

(5) I'm hungry.. ima cry

A.4.3 Health

(1) ... just finished pushing my workout to the next level with <hashtag>

(2) Appointment to see the dentist today....oh joy!

(3) ... gets my mind right. <hashtag> <hashtag> <hashtag> <hashtag> #workout

(4) I've never been so tired

(5) Smooth ride to ... first doctor visit!

A.4.4 Sad in Cluster 1

(1) Tonight was seriously so depressing.

(2) I can't hold the tears back anymore...

(3) ¡username¿ this makes me so sad

(4) ... yesterday was a completely strange & weird day ¡hashtag¿ #¡city¿flood ¡hashtag¿ ¡hashtag¿ #sad

(5) Lost my ¡hashtag¿ ... poster from ¡hashtag¿ in last night's flooding. Had that thing for ... years.):

A.4.5 Sad in Cluster 2

(1) Feel so lonely ... I miss my ... cause they're there when I need them ¡sad emoji¿

(2) This is so sad they have to save people before saving animals # <city>flood

(3) Sigh I miss the sunny & hot <city>.

(4) This is the worst news I can wake up to. I'll miss you <name>.

(5) So sad to hear about the ... lost their lives during ... tornado in <state>. What a tragedy

A.4.6 Anger in Cluster 1

(1) Monday.. means I have school *shit*... my day is ruined

(2) Seriously don't tell me *shit* like that and then ignore me

(3) so many people give him *shit*, just because he's deaf. like what the fuck?

(4) <username>: First, no one left early. I was there. Second, stop talking about this *shit*

(5) buses are always up in my *shit* until I need to catch one

A.4.7 Anger in Cluster 2

(1) People in <state> only give a *fuck* about others when disasters happen.

(2) <username> it's *pissing* me the *fuck* off!!

(3) Ugh *hate* feeling helpless

(4) It's *frustrating* when people try to politicize a horrific disaster

(5) I *hate* the way <state> handles disasters #Unorganized. Lack of communication makes my job easier not harder

A.4.8 Anxiety in Cluster 1

(1) I've never been so *scared* in my life from thunder.

(2) <username> I'm starting to get *scared* now

(3) RT <username> Most attention being paid to Sunday's heavy rain *threat*, but locally heavy rain is possible today too...

(4) I was so depressed and upset yesterday, I went to bed by 5:30 and slept until 4am. Woke up still angry

(5) Overwhelmed trying to figure out how I'll catch morning flight w/o transit or shuttle to airport. Here's hoping service resumes.

A.4.9 Anxiety in Cluster 2

(1) <username> Yeah, my blood *pressure* is up

(2) Days like this can really raise my *pressure*

(3) Idk what's going on but I am lost and *confused* :(#badday

(4) Really *upset* right now. I miss <country>. I miss warmer weather not shitty rain and the cold. I miss my family!

(5) News showing a tornado hit right where my boyfriend is at work. About to have a *panic* attack. :(

A.4.10 Body in Cluster 1

(1) While you were *asleep* ...#<state>flood

(2) Thank you flash flood warning for giving me a *heart* attack as I am trying to fall *asleep*! I am now wide awake

(3) I am so excited about getting to sleep in MY OWN BED tomorrow

(4) Its not incense, its a heating pad filled with seeds and lavender to help sore *muscles* and headaches

(5) This chilly weather and painful ribs just make me want to crawl in bed and snuggle all afternoon

A.4.11 Body in Cluster 2

(1) I've had one shower and only 8hrs *sleep* maybe in three days

(2) Gooood news! My sunburn is gone and I'm not *peeling*. perks of flooding your *body* with water and 18+ layers of lotion applied!

(3) house flooded and alone at a hotel? take selfies with your bare *face* <url>

(4) RT <user>: My *heart* goes out to everyone impacted by tornadoes in the <city> area. Stay safe and #PrayFor<state>.

(5) Don't think I'll be going to sleep anytime soon

A.5 Forecasting Community Response using Cluster Trajectories

We adapted an approach inspired by cross-validation for statistical models to evaluate the usefulness of the clustered community response trajectories in forecasting the community response for future events.

We repeated the time-series clustering N (=203) times, each time leaving one event out. As with the original results discussed in the paper, each time, we used hierarchical clustering with DTW distance and set the number of clusters to 2. Then, for each clustering result, we estimate the DTW distance between different segments of the unobserved response that begin from the start of the study period and the representative signatures of the two clusters. Specifically, we considered three segments of length 1,2, and 3 weeks (i.e. $0 \leq D < 7$, $0 \leq D < 14$, $\leq D < 21$). We select the cluster that has the smaller DTW distance to the observed segment of the test event as our prediction of how the community response of that event would have evolved in the unobserved period.

For each observed segment length, we compare the similarity of the unobserved segment of community response to the selected cluster, the remaining cluster, and a baseline prediction. The baseline prediction was generated by individually fitting ARIMA models to the different dimensions (23) of the observed segment [1]. As figure A.15 shows, the response cluster that was most similar to the observed segment of the test event is a better predictor of how the community response evolved from there during the unobserved period consistently for more than 75% of the events across different lengths of the observed period. The performance of the cluster prediction relative to the baseline model is best when a moderate amount of the community response has been observed. In Figure A.15, these correspond to training horizons of Day 7 and Day 14 (i.e. lengths of the observed period is 2 and 3 weeks, respectively). This suggests that in the early stages of a community response to disaster, cluster assignment benefits more from incremental information. The comparative decline in cluster-based predictive performance at the training horizon of Day 21 (i.e., a 4 week observed period and predictions are made for only 7 days), is likely to due to the baseline model continuing to be fit to the new data, whereas on average the influence of those data on cluster assignment decreases. Also, the cluster trajectories themselves are fixed and do not benefit from additional information. It also remains an open question as to whether the baseline model would do better or worse than it is now (winning less than 25% of the one-versus-one comparisons against cluster prediction), if the time for which community response was longer than one week while holding the observed period at 4 weeks.

In addition to comparing the forecasting performance of the assigned cluster against the baseline model, we also conduct a three-way comparison between both clusters and the baseline model. Figure A.16 shows that the baseline model performs even worse than before when compared with the forecasts of both clusters. i.e., even when the baseline model provides a better forecast than the prediction of the assigned cluster, it is frequently beaten by the forecast offered by the cluster that was not chosen. This suggests that the two cluster-based predictions, when taken together, as primary forecast and an alternate forecast, may be more useful than only the forecast from the assigned cluster.

[1] We used the **ARIMA** function from the R package *fable*

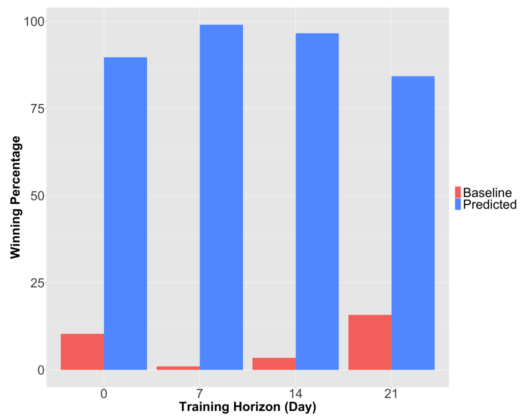

Figure A.15: Comparison of forecast performance between the predicted cluster trajectory and the baseline forecasting method. The heights of the bars indicate the percentage of times each method provided a better forecast of community response to a particular event for different lengths of observed segments

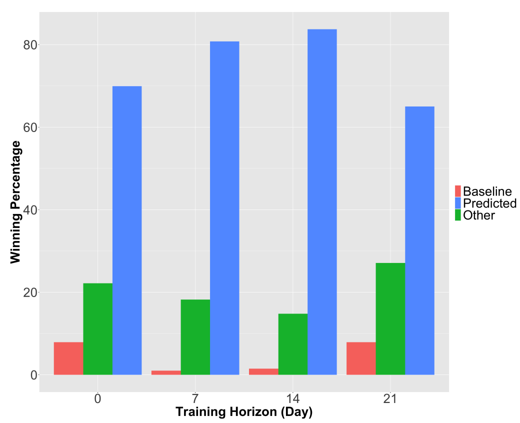

Figure A.16: Three-way comparison of forecasting performance across the predicted cluster trajectory, the alternative cluster trajectory (i.e., the one that was not chosen), and the baseline forecasting method. The heights of the bars indicate the percentage of times each method provided a better forecast of the community response to a particular event for different lengths of observed segments

A.6 Validation of Observed Community Response Time Series

Among the results of this study, the lexicon time series shown in Figure 2.5 constitute community response patterns that generalize well across the disasters in our data. In this section, we validate that these patterns are not artifacts of data generation or processing by comparing them against the corresponding time-series generated from a null model.

We generate a new time series for each dimension of each time series by shuffling the daily observations. We only shuffle observations within the study period once they have been normalized w.r.t. the behavior from the 3 week baseline period). We repeat this procedure 200 times across all disasters to generate as many sets of community responses. Then we use each synthesized set of community responses to generate representative time series for each of the dimensions in Figure 2.5. The result is a distribution of 200 synthetic representative time series for each dimension. In Figure A.17, we replicate Figure 2.5, by using the synthetic distributions. It shows that shuffled observations result in synthetic community responses that, for the most part, do not vary with time and have none of the rich temporal structure observed in the actual data. For all dimensions, the community responses correspond to a fixed intensity which corresponds to the average level in relation to the baseline period.

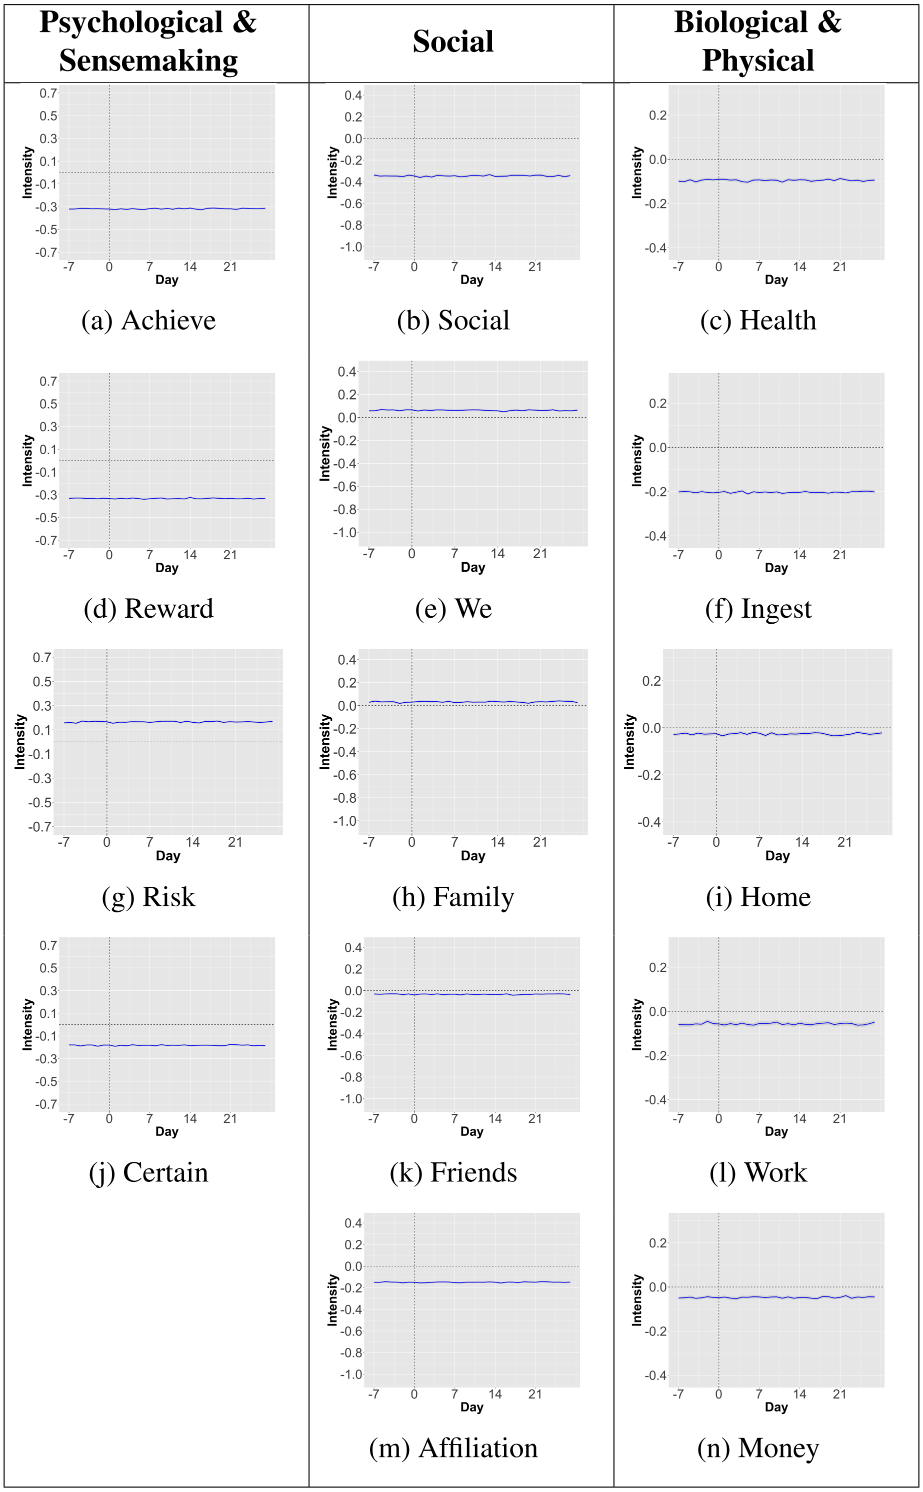

Figure A.17: Community response patterns using shuffled observations for dimensions with universal trajectories across disasters in the original data. All figures in a column have the same intensity (Y) scales which are identical to those of Figure 2.5. The shaded areas are 95% bootstrap confidence intervals. If a confidence interval is not visible for a particular dimension, it is narrower than width of the line showing the mean trajectory.

APPENDIX B

Supplemental Materials for Chapter 4

B.1 Relocation and Return over Different Post-Disaster Periods

	12 Weeks	24 Weeks	36 Weeks
Relocated	2094	645	466
Returned	-	273	155
# Users	16,372	16,667	16,797

Table B.1: Number of people who newly relocated from or returned to their community during different post-disaster periods

B.2 Robustness of Findings to Variations in Relocation Measurement

Figure B.1 summarizes results for variations in the time period of post-disaster observation and the measurement of relocation. They show that, generally, the effects of bonding and bridging capital are consistent across these differences. There are a few differences in the magnitude and statistical significance of the effects observed for some measures. In particular, ego-network density, which did not have a significant effect on relocation likelihood, is associated with an increase in the rank distance of the relocation. While not statistically significant, observations for the effect of density over time for the boolean and distance measurements of relocation follow a similar trajectory to those for rank distance. This observation contradicts our hypothesis that having more bonding capital decreases relocation likelihood and requires further study.

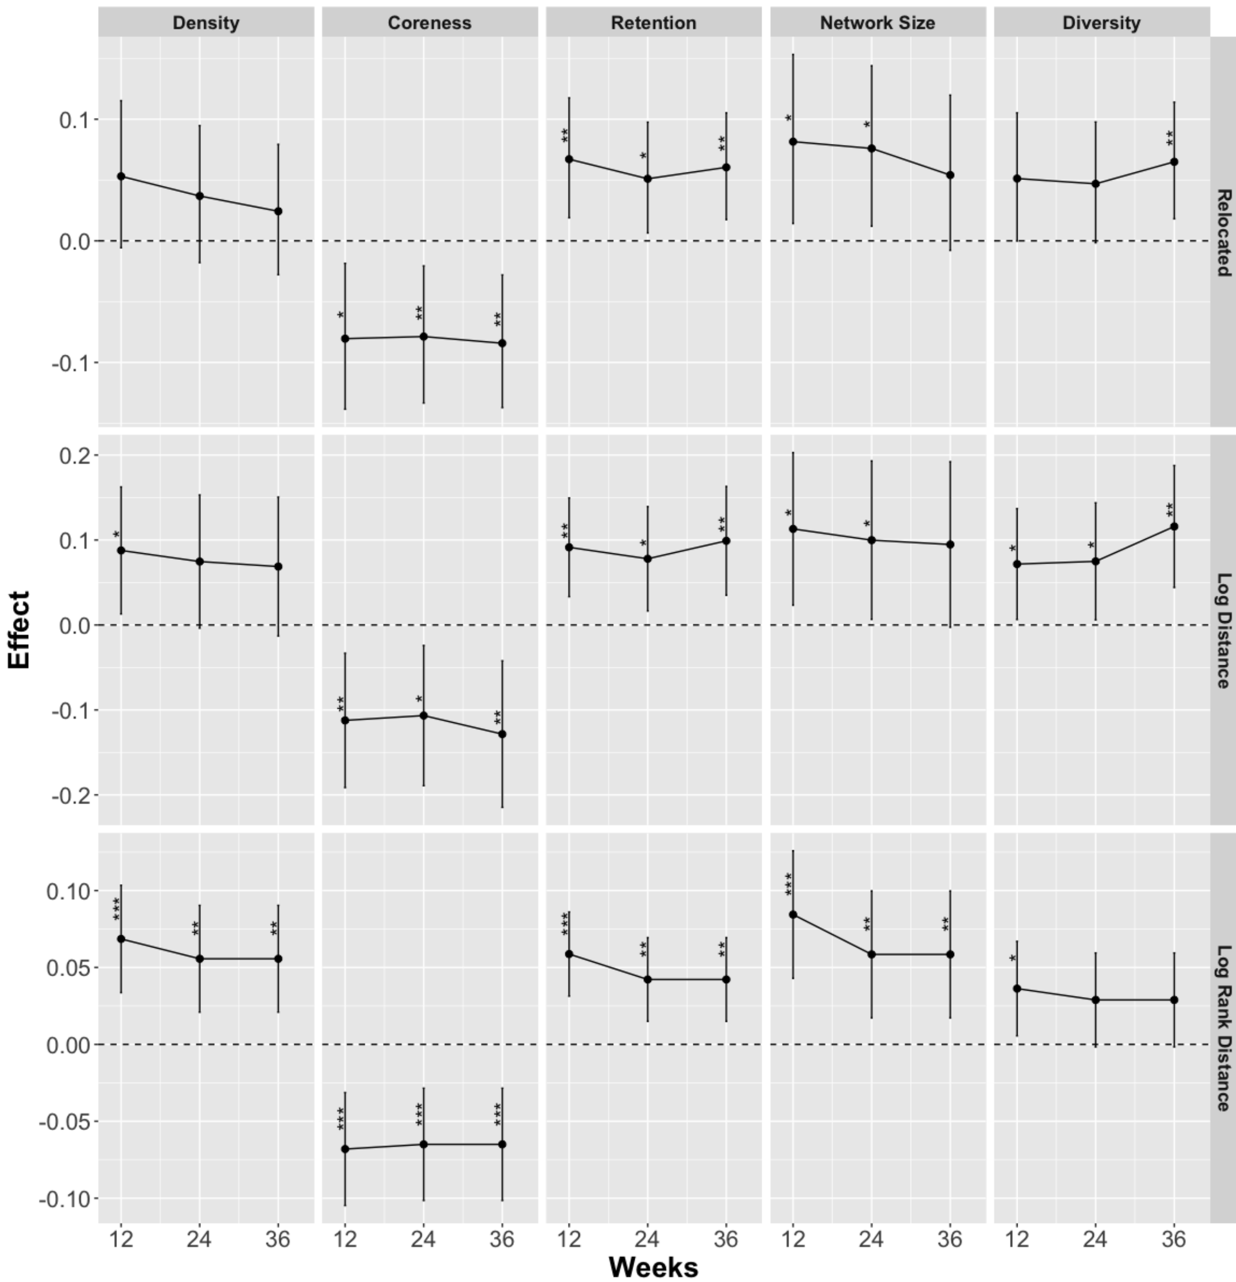

Figure B.1: Effects of social capital on post-disaster relocation across different measurements of relocation (boolean, log distance, and log rank distance) over different measurement periods (12, 24, and 36 weeks). For boolean measurement of relocation (i.e., **Relocated**), the units of the effects is change in odds, while for log distance and log rank distance the units are corresponding magnitudes (i.e., meters (log), change in rank respectively). The error bars show 95% confidence intervals. The stars show significance (*** p-val < 0.001, ** p-val < 0.01, * p-val < 0.05).

BIBLIOGRAPHY

Alberto Abadie. 2005. Semiparametric difference-in-differences estimators. *The Review of Economic Studies* 72, 1 (2005), 1–19.

Francis O Adeola and J Steven Picou. 2014. Social capital and the mental health impacts of Hurricane Katrina: Assessing long-term patterns of psychosocial distress. *International Journal of Mass Emergencies & Disasters* 32, 1 (2014), 121–156.

Emily Aiken, Suzanne Bellue, Dean Karlan, Chris Udry, and Joshua E Blumenstock. 2022. Machine learning and phone data can improve targeting of humanitarian aid. *Nature* 603, 7903 (2022), 864–870.

Christopher A Airriess, Wei Li, Karen J Leong, Angela Chia-Chen Chen, and Verna M Keith. 2008. Church-based social capital, networks and geographical scale: Katrina evacuation, relocation, and recovery in a New Orleans Vietnamese American community. *Geoforum* 39, 3 (2008), 1333–1346.

Daniel P. Aldrich. 2012a. *Building resilience: social capital in post-disaster recovery*. The University of Chicago Press, Chicago.

Daniel P Aldrich. 2012b. Social capital in post disaster recovery: towards a resilient and compassionate East Asian community. *Economic and welfare impacts of disasters in East Asia and policy responses* (2012), 157–178.

Daniel P Aldrich. 2017. The importance of social capital in building community resilience. *Rethinking resilience, adaptation and transformation in a time of change* (2017), 357–364.

Daniel P Aldrich and Michelle A Meyer. 2015. Social capital and community resilience. *American behavioral scientist* 59, 2 (2015), 254–269.

David Alexander. 1993. *Natural disasters*. UCL Press.

T Jennings Anderson, Marina Kogan, Melissa Bica, Leysia Palen, Kenneth M Anderson, Rebecca Morss, Julie Demuth, Heather Lazrus, Olga Wilhelmi, and J Henderson. 2016. Far Far Away in Far Rockaway: Responses to Risks and Impacts during Hurricane Sandy through First-Person Social Media Narratives.. In *ISCRAM*.

Ofer Arazy and Oded Nov. 2010. Determinants of Wikipedia Quality: The Roles of Global and Local Contribution Inequality. In *Proceedings of the 2010 ACM Conference on Computer Supported Cooperative Work (CSCW '10)*. ACM, New York, NY, USA, 233–236.

Valerio Arnaboldi, Marco Conti, Massimiliano La Gala, Andrea Passarella, and Fabio Pezzoni. 2016. Ego network structure in online social networks and its impact on information diffusion. *Computer Communications* 76 (2016), 26–41.

American Psychiatric Association. 2013. *Diagnostic and Statistical Manual of Mental Disorders* (5th ed. ed.). American Psychiatric Association.

James P Bagrow, Dashun Wang, and Albert-Laszlo Barabasi. 2011. Collective response of human populations to large-scale emergencies. *PloS one* 6, 3 (2011), e17680.

Douglas D Baker and John B Cullen. 1993. Administrative reorganization and configurational context: The contingent effects of age, size, and change in size. *Academy of Management Journal* 36, 6 (1993), 1251–1277.

Andrew Begel, Jan Bosch, and Margaret-Anne Storey. 2013. Social networking meets software development: Perspectives from github, msdn, stack exchange, and topcoder. *IEEE Software* 30, 1 (2013), 52–66.

Yochai Benkler. 2006. *The wealth of networks: How social production transforms markets and freedom*. Yale University Press.

Yochai Benkler, Aaron Shaw, and Benjamin Mako Hill. 2015. Peer production: A form of collective intelligence. *Handbook of collective intelligence* 175 (2015).

Gary Blau. 2006. A process model for understanding victim responses to worksite/function closure. *Human Resource Management Review* 16, 1 (2006), 12–28.

Peter M Blau. 1970. A formal theory of differentiation in organizations. *American sociological review* (1970), 201–218.

The GitHub Blog. 2013. Explore what is Trending on GitHub. https://github.blog/2013-08-13-explore-what-is-trending-on-github/. Accessed: 2019-10-05.

Joshua Blumenstock, Gabriel Cadamuro, and Robert On. 2015. Predicting poverty and wealth from mobile phone metadata. *Science* 350, 6264 (2015), 1073–1076.

George A Bonanno. 2004. Loss, trauma, and human resilience: Have we underestimated the human capacity to thrive after extremely aversive events? *American psychologist* 59, 1 (2004), 20.

George A Bonanno, Chris R Brewin, Krzysztof Kaniasty, and Annette M La Greca. 2010. Weighing the costs of disaster: Consequences, risks, and resilience in individuals, families, and communities. *Psychological science in the public interest* 11, 1 (2010), 1–49.

George A Bonanno, Samuel MY Ho, Jane CK Chan, Rosalie SY Kwong, Celia KY Cheung, Claudia PY Wong, and Vivian CW Wong. 2008. Psychological resilience and dysfunction among hospitalized survivors of the SARS epidemic in Hong Kong: a latent class approach. *Health Psychology* 27, 5 (2008), 659.

George A Bonanno, Paul Kennedy, Isaac R Galatzer-Levy, Peter Lude, and Mangus L Elfström. 2012. Trajectories of resilience, depression, and anxiety following spinal cord injury. *Rehabilitation psychology* 57, 3 (2012), 236.

Linda B Bourque, Judith M Siegel, Megumi Kano, and Michele M Wood. 2006. Weathering the storm: The impact of hurricanes on physical and mental health. *The Annals of the American Academy of Political and Social Science* 604, 1 (2006), 129–151.

Peter Brown. 1971. *The World of Late Antiquity AD 150-750*. Harcourt Brace Jovanovich.

Jed Brubaker, Funda Kivran-Swaine, Lee Taber, and Gillian Hayes. 2012. Grief-stricken in a crowd: The language of bereavement and distress in social media. In *Proceedings of the International AAAI Conference on Web and Social Media*, Vol. 6. 42–49.

Moira Burke and Robert Kraut. 2013. Using Facebook after losing a job: Differential benefits of strong and weak ties. In *Proceedings of the 2013 conference on Computer supported cooperative work*. 1419–1430.

Ronald S Burt. 2000. The network structure of social capital. *Research in organizational behavior* 22 (2000), 345–423.

Ronald S Burt. 2007. *Brokerage and closure: An introduction to social capital*. OUP Oxford.

Francesco Calabrese, Giusy Di Lorenzo, Liang Liu, and Carlo Ratti. 2011. Estimating Origin-Destination flows using opportunistically collected mobile phone location data from one million users in Boston Metropolitan Area. (2011).

David Card and Alan B Krueger. 1993. Minimum wages and employment: A case study of the fast food industry in New Jersey and Pennsylvania.

David G Casagrande, Heather McIlvaine-Newsad, and Eric C Jones. 2015. Social networks of help-seeking in different types of disaster responses to the 2008 Mississippi river floods. *Human Organization* 74, 4 (2015), 351–361.

Casey Casalnuovo, Bogdan Vasilescu, Premkumar Devanbu, and Vladimir Filkov. 2015. Developer onboarding in GitHub: the role of prior social links and language experience. In *Proceedings of the 2015 10th joint meeting on foundations of software engineering*. ACM, 817–828.

Pew Research Center. 2023. Social Media Fact Sheet. https://www.pewresearch.org/internet/fact-sheet/social-media/

Emily Chamlee-Wright and Virgil Henry Storr. 2011. Social capital as collective narratives and post-disaster community recovery. *The sociological review* 59, 2 (2011), 266–282.

Katie E Cherry, Matthew R Calamia, Traci Birch, and Aimee Moles. 2021. Assessing Mental Health After a Disaster: Flood Exposure, Recovery Stressors, and Prior Flood Experience. In *The Intersection of Trauma and Disaster Behavioral Health*. Cham: Springer International Publishing, 271–283. https://doi.org/10.1007/978-3-030-51525-6_16

John Child and Alfred Kieser. 1981. Development of organizations over time. *Handbook of organizational design* 1 (1981), 28–64.

Minje Choi, Ceren Budak, Daniel M Romero, and David Jurgens. 2021. More than Meets the Tie: Examining the Role of Interpersonal Relationships in Social Networks. *Proceedings of the International AAAI Conference on Web and Social Media* 15 (2021), 105–116.

Minje Choi, David Jurgens, and Daniel M. Romero. 2023. Analyzing the Engagement of Social Relationships during Life Event Shocks in Social Media. *Proceedings of the International AAAI Conference on Web and Social Media* 17, 1 (Jun. 2023), 149–160. https://doi.org/10.1609/icwsm.v17i1.22134

TINA A Coffelt, FRANCES LM Smith, Michael Sollitto, and April R Payne. 2010. Using sense-making to understand victims' responses to a natural disaster. *Northwest Journal of Communication* 39, 1 (2010), 11–35.

Jacob Cohen. 1988. *Statistical power analysis for the behavioral sciences* (2nd ed. ed.). L. Erlbaum Associates.

Michael A Cohn, Matthias R Mehl, and James W Pennebaker. 2004. Linguistic markers of psychological change surrounding September 11, 2001. *Psychological science* 15, 10 (2004), 687–693.

James S Coleman. 1988. Social capital in the creation of human capital. *American journal of sociology* 94 (1988), S95–S120.

Zhen Cong, Ali Nejat, Daan Liang, Yaolin Pei, and Roxana J Javid. 2018. Individual relocation decisions after tornadoes: a multi-level analysis. *Disasters* 42, 2 (2018), 233–250.

Mackenzie Consoer and Anita Milman. 2016. The dynamic process of social capital during recovery from Tropical Storm Irene in Vermont. *Natural Hazards* 84 (2016), 155–174.

Thomas D Cook, Donald Thomas Campbell, and William Shadish. 2002. *Experimental and quasi-experimental designs for generalized causal inference*. Vol. 1195. Houghton Mifflin Boston, MA.

Leonardo Corbo, Raffaele Corrado, and Simone Ferriani. 2016. A new order of things: Network mechanisms of field evolution in the aftermath of an exogenous shock. *Organization Studies* 37, 3 (2016), 323–348.

Matteo Coronese, Francesco Lamperti, Klaus Keller, Francesca Chiaromonte, and Andrea Roventini. 2019. Evidence for sharp increase in the economic damages of extreme natural disasters. *Proceedings of the National Academy of Sciences* 116, 43 (2019), 21450–21455.

National Research Council. 2006. *Facing Hazards and Disasters: Understanding Human Dimensions*. The National Academies Press, Washington, DC.

Peter Craig, Srinivasa Vittal Katikireddi, Alastair Leyland, and Frank Popham. 2017. Natural experiments: an overview of methods, approaches, and contributions to public health intervention research. *Annual review of public health* 38 (2017), 39–56.

N Crossley, D Bellotti, G Edwards, MG Everett, J Koskinen, and M Tranmer. 2015. Social capital and small worlds: A primer. *Social network analysis for ego-nets. Thousand Oaks, CA: Sage* (2015), 23–43.

Laura Dabbish, Colleen Stuart, Jason Tsay, and Jim Herbsleb. 2012. Social coding in GitHub: transparency and collaboration in an open software repository. In *Proceedings of the ACM 2012 conference on computer supported cooperative work*. ACM, Association for Computing Machinery, 1277–1286.

Clarissa C David, Jonathan Corpus Ong, and Erika Fille T Legara. 2016. Tweeting Supertyphoon Haiyan: Evolving Functions of Twitter during and after a Disaster Event. *PloS one* 11, 3 (2016), e0150190.

Munmun De Choudhury and Scott Counts. 2013. Understanding affect in the workplace via social media. In *Proceedings of the 2013 conference on Computer supported cooperative work*. ACM, Association for Computing Machinery, 303–316.

Munmun De Choudhury and Emre Kiciman. 2017. The language of social support in social media and its effect on suicidal ideation risk. In *Proceedings of the international AAAI conference on web and social media*, Vol. 11. 32–41.

Munmun De Choudhury, Andres Monroy-Hernandez, and Gloria Mark. 2014. Narco emotions: affect and desensitization in social media during the mexican drug war. In *Proceedings of the 32nd annual ACM conference on Human factors in computing systems*. ACM, 3563–3572.

Anais Delilah Roque, David Pijawka, and Amber Wutich. 2020. The role of social capital in resiliency: Disaster recovery in Puerto Rico. *Risk, Hazards & Crisis in Public Policy* 11, 2 (2020), 204–235.

Deborah J DeWolfe. 2000. Training manual for mental health and human service workers in major disasters. (2000).

Fernando Diaz, Michael Gamon, Jake M Hofman, Emre Kıcıman, and David Rothschild. 2016. Online and social media data as an imperfect continuous panel survey. *PloS one* 11, 1 (2016), e0145406.

Sarah Dobie, Jennifer Schneider, and Avery Szafranski. 2019. Going beyond the Waffle House index: Using food systems as an indicator of community health and sustainability. In *2019 IEEE International Symposium on Technologies for Homeland Security (HST)*. IEEE, 1–6.

Bruce Doré, Leonard Ort, Ofir Braverman, and Kevin N Ochsner. 2015. Sadness shifts to anxiety over time and distance from the national tragedy in Newtown, Connecticut. *Psychological science* 26, 4 (2015), 363–373.

Thomas E Drabek. 2018. Community Processes: Coordination. New York, NY: Springer New York, 217–233. https://doi.org/10.1007/978-0-387-32353-4_13

Robin IM Dunbar, Valerio Arnaboldi, Marco Conti, and Andrea Passarella. 2015. The structure of online social networks mirrors those in the offline world. *Social networks* 43 (2015), 39–47.

Jennifer A Dunne, Richard J Williams, and Neo D Martinez. 2002. Network structure and biodiversity loss in food webs: robustness increases with connectance. *Ecology letters* 5, 4 (2002), 558–567.

Thad Dunning. 2012. *Natural experiments in the social sciences: A design-based approach*. Cambridge University Press.

Russell R. Dynes. 1970. Organized behavior in disaster. (1970).

Achim Edelmann, Tom Wolff, Danielle Montagne, and Christopher A Bail. 2020. Computational social science and sociology. *Annual Review of Sociology* 46 (2020), 61–81.

Lesley Edgemon, C Freeman, C Burdi, J Trail, K Marsh, and K Pfeiffer. 2020. Community Resilience Indicator Analysis: County-Level Analysis of Commonly Used Indicators from Peer-Reviewed Research, 2020 Update.

Jared F Edgerton, Skyler J Cranmer, and Victor Finomore. 2022. How teams adapt to exogenous shocks: Experimental evidence with node knockouts of central members. *Network Science* 10, 3 (2022), 261–282.

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.

Stefanie Engel and Ana María Ibáñez. 2007. Displacement due to violence in Colombia: A household-level analysis. *Economic development and cultural change* 55, 2 (2007), 335–365.

AJ Faas and Eric C Jones. 2017. Social network analysis focused on individuals facing hazards and disasters. *Social network analysis of disaster response, recovery, and adaptation* (2017), 11–23.

Hongbo Fang, Hemank Lamba, James Herbsleb, and Bogdan Vasilescu. 2022. " This is damn slick!" estimating the impact of tweets on open source project popularity and new contributors. In *Proceedings of the 44th International Conference on Software Engineering*. 2116–2129.

Charles J Fombrun and Stefan Wally. 1989. Structuring small firms for rapid growth. *Journal of Business Venturing* 4, 2 (1989), 107–122.

David Forbes, Nathan Alkemade, Elizabeth Waters, Lisa Gibbs, Colin Gallagher, Phillipa Pattison, Dean Lusher, Colin MacDougall, Louise Harms, Karen Block, et al. 2015. The role of anger and ongoing stressors in mental health following a natural disaster. *Australian & New Zealand Journal of Psychiatry* 49, 8 (2015), 706–713.

David Garcia, Max Pellert, Jana Lasser, and Hannah Metzler. 2021. Social media emotion macroscopes reflect emotional experiences in society at large. *arXiv preprint arXiv:2107.13236* (2021).

David Garcia and Bernard Rimé. 2019. Collective emotions and social resilience in the digital traces after a terrorist attack. *Psychological science* 30, 4 (2019), 617–628.

David Garcia and Bernard Rimé. 2019. Collective Emotions and Social Resilience in the Digital Traces After a Terrorist Attack. *Psychological Science* 30, 4 (2019), 617–628.

Venkata Rama Kiran Garimella, Ingmar Weber, and Sonya Dal Cin. 2014. From "i love you babe" to "leave me alone"-Romantic relationship breakups on Twitter. In *Social Informatics: 6th International Conference, SocInfo 2014, Barcelona, Spain, November 11-13, 2014. Proceedings 6*. Springer, 199–215.

Rui Gaspar, Cláudia Pedro, Panos Panagiotopoulos, and Beate Seibt. 2016. Beyond positive or negative: Qualitative sentiment analysis of social media reactions to unexpected stressful events. *Computers in Human Behavior* 56 (2016), 179–191.

Franziska Gaupp. 2020. Extreme events in a globalized food system. *One Earth* 2, 6 (2020), 518–521.

Rob Gordon. 2004. The social system as site of disaster impact and resource for recovery. *Australian Journal of Emergency Management, The* 19, 4 (2004), 16–22.

Eva-Maria Gortner and James W Pennebaker. 2003. The archival anatomy of a disaster: Media coverage and community-wide health effects of the Texas A&M bonfire tragedy. *Journal of Social and Clinical Psychology* 22, 5 (2003), 580–603.

Mark S Granovetter. 1973. The strength of weak ties. *American journal of sociology* 78, 6 (1973), 1360–1380.

Robert A Greevy Jr, Carlos G Grijalva, Christianne L Roumie, Cole Beck, Adriana M Hung, Harvey J Murff, Xulei Liu, and Marie R Griffin. 2012. Reweighted Mahalanobis distance matching for cluster-randomized trials with missing data. *Pharmacoepidemiology and drug safety* 21 (2012), 148–154.

Anna Grimm, Lynn Hulse, Marek Preiss, and Silke Schmidt. 2014. Behavioural, emotional, and cognitive responses in European disasters: results of survivor interviews. *Disasters* 38, 1 (2014), 62–83.

Anna Grimm, Lynn Hulse, and Silke Schmidt. 2012. Human responses to disasters: A pilot study on peritraumatic emotional and cognitive processing. *Europe's Journal of Psychology* 8, 1 (2012), 112.

Peter Gronn. 2002. Distributed leadership as a unit of analysis. *The leadership quarterly* 13, 4 (2002), 423–451.

Stefanie Habersang, Jill Küberling-Jost, Markus Reihlen, and Christoph Seckler. 2019. A Process Perspective on Organizational Failure: A Qualitative Meta-Analysis. *Journal of Management Studies* 56, 1 (2019), 19–56.

Kirsten Hagon, Vincent Turmine, Guido Pizzini, Alison Freebairn, and Roop Singh. 2020. Hazards everywhere – climate and disaster trends and impacts. In *World Disasters Report: come heat or high water.* International Federation of Red Cross, Geneva, Chapter 2. OCLC: 1232500224.

Donald C Hambrick and Lynn M Crozier. 1985. Stumblers and stars in the management of rapid growth. *Journal of business venturing* 1, 1 (1985), 31–45.

Benjamin Mako Hill and Aaron Shaw. 2021. The Hidden Costs of Requiring Accounts: Quasi-Experimental Evidence From Peer Production. *Communication Research* 48, 6 (2021), 771–795.

William R Hobbs and Moira K Burke. 2017. Connective recovery in social networks after the death of a friend. *Nature Human Behaviour* 1, 5 (2017), 0092.

Tom Horlick-Jones. 1995. Modern disasters as outrage and betrayal. *International Journal of Mass Emergencies and Disasters* 13, 3 (1995), 305–315.

J Brian Houston, Joshua Hawthorne, Mildred F Perreault, Eun Hae Park, Marlo Goldstein Hode, Michael R Halliwell, Sarah E Turner McGowen, Rachel Davis, Shivani Vaid, Jonathan A McElderry, et al. 2015. Social media and disasters: a functional framework for social media use in disaster planning, response, and research. *Disasters* 39, 1 (2015), 1–22.

Hsin-Yi Hsueh. 2019. The role of household social capital in post-disaster recovery: An empirical study in Japan. *International journal of disaster risk reduction* 39 (2019), 101199.

Kate Hunt and Mike Gruszczynski. 2023. "Horizontal" Two-Step Flow: The Role of Opinion Leaders in Directing Attention to Social Movements in Decentralized Information Environments. *Mass Communication and Society* (2023), 1–24.

Guido W Imbens and Donald B Rubin. 2015. *Causal inference in statistics, social, and biomedical sciences*. Cambridge University Press.

Muhammad Imran, Carlos Castillo, Fernando Diaz, and Sarah Vieweg. 2015. Processing social media messages in mass emergency: A survey. *ACM Computing Surveys (CSUR)* 47, 4 (2015), 67.

Muhammad Imran, Shady Elbassuoni, Carlos Castillo, Fernando Diaz, and Patrick Meier. 2013. Practical extraction of disaster-relevant information from social media. In *Proceedings of the 22nd International Conference on World Wide Web*. ACM, 1021–1024.

Rabiul Islam and Greg Walkerden. 2014. How bonding and bridging networks contribute to disaster resilience and recovery on the Bangladeshi coast. *International journal of disaster risk reduction* 10 (2014), 281–291.

Swami Iyer, Timothy Killingback, Bala Sundaram, and Zhen Wang. 2013. Attack robustness and centrality of complex networks. *PloS one* 8, 4 (2013), e59613.

Jing Jiang, David Lo, Jiahuan He, Xin Xia, Pavneet Singh Kochhar, and Li Zhang. 2017. Why and how developers fork what from whom in GitHub. *Empirical Software Engineering* 22, 1 (2017), 547–578.

Nickolas M Jones, Rebecca R Thompson, Christine Dunkel Schetter, and Roxane Cohen Silver. 2017. Distress and rumor exposure on social media during a campus lockdown. *Proceedings of the National Academy of Sciences* 114, 44 (2017), 11663–11668.

Muneo Kaigo. 2012. Social media usage during disasters and social capital: Twitter and the Great East Japan earthquake. *Keio Communication Review* 34, 1 (2012), 19–35.

Eirini Kalliamvakou, Georgios Gousios, Kelly Blincoe, Leif Singer, Daniel M German, and Daniela Damian. 2014. The promises and perils of mining GitHub. In *Proceedings of the 11th working conference on mining software repositories*. ACM, 92–101.

Nilamadhab Kar and Binaya Kumar Bastia. 2006. Post-traumatic stress disorder, depression and generalised anxiety disorder in adolescents after a natural disaster: a study of comorbidity. *Clinical Practice and Epidemiology in Mental Health* 2, 1 (2006), 1–7.

Ananda Y Karunarathne and Gunhak Lee. 2019. Traditional social capital and socioeconomic networks in response to flood disaster: A case study of rural areas in Sri Lanka. *International Journal of Disaster Risk Reduction* 41 (2019), 101279.

Robert W Kates and David Pijawka. 1977. From rubble to monument: the pace of reconstruction. *Reconstruction following disaster* 1 (1977), 1–23.

Kiyomi Kawamoto and Karl Kim. 2019. Efficiencies of bonding, bridging and linking social capital: Cleaning up after disasters in Japan. *International Journal of disaster risk reduction* 33 (2019), 64–73.

Brian Keegan, Darren Gergle, and Noshir Contractor. 2013. Hot off the wiki: Structures and dynamics of Wikipedia's coverage of breaking news events. *American Behavioral Scientist* 57, 5 (2013), 595–622.

Ann P Kinzig, Paul Ryan, Michel Etienne, Helen Allison, Thomas Elmqvist, and Brian H Walker. 2006. Resilience and regime shifts: assessing cascading effects. *Ecology and society* 11, 1 (2006).

Aniket Kittur and Robert E. Kraut. 2008. Harnessing the Wisdom of Crowds in Wikipedia: Quality Through Coordination. In *Proceedings of the 2008 ACM Conference on Computer Supported Cooperative Work (CSCW '08)*. ACM, New York, NY, USA, 37–46.

Niklas Kiviluoto. 2013. Growth as evidence of firm success: myth or reality? *Entrepreneurship & Regional Development* 25, 7-8 (2013), 569–586.

Marina Kogan and Leysia Palen. 2018. Conversations in the eye of the storm: At-scale features of conversational structure in a high-tempo, high-stakes microblogging environment. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems*. 1–13.

Marina Kogan, Leysia Palen, and Kenneth M Anderson. 2015. Think local, retweet global: Retweeting by the geographically-vulnerable during Hurricane Sandy. In *Proceedings of the 18th ACM conference on computer supported cooperative work & social computing*. ACM, 981–993.

Robert Kohn and Itzhak Levav. 1990. Bereavement in disaster: an overview of the research. *International Journal of Mental Health* 19, 2 (1990), 61–76.

Dimitrios Kotsakos, Goce Trajcevski, Dimitrios Gunopulos, and Charu C Aggarwal. 2013. Time-Series Data Clustering.

Gabriel E Kreindler and Yuhei Miyauchi. 2021. Measuring commuting and economic activity inside cities with cell phone records. *Review of Economics and Statistics* (2021), 1–48.

William H Kruskal and W Allen Wallis. 1952. Use of ranks in one-criterion variance analysis. *Journal of the American statistical Association* 47, 260 (1952), 583–621.

Roeline G Kuijer and Jessica A Boyce. 2012. Emotional eating and its effect on eating behaviour after a natural disaster. *Appetite* 58, 3 (2012), 936–939.

Dean Kyne and Daniel P Aldrich. 2020. Capturing bonding, bridging, and linking social capital through publicly available data. *Risk, Hazards & Crisis in Public Policy* 11, 1 (2020), 61–86.

David Laniado and Riccardo Tasso. 2011. Co-authorship 2.0: Patterns of collaboration in Wikipedia. In *Proceedings of the 22nd ACM conference on Hypertext and hypermedia*. ACM, 201–210.

David Lazer, Alex Pentland, Lada Adamic, Sinan Aral, Albert-László Barabási, Devon Brewer, Nicholas Christakis, Noshir Contractor, James Fowler, Myron Gutmann, et al. 2009. Computational social science. *Science* 323, 5915 (2009), 721–723.

Seungyoon Lee, Laura K Siebeneck, Bailey C Benedict, Takahiro Yabe, Caitlyn M Jarvis, and Satish V Ukkusuri. 2022. Patterns of Social Support and Trajectories of Household Recovery after Superstorm Sandy: Contrasting Influences of Bonding and Bridging Social Capital. *Natural Hazards Review* 23, 2 (2022), 04022002.

Young-Jun Lee, Hiroaki Sugiura, and Ingrida Gečienė. 2017. Stay or relocate: The roles of networks after the Great East Japan earthquake. In *Social network analysis of disaster response, recovery, and adaptation*. Elsevier, 223–238.

Antonio Lima, Luca Rossi, and Mirco Musolesi. 2014. Coding together at scale: GitHub as a collaborative social network. In *Eighth International AAAI Conference on Weblogs and Social Media*.

Jie Lin and Robert G Cromley. 2018. Inferring the home locations of Twitter users based on the spatiotemporal clustering of Twitter data. *Transactions in GIS* 22, 1 (2018), 82–97.

Yu-Ru Lin and Drew Margolin. 2014. The ripple of fear, sympathy and solidarity during the Boston bombings. *EPJ Data Science* 3, 1 (2014), 31.

Yu-Ru Lin, Drew Margolin, and Xidao Wen. 2017. Tracking and analyzing individual distress following terrorist attacks using social media streams. *Risk analysis* 37, 8 (2017), 1580–1605.

Michael K Lindell. 2013. Recovery and Reconstruction After Disaster. Dordrecht: Springer Netherlands, 812–824. https://doi.org/10.1007/978-1-4020-4399-4_285

Ronald W Lindell, Michael K; Perry. 2003. *Communicating Environmental Risk in Multiethnic Communities*. Vol. 7. Thousand Oaks: SAGE Publications. https://doi.org/10.4135/9781452229188

Stephan Liozu, Andreas Hinterhuber, and Toni Somers. 2014. Organizational design and pricing capabilities for superior firm performance. *Management Decision* 52, 1 (2014), 54–78.

Ano Lobb, Nancy Mock, and Paul L Hutchinson. 2012. Traditional and social media coverage and charitable giving following the 2010 earthquake in Haiti. *Prehospital and disaster medicine* 27, 04 (2012), 319–324.

Xuan Lu, Wei Ai, Yixin Wang, and Qiaozhu Mei. 2023. Team Resilience under Shock: An Empirical Analysis of GitHub Repositories during Early COVID-19 Pandemic. *arXiv preprint arXiv:2301.12326* (2023).

Kristen Magis. 2010. Community resilience: An indicator of social sustainability. *Society and natural resources* 23, 5 (2010), 401–416.

Momin Malik and Jürgen Pfeffer. 2016. Identifying platform effects in social media data. In *Proceedings of the International AAAI Conference on Web and Social Media*, Vol. 10. 241–249.

Adriana M Manago, Tamara Taylor, and Patricia M Greenfield. 2012. Me and my 400 friends: the anatomy of college students' Facebook networks, their communication patterns, and well-being. *Developmental psychology* 48, 2 (2012), 369.

Jennifer Marlow, Laura Dabbish, and Jim Herbsleb. 2013. Impression formation in online peer production: activity traces and personal profiles in github. In *Proceedings of the 2013 conference on Computer supported cooperative work*. ACM, 117–128.

Abraham H. Maslow. 1954. Motivation and personality. (1954).

Nora McDonald, Kelly Blincoe, EVA Petakovic, and Sean Goggins. 2014. Modeling distributed collaboration on Github. *Advances in Complex Systems* 17, 07n08 (2014), 1450024.

Danaë Metaxa-Kakavouli, Paige Maas, and Daniel P Aldrich. 2018. How social ties influence hurricane evacuation behavior. *Proceedings of the ACM on human-computer interaction* 2, CSCW (2018), 1–16.

Alan D Meyer. 1982. Adapting to environmental jolts. *Administrative science quarterly* (1982), 515–537.

Alan D Meyer, Geoffrey R Brooks, and James B Goes. 1990. Environmental jolts and industry revolutions: Organizational responses to discontinuous change. *Strategic Management Journal* (1990), 93–110.

Alan D Meyer, Vibha Gaba, and Kenneth A Colwell. 2005. Organizing far from equilibrium: Nonlinear change in organizational fields. *Organization science* 16, 5 (2005), 456–473.

Jason Miklian and Kristian Hoelscher. 2022. SMEs and exogenous shocks: A conceptual literature review and forward research agenda. *International Small Business Journal* 40, 2 (2022), 178–204.

Danny Miller. 1994. What happens after success: The perils of excellence. *Journal of Management Studies* 31, 3 (1994), 325–358.

Sanchayeeta Misra, Rupak Goswami, Tandra Mondal, and Rabindranath Jana. 2017. Social networks in the context of community response to disaster: Study of a cyclone-affected community in Coastal West Bengal, India. *International journal of disaster risk reduction* 22 (2017), 281–296.

Saif M Mohammad and Peter D Turney. 2013. Crowdsourcing a word–emotion association lexicon. *Computational Intelligence* 29, 3 (2013), 436–465.

Joshua C Morganstein, James C West, Lester A Huff, Brian W Flynn, Carol S Fullerton, David M Benedek, and Robert J Ursano. 2016. Psychosocial responses to disaster and exposures: Distress reactions, health risk behavior, and mental disorders. *Mental health and social issues following a nuclear accident: The case of Fukushima* (2016), 99–117.

Lydia Morris and Sarah Irwin. 1992. Unemployment and informal support: dependency, exclusion, or participation? *Work, Employment and Society* 6, 2 (1992), 185–207.

Brenda L Murphy. 2007. Locating social capital in resilient community-level emergency management. *Natural Hazards* 41 (2007), 297–315.

Dhiraj Murthy and Scott A Longwell. 2013. Twitter and disasters: The uses of Twitter during the 2010 Pakistan floods. *Information, Communication & Society* 16, 6 (2013), 837–855.

Diane Garaventa Myers. 1994. *Disaster response and recovery: A handbook for mental health professionals*. DIANE Publishing.

Seth A Myers and Jure Leskovec. 2014. The bursty dynamics of the twitter information network. In *Proceedings of the 23rd international conference on World wide web*. 913–924.

Akira Namatame and Takanori Komatsu. 2011. Management of systemic risks and cascade failures in a networked society. *Inf. Knowl. Syst. Manag.* 10, 1-4 (2011), 111–133.

NOAA NCEI. 2020. NOAA National Centers for Environmental Information (NCEI) US billion-dollar weather and climate disasters.

Mark EJ Newman. 2001. Scientific collaboration networks. II. Shortest paths, weighted networks, and centrality. *Physical review E* 64, 1 (2001), 016132.

Charlene L Nicholls-Nixon. 2005. Rapid growth and high performance: The entrepreneur's "impossible dream?". *Academy of Management Perspectives* 19, 1 (2005), 77–89.

Friedrich Nietzsche and Oscar Levy. 1909. *The twilight of the idols*. Vol. 17. Allen & Unwin.

Fran H Norris, Matthew J Friedman, Patricia J Watson, Christopher M Byrne, Eolia Diaz, and Krzysztof Kaniasty. 2002. 60,000 disaster victims speak: Part I. An empirical review of the empirical literature, 1981–2001. *Psychiatry: Interpersonal and biological processes* 65, 3 (2002), 207–239.

Fran H. Norris, James F. Phifer, and Krzysztof Kaniasty. 1994. Individual and Community Reactions to the Kentucky Floods: Findings from a Longitudinal Study of Older Adults. (1994), 378–400. https://doi.org/10.1017/CBO9780511570162.019

Fran H Norris, Susan P Stevens, Betty Pfefferbaum, Karen F Wyche, and Rose L Pfefferbaum. 2008. Community resilience as a metaphor, theory, set of capacities, and strategy for disaster readiness. *American journal of community psychology* 41, 1 (2008), 127–150.

Hüseyin Oktay, Brian J Taylor, and David D Jensen. 2010. Causal discovery in social media using quasi-experimental designs. In *Proceedings of the first workshop on social media analytics*. 1–9.

Alexandra Olteanu, Carlos Castillo, Fernando Diaz, and Sarah Vieweg. 2014. CrisisLex: A Lexicon for Collecting and Filtering Microblogged Communications in Crises.. In *ICWSM*.

Alexandra Olteanu, Sarah Vieweg, and Carlos Castillo. 2015. What to Expect When the Unexpected Happens: Social Media Communications Across Crises. In *Proceedings of the 18th ACM Conference on Computer Supported Cooperative Work & Social Computing (CSCW '15)*. Association for Computing Machinery, New York, NY, USA, 994–1009. https://doi.org/10.1145/2675133.2675242

Felipe Ortega, Jesus M Gonzalez-Barahona, and Gregorio Robles. 2008. On the inequality of contributions to Wikipedia. In *Proceedings of the 41st Annual Hawaii International Conference on System Sciences (HICSS 2008)*. IEEE, 304–304.

Courtney Page-Tan. 2021. Bonding, bridging, and linking social capital and social media use: How hyperlocal social media platforms serve as a conduit to access and activate bridging and linking ties in a time of crisis. *Natural Hazards* 105, 2 (2021), 2219–2240.

Róbert Pálovics, Primož Dolenc, and Jure Leskovec. 2021. Companies under stress: the impact of shocks on the production network. *EPJ Data Science* 10, 1 (2021), 57.

Sarita Panday, Simon Rushton, Jiban Karki, Julie Balen, and Amy Barnes. 2021. The role of social capital in disaster resilience in remote communities after the 2015 Nepal earthquake. *International Journal of Disaster Risk Reduction* 55 (2021), 102112.

Max Pellert, Hannah Metzler, Michael Matzenberger, and David Garcia. 2022. Validating daily social media macroscopes of emotions. *Scientific Reports* 12, 1 (2022), 11236.

Ronald W Perry. 2018. Defining disaster: An evolving concept. In *Handbook of disaster research*. Springer International Publishing, Cham, 3–22.

François Petitjean, Alain Ketterlin, and Pierre Gançarski. 2011. A global averaging method for dynamic time warping, with applications to clustering. *Pattern Recognition* 44, 3 (2011), 678–693.

Betty Pfefferbaum, Richard L Van Horn, and Rose L Pfefferbaum. 2017. A conceptual framework to enhance community resilience using social capital. *Clinical Social Work Journal* 45, 2 (2017), 102–110.

Tuan Q Phan and Edoardo M Airoldi. 2015. A natural experiment of social network formation and dynamics. *Proceedings of the National Academy of Sciences* 112, 21 (2015), 6595–6600.

María Piquer-Rodríguez, Cecilie Friis, R Ntsiva N Andriatsitohaina, Sébastien Boillat, Paula Roig-Boixeda, Chiara Cortinovis, Davide Geneletti, Maria-Jose Ibarrola-Rivas, Lisa C Kelley, Jorge C Llopis, et al. 2023. Global shocks, cascading disruptions, and (re-) connections: viewing the COVID-19 pandemic as concurrent natural experiments to understand land system dynamics. *Landscape ecology* (2023), 1–15.

E Platt and DM Romero. 2018. Network Structure, Efficiency, and Performance in WikiProjects. In *Proceedings of International AAAI Conference on Web and Social Media (ICWSM)*.

Robert D Putnam. 2000. *Bowling alone: The collapse and revival of American community*. Simon and schuster.

Barry G Rabe. 2004. *Statehouse and greenhouse: The emerging politics of American climate change policy*. Brookings Institution Press.

Beverley Raphael. 1986. When disaster strikes: how individuals and communities cope with catastrophe. (1986).

Eric Raymond. 1999. The cathedral and the bazaar. *Knowledge, Technology & Policy* 12, 3 (1999), 23–49.

Dahlia K. Remler and Van Ryzin Gregg G. 2014. *Natural and Quasi Experiments* (2nd ed.). SAGE Publishing, 427–464.

Brian K Richardson, Laura K Siebeneck, Sara Shaunfield, and Elizabeth Kaszynski. 2014. From" No Man's Land" to a "Stronger Community": Communitas as a Theoretical Framework for Successful Disaster Recovery. *International Journal of Mass Emergencies & Disasters* 32, 1 (2014).

Hannah Ritchie, Pablo Rosado, and Max Roser. 2022. Natural Disasters. *Our World in Data* (2022). https://ourworldindata.org/natural-disasters.

Daniel M Romero, Dan Huttenlocher, and Jon Kleinberg. 2015. Coordination and efficiency in decentralized collaboration. In *Ninth International AAAI Conference on Web and Social Media*.

Daniel M Romero, Brian Uzzi, and Jon Kleinberg. 2016. Social networks under stress. In *Proceedings of the 25th International Conference on World Wide Web*. 9–20.

Paul R Rosenbaum and Donald B Rubin. 1985. Constructing a control group using multivariate matched sampling methods that incorporate the propensity score. *The American Statistician* 39, 1 (1985), 33–38.

Michael Rutter. 2012. Resilience as a dynamic concept. *Development and Psychopathology* 24, 2 (2012), 335–344. https://doi.org/10.1017/S0954579412000028

Fabio Sabatini. 2009. Social capital as social networks: A new framework for measurement and an empirical analysis of its determinants and consequences. *The Journal of Socio-Economics* 38, 3 (2009), 429–442.

Koustuv Saha and Munmun De Choudhury. 2017. Modeling stress with social media around incidents of gun violence on college campuses. *Proceedings of the ACM on Human-Computer Interaction* 1, CSCW (2017), 92.

Koustuv Saha, Benjamin Sugar, John Torous, Bruno Abrahao, Emre Kıcıman, and Munmun De Choudhury. 2019. A social media study on the effects of psychiatric medication use. In *Proceedings of the International AAAI Conference on Web and Social Media*, Vol. 13. 440–451.

Koustuv Saha, John Torous, Eric D Caine, and Munmun De Choudhury. 2020. Psychosocial effects of the COVID-19 pandemic: large-scale quasi-experimental study on social media. *Journal of medical internet research* 22, 11 (2020), e22600.

Haji Mohammad Saleem, Yishi Xu, and Derek Ruths. 2014. Effects of disaster characteristics on Twitter event signature. *Procedia engineering* 78 (2014), 165–172.

Ellyn Satter. 2007. Hierarchy of food needs. *Journal of nutrition education and behavior* 39, 5 (2007), S187–S188.

Melissa A Schilling. 2015. Technology shocks, technological collaboration, and innovation outcomes. *Organization Science* 26, 3 (2015), 668–686.

Fiona Shirani and Karen Henwood. 2011. Taking one day at a time: Temporal experiences in the context of unexpected life course transitions. *Time & Society* 20, 1 (2011), 49–68.

Irina Shklovski, Leysia Palen, and Jeannette Sutton. 2008. Finding community through information and communication technology in disaster response. In *Proceedings of the 2008 ACM conference on Computer supported cooperative work*. ACM, 127–136.

Luke Sloan and Jeffrey Morgan. 2015. Who tweets with their location? Understanding the relationship between demographic characteristics and the use of geoservices and geotagging on Twitter. *PloS one* 10, 11 (2015), e0142209.

Chris M Smith. 2020. Exogenous shocks, the criminal elite, and increasing gender inequality in chicago organized crime. *American Sociological Review* 85, 5 (2020), 895–923.

Kate Starbird and Leysia Palen. 2010. Pass it on?: Retweeting in mass emergency. In *7th International Conference on Information Systems for Crisis Response and Management: Defining Crisis Management 3.0, Proceedings*. ISCRAM.

Jeannette Sutton, C Ben Gibson, Nolan Edward Phillips, Emma S Spiro, Cedar League, Britta Johnson, Sean M Fitzhugh, and Carter T Butts. 2015. A cross-hazard analysis of terse message retransmission on Twitter. *Proceedings of the National Academy of Sciences* 112, 48 (2015), 14793–14798.

Ali Sydney, Caterina Scoglio, Mina Youssef, and Phillip Schumm. 2010. Characterising the robustness of complex networks. *International Journal of Internet Technology and Secured Transactions* 2, 3-4 (2010), 291–320.

Nassim Nicholas Taleb. 2007. *The black swan: The impact of the highly improbable*. Vol. 2. Random house.

TechCrunch. 2013. GitHub Adds Trending Page To Filter By Project, Programming Languages And Developers. http://social.techcrunch.com/2013/08/13/github-adds-trending-page-to-filter-by-project-programming-languages-and- Accessed: 2019-10-05.

Kathleen J Tierney, Michael K Lindell, and Ronald W Perry (Eds.). 2001. *Facing the unexpected: Disaster preparedness and response in the United States*. The National Academies Press, Washington, DC.

Ivan Townshend, Olu Awosoga, Judith Kulig, and HaiYan Fan. 2015. Social cohesion and resilience across communities that have experienced a disaster. *Natural Hazards* 76, 2 (2015), 913–938.

Stojan Trajanovski, Javier Martín-Hernández, Wynand Winterbach, and Piet Van Mieghem. 2013. Robustness envelopes of networks. *Journal of Complex Networks* 1, 1 (2013), 44–62.

Benjamin D Trump, Kelsey Poinsatte-Jones, Meir Elran, Craig Allen, Bojan Srdjevic, Myriam Merad, Dejan M Vasovic, and José Manuel Palma-Oliveira. 2017. Social resilience and critical infrastructure systems. In *Resilience and Risk: Methods and Application in Environment, Cyber and Social Domains*. Springer, 289–299.

Johan Ugander, Lars Backstrom, Cameron Marlow, and Jon Kleinberg. 2012. Structural diversity in social contagion. *Proceedings of the national academy of sciences* 109, 16 (2012), 5962–5966.

Robert J Ursano, Carol S Fullerton, Lars Weisaeth, and Beverley Raphael. 2007. Individual and community responses to disasters. *Textbook of disaster psychiatry* (2007), 3–26.

Bellis van den Berg, Linda Grievink, Joris Yzermans, and Erik Lebret. 2005. Medically unexplained physical symptoms in the aftermath of disasters. *Epidemiologic Reviews* 27, 1 (2005), 92–106.

BN Vasilescu. 2014. Social aspects of collaboration in online software communities. (2014).

Bogdan Vasilescu, Kelly Blincoe, Qi Xuan, Casey Casalnuovo, Daniela Damian, Premkumar Devanbu, and Vladimir Filkov. 2016. The sky is not the limit: multitasking across GitHub projects. In *2016 IEEE/ACM 38th International Conference on Software Engineering (ICSE)*. IEEE, 994–1005.

Bogdan Vasilescu, Vladimir Filkov, and Alexander Serebrenik. 2015a. Perceptions of diversity on GitHub: A user survey. In *Proceedings of the Eighth International Workshop on Cooperative and Human Aspects of Software Engineering*. IEEE Press, 50–56.

Bogdan Vasilescu, Daryl Posnett, Baishakhi Ray, Mark GJ van den Brand, Alexander Serebrenik, Premkumar Devanbu, and Vladimir Filkov. 2015b. Gender and tenure diversity in GitHub teams. In *Proceedings of the 33rd annual ACM conference on human factors in computing systems*. ACM, 3789–3798.

Bogdan Vasilescu, Yue Yu, Huaimin Wang, Premkumar Devanbu, and Vladimir Filkov. 2015c. Quality and productivity outcomes relating to continuous integration in GitHub. In *Proceedings of the 2015 10th Joint Meeting on Foundations of Software Engineering*. ACM, 805–816.

Fiona Verity and Gwyn Jolley. 2008. Closure of an automotive plant: transformation of a work-based 'community'. *Policy Studies* 29, 3 (2008), 331–341.

Mason Walker and Katerina Eva Matsa. 2021. News consumption across social media in 2021. (2021).

Xidao Wen and Yu-Ru Lin. 2016. Sensing Distress Following a Terrorist Event. In *International Conference on Social Computing, Behavioral-Cultural Modeling and Prediction and Behavior Representation in Modeling and Simulation*. Springer, 377–388.

Luke Wertis, Jennifer D Runkle, Margaret M Sugg, and Devyani Singh. 2023. Examining Hurricane Ida's impact on mental health: results from a quasi-experimental analysis. *GeoHealth* 7, 2 (2023), e2022GH000707.

David A Whetten. 1987. Organizational growth and decline processes. *Annual review of sociology* 13, 1 (1987), 335–358.

Thomas Wicke and Roxane Cohen Silver. 2009. A community responds to collective trauma: An ecological analysis of the James Byrd murder in Jasper, Texas. *American Journal of Community Psychology* 44 (2009), 233–248.

Dai Williams. 1999. Human responses to change. *Futures* 31, 6 (1999), 609–616.

Trenton A Williams, Daniel A Gruber, Kathleen M Sutcliffe, Dean A Shepherd, and Eric Yanfei Zhao. 2017. Organizational response to adversity: Fusing crisis management and resilience research streams. *Academy of Management Annals* 11, 2 (2017), 733–769.

Donald Worster. 2004. *Dust bowl: the southern plains in the 1930s*. Oxford University Press.

Takahiro Yabe, P Suresh C Rao, Satish V Ukkusuri, and Susan L Cutter. 2022. Toward data-driven, dynamical complex systems approaches to disaster resilience. *Proceedings of the National Academy of Sciences* 119, 8 (2022), e2111997119.

William Yule, Derek Bolton, Orlee Udwin, Stephanie Boyle, Dominic O'Ryan, and Julie Nurrish. 2000. The long-term psychological effects of a disaster experienced in adolescence: I: The incidence and course of PTSD. *Journal of child psychology and psychiatry* 41, 4 (2000), 503–511.

Emilio Zagheni, Venkata Rama Kiran Garimella, Ingmar Weber, and Bogdan State. 2014. Inferring international and internal migration patterns from twitter data. In *Proceedings of the 23rd international conference on world wide web*. 439–444.

Ark Fangzhou Zhang, Danielle Livneh, Ceren Budak, Lionel Robert, and Daniel Romero. 2017a. Shocking the crowd: The effect of censorship shocks on chinese wikipedia. In *Eleventh International AAAI Conference on Web and Social Media*.

Ark Fangzhou Zhang, Danielle Livneh, Ceren Budak, Lionel P. Robert, Jr., and Daniel M. Romero. 2017b. Crowd Development: The Interplay Between Crowd Evaluation and Collaborative Dynamics in Wikipedia. *Proc. ACM Hum.-Comput. Interact.* 1, CSCW, Article 119 (Dec. 2017), 21 pages.

Ark Fangzhou Zhang, Ruihan Wang, Eric Blohm, Ceren Budak, Lionel P Robert Jr, and Daniel M Romero. 2019. Participation of New Editors after Times of Shock on Wikipedia. In *Proceedings of the International AAAI Conference on Web and Social Media*, Vol. 13. 560–571.

Philip Ziegler. 1970. The black death. *The Medical journal of Australia* 2 16 (1970), 782–3.

Matteo Zignani, Sabrina Gaito, Gian Paolo Rossi, Xiaohan Zhao, Haitao Zheng, and Ben Zhao. 2014. Link and triadic closure delay: Temporal metrics for social network dynamics. In *Proceedings of the International AAAI Conference on Web and Social Media*, Vol. 8. 564–573.