DISSERTATION

DEVELOPMENT OF DATA INTEGRATION STRATEGIES TO IMPROVE INTERDISCIPLINARITY IN HAZARDS RESEARCH

Submitted by:

Blythe Johnston

Department of Civil and Environmental Engineering

In partial fulfillment of the requirements

for the Degree of Doctor of Philosophy

Colorado State University

Fort Collins, Colorado

Spring 2025

Doctoral Committee:

Advisor: John W. van de Lindt

Yanlin Guo
Hussam Mahmoud
Martin Shields

ABSTRACT

DEVELOPMENT OF DATA INTEGRATION STRATEGIES TO IMPROVE INTERDISCIPLINARITY IN HAZARDS RESEARCH

Natural hazards are inherently interdisciplinary problems that pose risk to human life, property and prosperity. To provide holistic and actionable solutions in the face of these hazards, a more integrated approach to hazards research is needed. The current state of the hazards and disaster research continues to work more in disciplinary silos with progress being made around the world. This progress is ongoing, and this dissertation contributes to the investigation of these cross- and trans-disciplinary spaces in the context of natural hazards research and how they can be further fused and progressed, with a specific focus on data integration and modeling techniques that inform the complex problem of outmigration characterization following a hazard event. For contextualization, this dissertation first presents prior attempts at data integration. With the commonly echoed best practice of data integration from the earliest stages of data creation, a set of tools are developed for more interdisciplinary data collection in geographically large field studies. These tools are then implemented for the creation of a multi-community dataset tracking damage and recovery following the December 2021 Midwest Tornado Outbreak. This data can be utilized in the training and parameterization of long-term post-event models such as outmigration prediction. Modeling techniques for using the knowledge and data acquired in this field study are explored to arrive at actionable and predictive data for enhanced interdisciplinary hazards research. These modeling techniques include the combination of top-down and bottom-up approaches, linear multi-regression modeling, agent-based modeling, and hindcasting. Some or all these techniques

are used to first develop a sheltering model to determine the viability of community tornado shelters during an event similar to that seen in the field study, and establish the knowledge needed to undertake the more complex outmigration model. The datasets and modeling techniques created and acquired are then leveraged to develop a top-down and bottom-up outmigration model after a hazard event that predicts rate of gross outmigration, gross inmigration, net migration, and demographic change following a hazard event. With this set of tools and resources, this dissertation aims to tangibly propel the task of interdisciplinarity in disaster and natural hazards research with the set of tools and resources provided here culminating in the development of a model for predicting long-term population flow following an event.

ACKNOWLEDGEMENTS

I would first and foremost like to thank my advisor Professor John W. van de Lindt for supporting the creation of the work presented herein. His guidance, knowledge, and perspective has been invaluable in the pursuit of my PhD. Furthermore, his willingness to let me explore those topics which are of greatest interest to me has helped in my search for a career path that I am excited to embark upon in life's next chapter. I always felt his unwavering trust in my ability to undertake and complete the work at hand even when I was less sure of my ability to do so. And lastly, I would like to thank him for the opportunity to do all this work at Colorado State University so that I could be fortified by the support of my family nearby and form many cherished memories with them over the past few years. For all of this, I am extremely grateful.

I would also like to thank my committee members, Professor Hussam Mahmoud, Professor Yanlin Guo, and Professor Martin Shields, for their willingness to take time out of incredibly busy schedules to help strengthen this dissertation. Thank you also to all other professors from whom I learned during my time at Colorado State University as well as The University of Alabama. I would especially like to thank Professor Thang Dao at The University of Alabama for the introduction to natural hazards research during my undergraduate studies and for the serendipitous introduction he provided to Professor van de Lindt when I was hoping to begin my graduate studies.

So many other mentors have shaped my career and who I am personally, that I cannot possibly list them all here. However, I would be remiss to not acknowledge a few of the groups who provided me with opportunities and mentorship that I cannot begin to understand how I deserved. I will always be grateful to my colleagues at AT&T and especially my boss, Jess Filante, whose belief in me and relentless fight for equitable resilience was and remains such a convicting testament to what we can accomplish when we simply never stop trying. Through my work with AT&T, I also was afforded the incredible opportunity to work with earnest, brilliant, and gracious mentors at FEMA and Argonne National Laboratory from whom I learned so much and to whom I owe even more. As for the CoE, I am so grateful for Connie Hale's support and humor as well as Katie Skakel's empathy and excitement. Lastly on this point of mentors, I am beyond humbled and grateful to the communities that allowed me to work with them on resilience planning work, which resides outside of the scope of this dissertation but was monumental in shaping my professional interests and career goals during my PhD.

I would also like to thank my incredible colleagues in Glover 212/218 and the Center for Risk-Based Community Resilience Planning. This incredible network of people has been such a bolster in what has been a thoroughly challenging yet incredibly fruitful portion of my life. For their support both professionally and personally, I will be forever grateful. I am also so indebted to the incredible undergraduate students with whom I have worked as part of the Midwest Field Study. Their steadfast commitment to this investigation and simultaneous pursuit of joy and friendship has been such a constant reminder for me that happiness can be found all along our journey and not just at finish lines.

And most whole-heartedly, I would like to thank my family. My PhD would not have been possible without my father's relentless enthusiasm, my mother's boundless heart, my brother's steadfast dependability, my sister's unwavering belief, and my fiancé's profound grace. The conviction they taught me has been the most essential characteristic for completing this dissertation, and I will forever be indebted to them for far more than what can be expressed here.

I would also like to thank the Colorado State Graduate School for the opportunity to study here, and the funding sources that made this research possible. This work was supported by the National

Science Foundation Graduate Research Fellowship under Grant No. 006784-00002, as well as The Center for Risk-Based Community Resilience Planning funded by the National Institute of Standards and Technology Science (NIST) under Grant No. 70NANB20H008. The work presented herein reflects the views and opinions of the author and not necessarily the NSF or NIST.

DEDICATION

To my family

TABLE OF CONTENTS

LIST OF TABLES

LIST OF FIGURES

INTRODUCTION

Motivations

As the predictive power of modeling grows, the interest in developing comprehensive and representative models grows as well. The completeness of models now largely relies on the degree to which interdisciplinarity can be achieved from project conception through completion. The primary drivers of interdisciplinary research include the investigation of overlaps and interaction between disciplines and incorporation of emerging technologies to fully understand the complexity of social and natural systems, especially as they impact overarching societal problems (National Academy of Sciences et al., 2004). Within the context of these drivers, fostering community resilience in disaster research necessitates an interdisciplinary approach due to several characteristic features: "many hazards and built and social systems are involved; many types of interacting impacts are involved; many interacting strategies exist to improve resilience; the problem is uncertain, spatial, and dynamic; many diverse stakeholders are involved; and effective communication about risk and resilience is difficult" (Davidson, 2015). Furthermore, Peek et al. (2020) suggest that interdisciplinarity is a requisite element of convergence research, however convergence research extends beyond interdisciplinary research in its explicit pursuit of solutions to identified problems. Indeed, disaster research can advance the implementation of convergence research, and by extension interdisciplinary practice, because it is so readily positioned at the intersection of multiple disciplines and has a self-evident problem to address (Peek et al., 2020). Thus, disaster research can serve as a platform for new integration strategies that can be applied to, or at least inform, a number of other research areas.

However, this path is not without its challenges. Previous efforts at interdisciplinary research from a variety of fields have noted several barriers to effective integration, most of which are universally felt across many disciplinary overlaps. As a grounding for this dissertation, these barriers will be discussed and the frequently encouraged solutions to alleviating these barriers will be investigated. While many of these suggestions seem to adequately address the challenges at hand, data integration remains an area for which there are several avenues forward, but none have been sufficiently applied to the case of disaster research yet. Thus, possible next steps for scalable and community-focused disaster research data integration are explored in this dissertation.

Barriers to Interdisciplinary Interaction

In working on interdisciplinary projects, past research has noted several barriers to complete integration. These barriers include differences in language. Bracken and Oughton (2006) categorized distinct facets of disciplinary language. The first challenge is that of dialect, or the variations in meaning of a word across disciplines as well as between specialized and everyday use. Secondly, disciplines often use contextually reduced or explanatory metaphors that are unique to that field and that perpetuate discipline-specific frames of thought. Solutions to these dialectical differences across disciplines are investigated by ontologists within the field of data science as well, though they typically use "semantics" in place of "dialects." Within this subfield, ontologists are interested in linking data from various sources. To do this they must first determine the meaning of the words in a given context to determine where overlaps exist and to understand the network of ideas in order to avoid the creation of inconsistencies when the information is integrated. The implications and potential of semantic technologies will be discussed further in the discussion on data integration due to the inherent link between data and the meaning assigned to it (Shvaiko & Euzenat, 2013). Also, differences in language pose a clerical issue. Research is difficult to find when different terminology is used across disciplines yet for the same topic (Sherman-Morris et al., 2021). These differences in terminology allude to the epistemological divides that have proven difficult to overcome.

Another key barrier is that of institutionally reinforced disciplinary boundaries. Because institutions are developed departmentally, there is often not the institutional social infrastructure, physical layout of facilities, or reward systems conducive to interdisciplinary research (Davidson, 2015). If there are not established avenues for interaction and agreed upon methodologies, then the time-sensitive work of post-disaster research may not have the necessary structure for interdisciplinarity (Ganapti & Mostafavi, 2018). Aside from the threats that not integrating disciplinary disaster research poses to the effectiveness of research in the field, it also draws into question their academic contribution. When this work is not institutionally supported and accepted, it can lead to duplicated and disjointed publications regarding the same topic but from other disciplinary lenses (Sherman-Morris et al., 2021). Current organizational structures preserve disciplinary boundaries even when to do so hampers the progress of representative, meaningful research. This leads to researchers from different fields struggling to agree on shared epistemologies for interdisciplinary research, resulting in a time-intensive integration process that is repeated every time a new interdisciplinary project occurs. This costly time commitment can dissuade some researchers from embarking on this type of research.

Another commonly reflected upon issue within interdisciplinary research is that of disciplinary superiority. According to Green & Andersen (2019), “disciplinary imperialism” describes the situation in which disciplines attempt to collaborate; however, one discipline, knowingly or not, will work to impose its methodology on the work of the other disciplines. The threat of this phenomena becomes more pronounced and is further exacerbated by one of the most fundamental issues that can arise in interdisciplinary research. That is the issue of distrust. When individuals do not have the social infrastructure mentioned previously to establish connections outside of the research, they do not afford themselves the opportunity to develop trust with the other researchers. When the researchers do not have complete knowledge of the other fields with which they are working, they will feel compelled to translate the information into their own field especially if that researcher has not first established professional and interpersonal trust with their disciplinarily disparate colleagues.

In working across disciplines, different methodologies often present points of difficulty, many of which have already been mentioned. Common differences not discussed above include the extent of community stakeholder involvement, reducibility of models, and the objectivity possible or desired in research (Eigenbrode et al., 2007). Additional differences include: "research design, sampling, data collection, analysis, and interpretation of results" (Jakobsen et al., 2004). A commonality in these differences is the difficulty in data collection and analysis. "Sampling to answer 'what,' 'how,' and 'why' questions using data from a single, integrated instrument remains an unresolved challenge" (Lynam et al., 2020). This issue of data integration is a common point of consideration for researchers in a wide range of interdisciplinary fields, not only disaster research. Within the field of biomathematics for example, researchers struggle to agree upon the data requirements as well as the scale and level of this data (Newgreen et al., 2019). Although disaster research is not unique in its struggle with data integration, it does have the added pressure of a short time horizon for collecting data post-event. The perishability and dimensionality of this data requires that common understandings of data requirements need to be reached for long-term theory and not on a case-by-case basis (Ge et al., 2021). Hence why more interdisciplinary work must be done specifically in the field of disaster research to develop robust and generalizable theory for data integration.

Not all the challenges presented here are addressed in this dissertation as some are far more systemic and thus will require the pivoting of policy and incentive structures. However, for those topics which will benefit from use cases and tool development for best practices, this dissertation strives to address these with a particular emphasis on the underpinning issue of data integration across different fields.

Outline of Objectives

As the title of this dissertation indicates, the objective is to develop a multi-faceted approach to data integration that can be viewed as a suite of integration tools and strategies. The primary areas of investigation are data collection strategies and integrated modeling techniques. These tasks are laid out with the central and driving goal of developing interdisciplinary techniques for hazards research in the implementation of a predictive outmigration model, but broader application is intended. This will be accomplished by exploring the following topics whose interconnections are characterized in Figure 1.

Chapter 1: Investigate previous attempts at interdisciplinary data integration: A literature review was undertaken to contextualize the primary aim of this dissertation and ensure that any solutions presented later were informed by lessons from previous attempts at interdisciplinarity and data integration.

Chapter 2: Develop a social susceptibility metric for rapid field deployments: A social susceptibility metric was developed to select communities in a longitudinal field study based on not only damage but also long-term community outcomes (Johnston & van de Lindt, 2024).

Chapter 3: Midwest Field Study Data Collection and Processing: A dataset tracking damage and recovery of the building stock in six communities impacted by the December 2021 tornado outbreak was developed through coordination with a team of researchers to increase the availability of high-granularity longitudinal data.

Chapter 4: Create a Generalized, Agent-Based Tornado Sheltering Model: As a means of exploring the usefulness of agent-based modeling in situations where the objective and necessary inputs are well-defined, an agent-based model was developed to determine the viability of community shelter locations based on travel time relative to warning time for tornadoes.

Chapter 5: Explore Agent-Based Modeling for Population Outmigration: In order to explore the usefulness of agent-based modeling in situations where the parameterization of the model is more difficult to define, a model was developed as a scenario analysis tool for population outmigration.

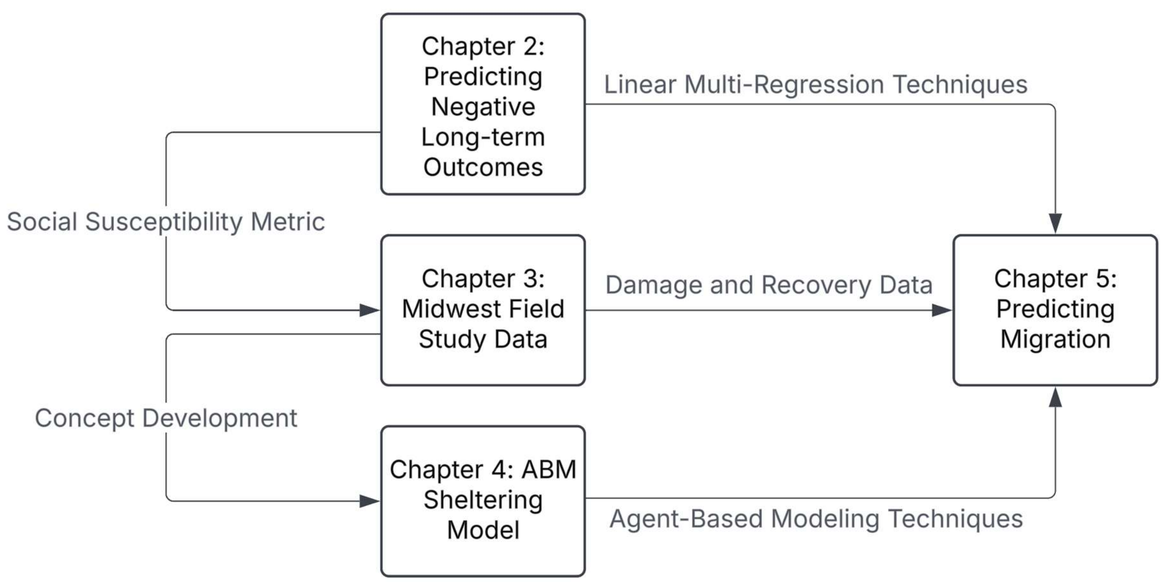

a

Figure 1: The flow of chapters and their interconnections presented in this dissertation

1. REVIEW OF GUIDANCE ON INTERDISCIPLINARY INTERACTION AND INTEGRATION (JOHNSTON & VAN DE LINDT, 2022)

1.1. Establishing Sustainable Interdisciplinarity

1.1.1. Facilitated Conversation

One of the primary and consistent tools suggested in bridging disciplinary gaps is sustained and intensive communication with colleagues from various disciplines. Several methods for communication have been recommended. Bracken and Oughton (2006) suggest the use of articulation to bridge disciplinary language barriers. Articulation involves creating a coherent message across disciplines with different dialects or metaphors, thus creating commonality to bridge differences in disciplinary language. Similarly, metaphors can be repurposed as a helpful bridging tool between disciplines, instead of a barrier to integration as mentioned previously, when appropriate context is applied. Additionally, the practice of simply having agreed upon definitions in collaborative research is an imperative tool for convenient communication (Newgreen et al., 2019).

Boundary objects are another commonly recommended method for facilitated conversation. Boundary objects are any points at which disciplines overlap, and cross-cutting conversations are easily fostered. One such boundary object is scenario development, which has been commonly implemented in a variety of fields including business, government, and sustainability planning. Scenario development, or scenario planning, centers on the idea of considering "a variety of possible futures that include many of the important uncertainties in the system rather than to focus on the accurate prediction of a single outcome" (Peterson et al., 2003). Hence this form of anticipatory planning is also referred to as "Futures Thinking" due to its usefulness in determining probable future scenarios, or simply futures. Because "scenarios are built from various assumptions, theories, and methods for studying the world" (Kröger & Schäfer, 2016), individuals can be brought together to develop and defend scenarios with individuals of disparate perspectives in order to find commonalities as well as the source of differences. This work found increased cognitive and social integration as researchers from various disciplines began to create common perspectives and derive agreed upon results (Kröger & Schäfer, 2016). Within another interesting scenario development study, five contrasting scenarios were used to predict the future of Brazil's agriculture based on a series of interdisciplinary variables, a method which will be discussed further in the data integration discussion (Gomes et al., 2020). This method of creating a few characteristic scenario outcomes from multiple natural and societal factors may be applicable in the field of disaster research as well.

A very similar boundary condition is that of disaster simulation, which is commonly used in disaster research today due to its inherent ability to create increasingly realistic results without having to wait for another disaster to occur to collect data. These simulations have been used for evacuation planning (Wang et al., 2016), nurse development training (Zapko et al., 2015), and a host of other applications to anticipate shortcomings of present systems in disaster scenarios and suggest modifications. It is often used in tandem with modeling software to increase predictive strength. Simulation has proven to be a useful tool for determining best practices, necessary policy changes, and physical infrastructure changes. Other researched boundary conditions include the sharing of experiential stories (Moezzi & Peek, 2021) and the alignment of desired outcomes (Lynam et al., 2020). These methods of social and academic integration can also provide the foundation for trust building, an imperative element of interdisciplinary work (Bracken & Oughton, 2006).

When considering communication between peers from different disciplines, it is also important for researchers to have some familiarity with the other disciplines with which they seek to collaborate. An interdisciplinary researcher requires two types of expertise: contributory expertise in their field of study and interactional expertise in fields of study with which they frequently collaborate. The first references the researcher's ability to create original work, and the second references their ability to understand work and communicate with experts from another field. As well as interactional expertise, it is advantageous to build a tacit knowledge in fields with which they plan to work consistently (Collins & Evans, 2007; Gilligan, 2021). Tacit knowledge is largely experiential, unwritten knowledge attained from working and interacting with experts from a given field. Thus, the conversations previously described provide an ideal space to develop and test this tacit knowledge and interactional expertise. Alternatively, some of this expertise can be built through interdisciplinary courses with instructors who team-teach the content for their individual areas of expertise (Gilligan, 2021; Lattuca, 2001). More long-term educational solutions for bridging disciplines are addressed in more detail in the following section.

1.1.2. Institutional Support

Providing more institutional and structural support is another means of increasing and improving interdisciplinary work. This institutional support reduces the time required to begin research and data collection efforts, as is especially important in the field of disaster research. When colleagues from various fields already have established means of communication and relationships, they are much more likely to rapidly establish the shared perspective necessary for disaster response research unlike newly formed interdisciplinary teams (Ge et al., 2021). Similarly, Faber et al. (2014) suggest that although communication is key in interdisciplinary work, it would be more effective to implement tracks of study that formally establish the frameworks necessary for successful transdisciplinary work. This strategy would eliminate the need to cultivate interdisciplinary connections anew for every project, in turn reducing the time investment required for a given project. Increasingly, collaborations that started as a series of individual projects are becoming established fields, thus remedying some of the issues previously mentioned as well as offering a level of longevity for the new field not afforded to all interdisciplinary efforts. Such is the case for the budding field of mathematical biology (Newgreen et al., 2019).

Further validating the barrier unnecessarily created by traditional institutional frameworks, a survey within the field of sustainability found that institutional responsibilities and traditional departmental frameworks sometimes prevent commitment to more interdisciplinary work. Meanwhile interdisciplinary research centers cultivated the collaborative projects that traditional institutions did not have the flexibility or interest in doing. By providing institutional frameworks that foster interdisciplinary work, it follows that more such work will be pursued (Nastar et al., 2018). It is important to acknowledge several key continued initiatives for interdisciplinary research. Along with some other organizations, the U.S. National Science Foundation sees the potential in these interdisciplinary efforts and have established dedicated funding and resources for researchers working in areas that exceed the bounds of a single discipline. They have included "Growing Convergent Research at NSF" as one of the 10 Big Ideas for Future NSF Investments, a portfolio of concepts "that will drive NSF's long-term research agenda" (National Science Foundation, 2016). Various universities, health organizations, and others also offer funding opportunities specifically for interdisciplinary projects.

Nevertheless, these initiatives for supporting interdisciplinary efforts do not serve to fully alleviate the increased difficulty in procuring funding for interdisciplinary projects. In fact, the rate of success in receiving funding and the degree of interdisciplinarity has been found to be inversely proportional (Bromham et al., 2016). This could be due in part to the difficulties funding organizations encounter when analyzing interdisciplinary projects. As noted in *Facilitating Interdisciplinary Research*, "Effective review of IDR proposals may not be possible with traditional peer review that relies primarily on experts in a single discipline" (National Academy of Sciences et al., 2004). For this same reason, interdisciplinary researchers may encounter barriers to publishing in journals with strong disciplinary scopes, which in turn could reduce their chances of receiving tenure (Bruzzese et al., 2020). Additionally, the traditional metrics of publication and funding success may be detrimental to the establishment of new interdisciplinary researchers if they have faced disciplinary gatekeeping when their interest areas exist in places of disciplinary overlap. Thus, it is important to simultaneously acknowledge the efforts that have been made thus far to bolster interdisciplinary efforts, while still understanding that interdisciplinary research is not yet enjoying the academic prestige of disciplinary work. With institutional support and restructuring, interdisciplinary research will gain a robustness, longevity, and academic trust that it cannot otherwise attain.

1.1.3. Additional Integration Methods

Research has documented several other best practices for alleviating interdisciplinary stresses. Metacognition is presented as a means of achieving more fruitful interdisciplinary work. By implementing practices that support "thinking about our thinking" throughout the planning and implementation of a project, many of the cognitive disciplinary barriers mentioned previously can be reduced if not eliminated. This use of metacognition also allows for more efficacy and creativity, improved team functioning, and increased adherence to project objectives (Ganapti & Mostafavi, 2018).

Another frequent recommendation for dissolving disciplinary silos is epistemological pluralism. This concept is founded on the idea that for most research tasks there are several means of viewing the topic of interest. Different disciplinary backgrounds offer these distinct and complimentary viewpoints. Furthermore, this pluralism suggests a need to collectively determine the objectives of an interdisciplinary project as well as the means that will be used to achieve these objectives. In turn this allows for a restructuring of academic understanding not limited to disciplinary boundaries (Miller et al., 2008). This method is largely universalizable, though by the same token has very little specificity in terms of how to apply such a pluralism aside from its adherence to iterative discussion and evaluation until a collective understanding is reached. Not surprisingly, recognition of epistemological pluralities allows for the development of more robust theories in disaster research (Sherman-Morris et al., 2021).

A final method encountered for combating interdisciplinary barriers is the iterative solution design articulated by Subedi et al. (2021), which involves epistemological pluralism, but it also includes the stakeholders in the iterative process. In adding stakeholders, the solution design becomes community-specific, yet the design process remains universalizable. The looped phases of this design consist of disciplinary data collection, the formation of a shared interdisciplinary interpretation of the data, and finally community feedback.

1.2. Investigating the Data Barrier

To more thoroughly understand the variety of methods used in data integration and their commonalities, a review of 29 papers that involve some level of data integration is conducted. The comprehensive list of these papers has been included as an addendum to the reference list provided at the end of this dissertation. These papers include a wide range of topics from socio-hydrology to systems biology and almost all deriving from disciplinary researchers who have begun efforts to produce collaborative integration across disciplines or interest areas. The breadth of papers sampled was deemed important based on the goal of finding the common features of research as it transcends any disciplinary boundaries. Furthermore, the goals of this research are drastically different. Some are aimed at producing a means of data storage or indexing, while others are interested in producing integrated, predictive models. By reviewing this variety of integration efforts, the trends presented in the subsequent graphs can be deemed more representative of data integration endeavors as a whole instead of being attributable to the unique features of a given focus area.

The selection process for these papers involved a general database query for data integration from which papers were analyzed for true data integration, and not merely compilation, between disciplines and/or stakeholders. From this query, papers that involved data integration, but which had minimal discussion of their method's shortcomings and widely applicable implications for future work, were also eliminated as they did not provide consequential information for the scope of this review. Additionally, the papers that documented original research were primarily selected for, with the inclusion of only three literature reviews which were shown to present insightful observations and actionable next steps for data integration based on the extent of their scope and strength of their conclusions.

It is worth noting here as well that many papers involved discussion of data integration methods that were deeply entrenched in the field of data science. These were occasionally within the scope of this review however the main interest of this review is to document cases in which disciplinary researchers had reached across disciplines to interact with other fields not where data scientists have taken the non-collaborative end results and meshed them together. Consequently, many of these data science papers did not offer conclusions with noteworthy implications for interdisciplinary efforts, so in alignment with the above, they were eliminated. Nevertheless, where a data integration system met the above requirements of collaboration and widely applicable suggestions, it is included. From these papers, data is compiled regarding the larger implications for data integration. This data has been collected based on a system involving the following steps. First, the areas of interest for this review were created as separate categories each of which will be presented individually in the following subsections. Then generalizable features of data integration were tabulated and sorted into the appropriate category; these features were condensed into groupings of relative commonality in order to explore trends more easily. Next, a binary system was applied in which each paper was given a value of one if that feature was present or a value of zero if it was not. The resulting summations for each feature resulted in the graphs throughout this section. By including a wide variety of topics and types of data integration, the resultant data can represent more holistically the prevalent features of data integration generally and not only through a single lens of interest.

The reasons for and benefits of data integration are shown in Figure 2. The most referenced benefit is the means by which data integration allows for more realistic portrayals of data and the ways that data interact across areas of study. These studies also frequently cite the benefit of data integration's ability to comprehensively analyze complex systems and offer an avenue to shared methodologies and data. The previously mentioned interest in merging epistemologies ties directly into this concept of using data integration as a means by which to attain a shared methodology and an authentically collaborative academic effort. The more cohesive, or at least noncontradictory, data from across disciplines can be, the more likely it is to yield meaningful and realistic results.

Some of the benefits found in Figure 2 also speak to the importance of epistemological pluralism in data collection and analysis. The incorporation of various epistemologies *offers knowledge left* *otherwise undetected* and suggests *increased reliability of convergent data*. Oftentimes data is incapable of fully explaining an event and requires another discipline's perspective to fully understand the interplay of objects in a system. Conversely if the same phenomena can be described via multiple theories or disciplines, this serves to prove a lack of bias in results.

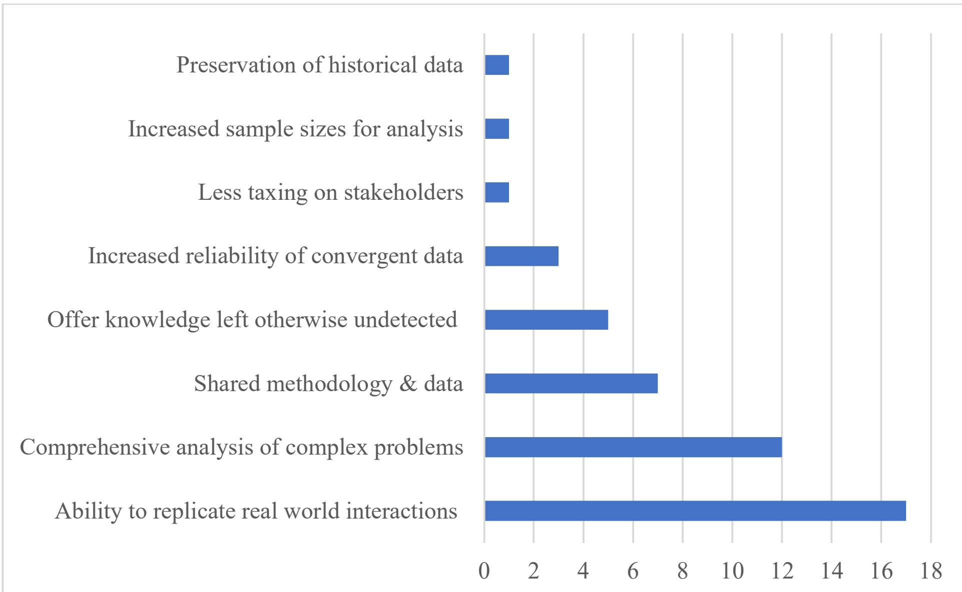

Figure 2: Reasons for and Benefits of Data Integration

Other stated benefits include *preservation of historical data*, *increased sample sizes for analysis*, and *less taxing on stakeholders*. These benefits are all associated with the centralizing of data in the form of databases or illustrative models. Creating this centralized data source allows for archival storage of previous data sets; increases sample sizes for improved reliability; and provides stakeholders with a single, comprehensive source for information.

1.2.1. Attempts at Data Integration

In the discussion of how data ought to be integrated, there are unsurprisingly different schools of thought. Within a reductionist framework, by linking more complex attributes of a system to simpler or more easily modeled elements of that system, researchers can develop constituent models from only the most simplified elements and their assigned links (Silberstein, 2002). This method offers increased certainty; however, it falls under scrutiny due to its susceptibility to disciplinary imperialism in establishing what matters in data and what links can reasonably be assumed (Eigenbrode et al., 2007). In one example of a reductionist framework, Lafuerza et al. (2016) suggests the use of a chain of simple models in order to generate near identical results to a complex model. They do this by implementing *retroductive validation*, also commonly referred to as *hindcasting* (see Figure 3), in which theories and models are developed that explain the outcomes of prior events, thus demonstrating validity of the model (Nastar et al., 2018). Lafuerza's technique assumes linearity as does any reductionist framework, which may not be accurate based on the nature of the complex model and the concerns present within complexity theory (Kallemeyn et al., 2020). The opposite of this reductionist method is the holistic method. Holistic science is interested in systems and their irreducible interactions. Representing a holistic approach, Nastar et al. (2018) doubt that integration is possible without "distortions, gaps, or inconsistencies" suggesting instead that interdisciplinary knowledge is complementary but impossible to reduce or integrate. However, they also recommend that the goal ought to be to cultivate a comprehensive understanding of a research problem rather than break it down into its simpler elements, suggesting the data then should not be integrated but kept separate in an effort to create a conceptually complete picture. These two approaches have limitations and strengths that complement each other. Where reductionism is seen as oversimplified and failing to capture the intricacies of the system interactions, holism can sometimes become too complex to be modeled with adequate certainty (Fang & Casadevall, 2011). Thus, it is often most advantageous to consider which framework offers more benefits for a given modeling scenario or if they can possibly be used in combination iteratively to achieve the most complete picture possible without sacrificing too much certainty.

As some of the methods presented in Figure 3 may not be commonly encountered, they will each be explained in order to provide adequate contextualization regarding the variety of methods formulated to combat the challenges of data integration. The largest group represented in this review are those studies that aim at data compilation, creating databases, generating illustrative models, or indexing of sources for data retrieval. These have been grouped together because they are useful sources of information, but they do not have predictive power on their own. Also noted in Figure 3 are the "top-down approaches" and "bottom-up approaches." Top-down models use a variety of statistical methods to represent systems-level behavior based on complex rules governing interactions. Bottom-up techniques rely on a method of understanding a complex system via the formation of isolated entities that are then connected through documented interactions. The bottom-up techniques presented in this review include matrices, network-based integration, modularity, and agent-based-modeling. Matrices compile information of disparate types and sources to determine correlations (Weckwerth, 2011). As a more visual representation of this, network-based integration involves individual objects of a system interacting with each other along "edges." This is best understood as nodes connected via lines in which lines are edges and nodes are objects in the system (Weighill et al., 2019). As an extension of this concept of networks, modularity typically operates such that a given discipline, focus area, or subtask constitutes a module, and that module can be connected to other modules via any number of realistic interdependencies in the systems of interest. As defined by Hinkel and Klein (2009) modularity is "the idea of encapsulating expert knowledge in the form of self-contained modules and making them available to others via well-defined interfaces." A common concern with modularity is the chance of contradicting methodologies across modules. A clear understanding of roles, acceptable inputs, and goals helps combat the creation of contradictions within the system. The last of these bottom-up approaches is agent-based modeling. This modeling creates agents that have simple rules of interaction with other agents and the environment around them. These rules then play out in iterations creating a complex system. Agent-based modeling has been used in the field of ecology and anthropology for many years and has more recently come into use in disaster research (Chen et al., 2006). The potential of this modeling technique has not yet been fully explored within the field of disaster research.

The remainder of the techniques presented in Figure 3, have primarily been used in combination with one of the three techniques already discussed as means of validating, simplifying, or further contextualizing the data. The process of "quantifying qualitative data" involves somehow codifying or classifying data that cannot otherwise be numerically modelled or visualized. In this review for instance, we classify the features of the papers being reviewed into a series of somewhat generalized descriptions and then assign values of 0 and 1 to indicate whether that feature is present in the study. The usefulness of this method is inherent in that patterns become markedly more recognizable; however, this does work to flatten the information, not expressing the extent to which a given feature is present in that paper. Hence why it is necessary to also conduct more in-depth reviews of the work presented here and speak to unique elements not fully captured in the graphs presented. In this way, this review also demonstrates the use of "case analysis/narrative" as means of exploring the nuances presented in these studies. Some papers reviewed also utilized "scenarios." In addition to facilitating conversation as discussed previously in this section, scenarios can also represent a nearly exhaustive list of plausible outcomes. Thus, given a series of logically cohesive inputs (the scenario) the model can generate a more certain outcome than when no knowns are input. For example, in Gomes et al. (2020), his concept of scenario development allows for social variables that often present sizable uncertainty to be manipulated to create plausible scenarios labeled: Sustainability (Green Road), Regional Rivalry (Rocky Road), Inequality, Fossil-Fueled Development, and Middle of the Road (Gomes et al., 2020). This land use example may be adaptable to disaster research in which there could be similar plausible scenarios implemented surrounding a disaster event, and it could similarly serve to alleviate some of the large uncertainties that have as of yet been irreconcilable across disciplines.

Other methods noted in this analysis include *triangulation or convergent validation*, *signal-based integration*, *portfolio attribute development*, and *diffraction*. *Triangulation* is a validation method, which validates a model by arriving at the same answer in multiple ways, often called convergent validation (Fielding, 2012). The second method above, *signal-based integration*, involves the tracking of a variable as it correlates to changing inputs (Weighill et al., 2019). For the third, Tobi (2014) proposed the formation of a single complex attribute based on the implementation of portfolio representation of measurement, noted in Figure 3 as *portfolio attribute development*. Note, this technique has data level requirements to operate effectively and thus may not be universally applicable. Lastly, Uprichard and Dawney (2019) warn against forcing data to integrate, suggesting instead the concept of *diffraction* which allows integration only when apt and meaningful. The methods presented in Figure 3 serve to both show the variety of attempts at data integration and orient the reader for the following sections in which challenges and shortcomings will be presented. Although there is a wide range of methods, the issues they encounter are commonly shared.

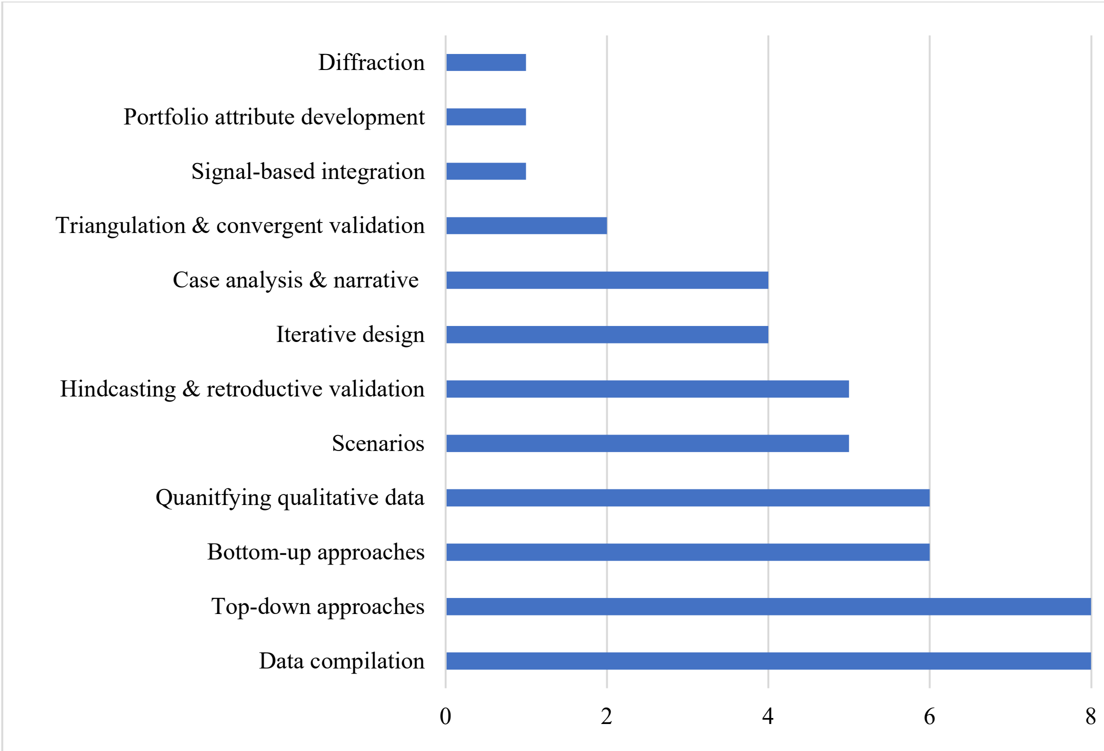

Figure 3: Data Integration Techniques

1.2.2. Feasibility and Appropriateness of Data Integration

Holistic frameworks suggest that systems are irreducible, but do not speak to the issue of how this suggestion would function within the context of modeling, thus ignoring the emergent technologies that could offer new solutions for disaster research (National Academy of Sciences et al., 2004). Similarly, Miller et al. (2008) recommends that any attempts at data integration are not generalizable and must be formed anew in every context: "… a reorganization of multiple, potentially equally valid ways of knowing requires a negotiation governed by the specifics of the question and the composition of the research team". However, there is an interest in defining an agreed upon interdisciplinary methodology for research (Tobi & Kampen, 2018), and there is presently a desire for actionable items and reduced abstraction of solutions as the result of research (Rodela & Alašević, 2017). These two goals are linked in that by accomplishing the first, the researcher is notably closer to achieving the second. Some work would also suggest that data integration should only be done when to do so is useful or illustrative (Lynam et al., 2020) or when it does not introduce an unacceptable amount of uncertainty (Nateghi et al., 2021). Another point of concern is reliance on potentially biased or inadequate expert opinion. The concept of who qualifies as an expert is unregulated and thus the results of expert opinions will likely not offer sufficient standardization for consistent interdisciplinary efforts (Brink et al., 2020). Additionally, an inability to discern between established fact and a previous conjecture based on an expert opinion also poses concern in that different disciplines may interrogate the validity of data differently (Newgreen et al., 2019). These are all necessary warnings to heed when attempting data integration, and it is within this context that this review attempts to evaluate the challenges faced with different integration methods.

As shown in Figure 4, the most commonly occurring challenge is indeed the differences in syntax, data format, data type, and data granularity. This category is admittedly large; however, to divide these elements into separate categories did not seem appropriate because they are all approaching the common difficulty of data heterogeneity. This issue consistently plagues data integration efforts, and as of yet, no universally accepted solution has been found. The next challenge is that of different available input and output requirements. Some of the research reviewed noted this challenge in the context of the scale, focus, and purpose of measured data versus desired outcome. Often data is being pulled from sources that collected data for a different purpose than the research being conducted, and so the data does not offer adequate scale or appropriate focus to conclusively answer the research question. Considerable efforts have been made to address these concerns of data generalizability through semantic technologies. As briefly mentioned in the discussion of dialects and semantics, semantic technologies allow machines to not only combine data but engage in contextualization from different sources and interpretation for different uses. This would in turn allow machines to assist in, if not automate, the process of integrating data in meaningful and consistent ways through the establishment of a linked data infrastructure. Within ontology, this linked data infrastructure is often referred to as the Semantic Web, a web of data that removes the barriers between knowledge sources. To pursue this as a possible solution for the barriers faced in disaster research, there will need to be improvements in data sharing and agreed upon data structures (Shvaiko & Euzenat, 2013).

Most of the other challenges presented here are somewhat self-explanatory or are alluded to elsewhere. However, a few points of contextualization and elaboration may be helpful. First, the challenge listed as *data integration must have robust rationale* in Figure 4 has two contexts: avoiding contradictions in the underpinnings of merged epistemologies and knowing when integration may be unnecessary. Second, it is worth noting that when top-down approaches are isolated in this analysis, the most frequently referenced challenges are *the unpredictability of human behavior* and the *different input/output requirements*. These challenges speak to the rigidity of the modeling requirements for top-down approaches and their inability to adequately represent the intricacies of social systems. As a last and lengthier point, Newgreen et al. (2019) urges practicality in interdisciplinary modeling, acknowledging the purpose as "to provide insight into relatively complex systems, not to produce facsimiles including every element." They continue by explaining that this requires omission, simplification, and approximation in data when appropriate, necessary, and such that all essentials will still be considered. However, this manipulation of data could once again be susceptible to disciplinary imperialism in that the members of one epistemology would decide what matters in the data and what is most useful in modeling (Green & Andersen, 2019). This may result in a less than comprehensive model, failing to reflect the perspective of all disciplines involved. The concepts of omission, simplification, and approximation are further problematized in the papers analyzed, contributing to the challenges in Figure 4 listed as *knowing when discrepancies are meaningful* and *integration must have robust rationale*. However, this frequently applied method of data handling also serves to alleviate some of the difficulties associated with system complexity and unpredictability of human behavior, noted in Figure 4 as well. This example is not unique in its ability to facilitate one element of integration while simultaneously further complicating another, demonstrating the balance that must be struck in sound data integration.

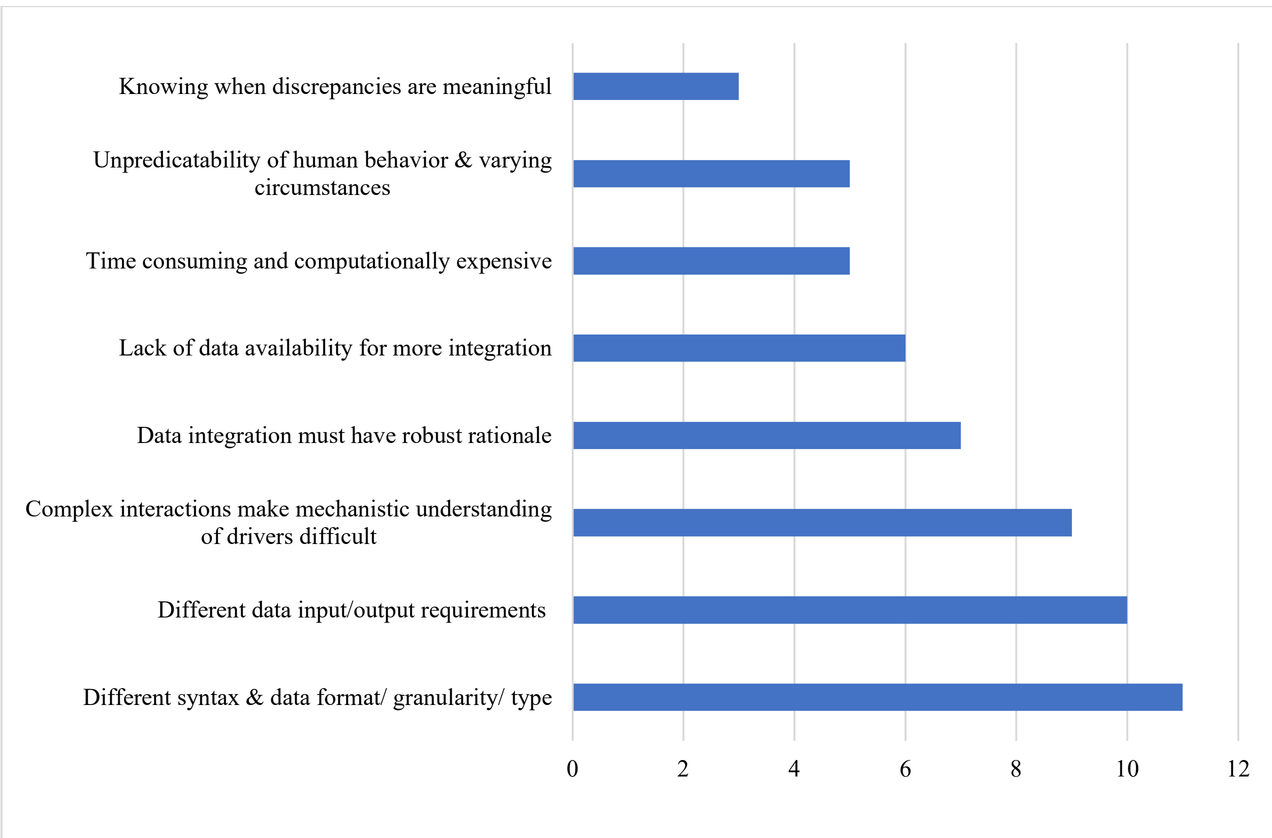

Figure 4: Challenges of Data Integration

1.2.3. Shortcomings of Previous Attempts and Recommendations for Future Research

In curating previous techniques or crafting new ones, it is useful to examine the points that these papers listed as shortcomings of their research, as well as their advice for future work. In Figure 5 and Figure 6, these features have been listed separately primarily for the purpose of increased readability and more cohesive phrasing.

As seen in Figure 5 the primary shortcoming noted in the studies reviewed was *model parameterization bias and substantial uncertainties*. These were grouped into a single category as they are frequently referenced together and contribute to difficulties in attaining accurate and actionable items from research. Sawada and Hanazaki (2020) connects these factors in stating, “The major limitation of socio-hydrological models is that they are often inaccurate due to the uncertainty in their input forcing, parameters, and descriptions of the processes.” Furthermore, situations are frequently encountered in data integration in which not all values are known for a given object in the system. The common solution for this is the omission of the object. This parameterization bias is problematized by Le Sueur et al. (2020) who explain, “One limitation of this approach is the inevitable loss of information either due to differences in granularity or data capture; another is that the final patient group may suffer from selection biases.” These difficulties in the creation of models due to incongruencies and uncertainties in data are a common and not easily solved problem in data integration. They also further exacerbate issues of epistemological difference. In working to explain a system, different disciplinary perspectives will likely value, and consequently, prioritize different parameters in the explanation of a system. Thus, this shortcoming serves as yet another warning to cultivate collaborative interdisciplinary teams grounded in academic trust.

The next noted shortcoming is *difficulty in capturing the intricacies in systems.* This is typically an issue associated with the frequently insurmountable limitations of data acquisition and analysis. Oftentimes this shortcoming occurs within the context of dynamic systems, in which case models may no longer be valid as the situation progresses as could be the case for some sea-level rise modeling (Kulp & Strauss, 2019). This leads to the next shortcoming of *insufficient data to ensure the nature of relationships*. This speaks to the challenges in determining the rules of interaction in a system. This difficulty has two sources. The first is once again epistemological biases in which different perspectives offer competing theories of what matters in a system. The second difficulty is simply a matter of not being able to accurately model some elements in systems nor fully extricate a single element in order to determine its interactions with the system. To the first, the solution is once again epistemological pluralism. To the second, there is no generalizable remedy, though a thorough review of systems science and how it ought to be applied may be helpful. This is of course the nature of the shortcoming listed as *no automated or generalizable solution*. It explains the frequent suggestion of best practices, as there are situations in which there are no other means of explaining how to combat the intricate systems involved in data integration and interdisciplinarity. However, some focus area specific solutions can and have been offered, hence why that shortcoming was only noted five times in the review. And finally, *sequential analysis* references those studies which problematize the use of procedures that act sequentially, not allowing for the realistic interaction of system elements. This can create skewed data based on the order in which analysis is conducted. This shortcoming is avoided through implementation of interdependencies and iteration in the model.

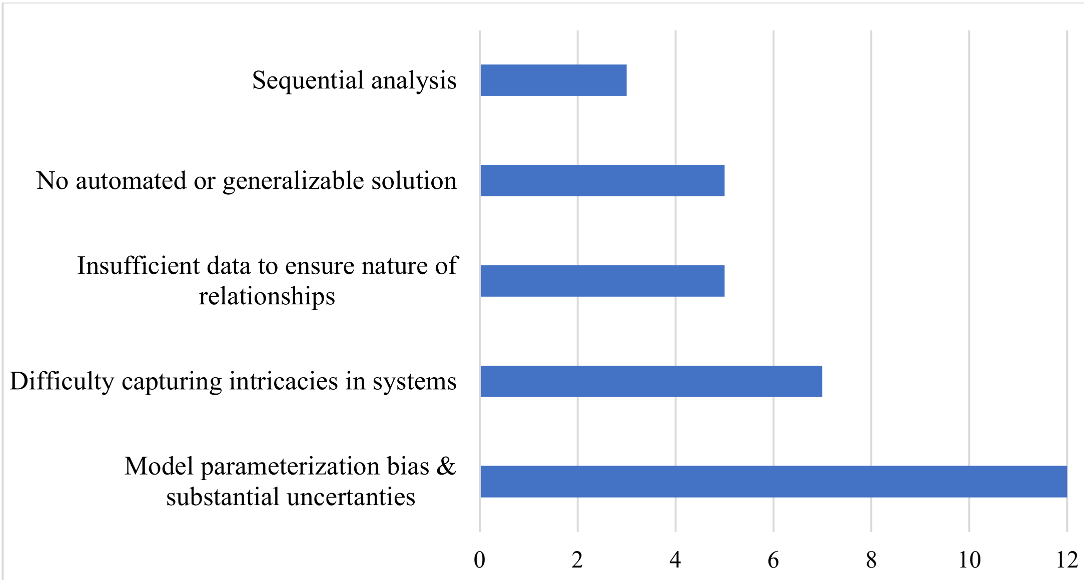

Figure 5: Shortcomings of Previous Research

Figure 6 notes the suggested future work from the studies reviewed. These papers are much more consistent in the suggestions. The first of which is *more standardization of data & methods*. This would undoubtedly aid in the more seamless processing of data and in the increase in data availability through the formation of more comprehensive databases. This could as mentioned previously be supported with semantic technologies, and vice versa, the usefulness of semantic technologies will increase as the data becomes more consistently shared in standardized formats. Then, there is the suggestion for *further integration with other networks*. This would create increasingly comprehensive and realistic models, thus bolstering their predictive power. *More quality shared data & models* speaks to the need for more reliable data. Models frequently require an extensive amount of data to function and even more data to be validated. The more readily quality data can be acquired at little to no cost the more powerful models can become.

There is also the matter of staying mindful of people in data integration to prevent inadequately capturing the social dynamics of the problem (Lynam et al., 2020). This recommendation is included in Figure 6 as *incorporation of more social dimensions*. These dimensions are sometimes, consciously or not, omitted due to the difficulties in modeling human behavior (noted in Figure 4). However, doing so introduces epistemological bias and fails to adequately represent the system. The final piece of advice offered is *more multi-scale and scalable analysis*. Frequently the different scales can modify the results of the study. By creating a scalable model, the results at different levels of granularity can be analyzed to provide more insight into the interactions within the system. In addition to these recommendations, some papers also more generally recommended continued efforts toward integration. This is not included in Figure 6 due to the lack of specificity and thus the lack of explanatory power. All these recommendations should be kept in mind when attempting data integration projects.

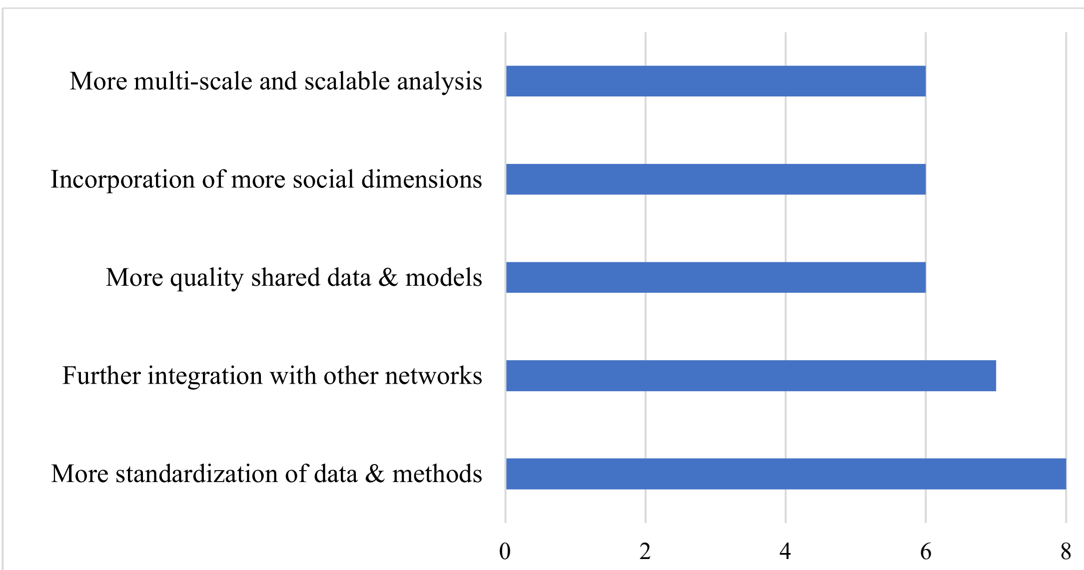

Figure 6: Recommendations for Future Research

1.3. Summary and Conclusions

Several barriers to interdisciplinary research and specifically interdisciplinary disaster research have been presented. However, most of these barriers can be effectively and adequately managed with the best practices and institutional shifts recommended in the "Establishing Sustainable Interdisciplinarity" subsection above. Upon review of the research, it seems that the primary irreconcilable barrier at present is that of data integration. The most documented challenges of data integration are those associated with the variety of data types, formats, and granularity levels that are needed for analyses. This can, in part, be alleviated by establishing standards for data collection and structures; however, this generates a greater point of contention, namely a discipline's data is inherently connected to its means of collecting that data and thus the core of its epistemology. Consequently, the matter of integrating data without invalidating any of the epistemologies involved in interdisciplinary work or oversimplifying the interactions between the complex systems during disasters presents a worthy and not easily alleviated issue. This difficulty is in alignment with the most commonly noted shortcoming of "model parametrization bias and substantial uncertainties."

Thus, this dissertation represents a renewed effort at finding means of integration that address uncertainty concerns without inadvertently introducing disciplinary imperialism through parameterization bias. The work presented herein strives to provide and implement data standards, data structures, and models that do not purge data of its interdisciplinary dimensionality. This task or rather set of tasks is well supported in the literature considering the most documented suggestion for future work was indeed "more standardization of data and methods." Addressing the issue of data integration will be the next step in providing comprehensive and predictive results for community resilience.

2. DEVELOPMENT OF A SOCIAL SUSCEPTIBILITY METRIC FOR RAPID FIELD DEPLOYMENTS (JOHNSTON & VAN DE LINDT, 2024)

2.1. Introduction

Following a spatially large hazard event, field study teams need to rapidly decide on locations and deployment strategies in order to gather perishable data. Care must be taken to not disrupt or negatively affect any rescue or other activities designed to assist the impacted population, but efforts must still be made to collect the highly perishable data after a hazard event so that lessons can be learned and in turn more adaptive strategies for community resilience can be established. In this work, the term "perishable data" should be taken to align with the definition provided by the NHERI CONVERGE Center, "Perishable data is highly transient data that may degrade in quality, be irrevocably altered, or be permanently lost if not collected soon after it is generated. Perishable data includes ephemeral information that exists before, during, or after a disaster that, if gathered, can characterize pre-existing hazardous conditions, near-miss or actual disaster events, and longer-term recovery processes. This data may need to be collected at multiple points in time across varying geographic scales to accurately characterize exposure, susceptibility to harm, and coping capacity" (Evans et al., 2021a). The development of a robust plan prior to deployment is the best way to ensure that perishable data is collected in an effective and minimally invasive way. The development of a robust plan prior to deployment is the best way to ensure that the data collection is optimally effective, minimally invasive, and ethically sound (Peek et al., 2021). However, the robustness required for this plan can sometimes make this a difficult time to spearhead interdisciplinary efforts since the complex combination of disciplines can often complicate communication and objective setting in the fast paced environment of post-disaster data collection (Ganapti & Mostafavi, 2018). The field study decision-making tool described in this chapter aims to aid in this integration for geographically large climatic hazard events focusing on longitudinal recovery studies. The underlying metric that this tool utilizes was developed using data from several different types of hazard events, consisting of tornadoes, floods, storm surges, wildfires, flash floods, coastal floods, and high wind events, as classified by NOAA. The illustrative case study included in this work documents the implementation of this framework following the December 2021 tornado outbreak, but the validation provided herein supports the use of this methodology to prioritize data collection efforts for any of the event types mentioned above. More widespread use of this tool may also be appropriate, but the validity of such alternative applications is not argued here. Also, regarding the appropriate application of the tool, a couple notes on implementation are included here for completeness. First, the exact timeline of when a team could implement this resource is somewhat dependent on the type of hazard and type of data to be collected. For instance, tornado damage can be surveyed right after the event whereas flood damage assessment may only be possible after flood waters have receded. Thus, it may be best acknowledged as a tool that can be implemented in the collection of perishable data as defined previously. If the perishable data of interest is still available, then application of the tool is still appropriate. This metric and the associated tool can be used both in initial community selection and as a screening tool for further reducing the list of initially selected communities for coordination of a longitudinal study. Second, it should be noted that the term “geographically large” is not intended to be an overly prescriptive or objective term. Rather, it should be taken to mean any event for which a particular field study team (or teams) cannot collect data for all impacted communities or areas, and thus, multi-criteria decisions must be made as to which communities will be included in the study. It is the hope of these authors that the tool proposed herein can aid in making these decisions.

Before going any further, it is important to acknowledge that the authors of this work are engineers by training. However, both authors are actively involved in interdisciplinary partnerships, and the more senior author on this work has spent almost two decades working closely with social scientists to develop interdisciplinary field studies and long-term collaborations. As such, some language and discussion throughout this work may indicate this positionality. The intended audience for this work is individuals who are endeavoring to develop interdisciplinary field studies that exist in the overlap between the social systems and the built environment. The development of this work by engineers was not intended to be exclusionary or overreaching but rather was an attempt at cross disciplinarity. By allowing for diverse perspectives in the development of disaster research protocols and metrics, the field can combat disciplinary bias, which serves to unnecessarily limit the scope of research to disciplinary boundaries (Repko, 2012, p.145).

2.2. Motivations

The motivation for this work is two-fold. First is the practical matter of providing a possible methodology for integrating built environment and social systems considerations into the decision-making process for field study deployments. The second motivation is a desire to provide a brief observation on the way vulnerability is used throughout the literature and a possible recommendation on how the field may be able to grapple with this moving forward. To the first point, the number and cost of costly extreme events has increased over the past few decades (A. B. Smith, 2020). Logically, the study of such events has grown as well (Tierney, 2018). This has been made evident as NIST, NSF, and others have funded more and more projects focused on or involved with collecting perishable data following a hazard event. These institutions have taken very clear steps to motivate the advancement of climate resilience and disaster research studies

(National Institute of Standards and Technology, 2023; National Science Foundation, n.d.; NHERI, n.d.). Some of this work is specifically aimed at collecting, storing, and making publicly available data following hazard events. This is incredibly vital work towards the goal of data democratization. However, as this work is undertaken, it is worth very intentionally considering how and when these teams will deploy.

Disaster research is inherently connected to many fields. As such, work in this field has been multidisciplinary for decades. Over the past decade, this work has become increasingly interdisciplinary, a characteristic distinct from multidisciplinary in its integration and "blending of diverse inputs that differs from and is more than the sum of the parts." Now disaster research is diving even deeper into this integration to become transdisciplinary, "transcending disciplinary approaches using more comprehensive frameworks" (Peek & Guikema, 2021). Based on this understanding of the well-established and warranted trend in disaster research to do away with disciplinary silos, it seems to follow that it will be necessary to integrate diverse knowledge systems from the very earliest stages of data collection to ensure that the data produced can address research questions that cross or better yet serve to remove disciplinary boundaries.

Setting this as the goal, we examine the state of such data collection endeavors. Some field studies have expressly focused on and succeeded in collecting data that has the capacity and necessary context to tell an interdisciplinary, and thus more holistic, story (van de Lindt et al., 2018). However, other reconnaissance and field study activities have not adopted a robust and replicable method for integrating diverse disciplinary considerations into their decision-making.

The tool presented in this work has primarily been developed as a boundary object for engineers and social scientists to explore the overlap of social systems and the built environment as explained by the Venn diagram in Peek and Guikema (2021). Previous work has already documented the fact that low resource populations are subjected to living in less structurally resilient housing than others in a community. These housing units are frequently built to low-code or pre-code standards, and as such, they are more likely to experience more significant structural damage than other houses (Peacock et al., 2007; Schmidtlein et al., 2011). However, with a changing climate not only are communities having to recover from more frequent hazard events, but the risk landscape on which that hazard occurs is changing. Put simply by Smith et al. (2018), "[c]ommunities face a… challenge to recover from more frequent extreme events considering that the spatial and design features common to human settlements tend to reflect a climate of the past." It will be increasingly important to document the disparate impact that this changing climate has on the built environment and social systems and to identify emergent characteristics of communities through interdisciplinary research. As the world's climate enters unprecedented territory, so too do the world's societies and the built environments they create.

Because the case study presented later in this work results from a tornado outbreak, a brief summary of the literature on tornado field studies may be helpful. Since the beginning of formal tornado recording in the United States beginning in 1950 several substantial tornado outbreaks have been recorded, including the "super outbreaks" of 1974 and 2011 (Corfidi et al., 2010; Flynn & Islam, 2019; *NCDC Storm Events Database*, 1996). However, it was only with this latter 2011 outbreak that researchers suggested more emphasis be placed on the social impacts of resulting structural damage (Prevatt et al., 2012). Prior to this, engineering-focused investigations into outbreaks exhibited the level of disciplinary isolation characteristic of the time with very little if any social contextualization of structural damage (Johnston & van de Lindt, 2022). As documented in the U.S. National Institute of Standards and Technology (NIST) Technical Investigation of the May 22, 2011, Tornado in Joplin, Missouri, disasters "provide unfortunate but important opportunities to learn from the performance of structures, emergency communications, and human behavior during catastrophic events" (Kuligowski et al., 2013). If there is interest in human behavior as well as in structural performance, then it follows that there should be interest in collecting data for not only the most structurally damaged communities but also communities with less damage but varying levels of social susceptibility to document disparities in community-level response and recovery following the event. This is especially the case for longitudinal studies that are well positioned to document disparities in recovery and link these to social as well as physical and economic disparities in the pre-disaster data.

For example, assume a hazard event impacts a region with three degrees of intensity as schematically represented in Figure 7 in which the shaded and numbered regions represent these intensity zones with "0" being the region with no perceptible event occurrence and "3" being the region of highest event intensity. Then it will be of interest to not only consider which communities experienced the most structural damage, but also which communities represent unique social outcomes. Consider the color-coded pins as communities for which the color assignment represents varying tiers of predicted social outcomes, and it is the interest of the research group that every tier of predicted social outcome be represented in the dataset. Then the selection of community protocol put forth in such work would recommend that a team select the communities shown in the right image of Figure 7. It should be noted that this example shows only one socially representative subset of communities from those in the left window of Figure 7; however, alternate groupings could be selected if, for example, the study requires a larger number of communities or if an emphasis is to be placed on those communities that are anticipated to have worse social outcomes.

3
2
1
0

Figure 7: Left: Schematic of all Communities in an impacted region. Right: Schematic of Selected Communities based on physical damage level and social susceptibility.

It should also be noted that an implicit first step in deciding what communities to survey is a filtering of only those communities which saw structural damage according to local sources such as news coverage. These local sources are likely to detail the structural damage experienced but are not likely to acknowledge the social context in which this damage is occurring. As such, the resource provided herein emphasizes the social context and less so the structural damage because researchers can explore local reports to perform a cursory evaluation of the damage, but they have less empirically validated data to get a similarly brief understanding of the community's social susceptibility to negative long-term outcomes. As illustrated in Figure 17 in this work, structural damage should remain a focus in the community selection process, but so should the social context in order to better understand the way social factors contribute to community recovery. These insights can further the understanding of how social systems and the built environment simultaneously shape each other.

By conducting research that considers social and physical impacts from the beginning of data collection through analysis, a more interdisciplinary, and thus holistic, concept of how these systems interact after a hazard event can be achieved. Furthermore, setting a precedent for how integrated, disciplinarily diverse teams can engage in areas of overlap allows for the dissipation of several barriers to interdisciplinary collaboration outlined by Sapat (2021). As explained by Ge et al. (2021), “One of the perceived strengths of an interdisciplinary collaboration is its potential for bringing a new, shared perspective to a problem. Because this cannot easily be accomplished by fitting a set of disparate pieces together at the end of a project, it implies a need to discuss how teams can learn to inhabit an interdisciplinary space from which to define and scope problems, identify the data types of interest to such problems, and identify their associated solutions.” At its core, the concept of beginning data collection with a team of researchers from a diverse range of fields and continuing that collaboration and integration throughout the data analysis is part of what differentiates multidisciplinary projects from interdisciplinary ones, the latter of which is required in community resilience and natural hazards field studies “to understand the factors shaping direct and indirect impacts, as well as restoration and recovery processes” (J. W. van de Lindt et al., 2018).

As alluded to previously, the second focus and motivation of this work is to propose and implement slightly altered language in the discussion of social vulnerability in order to avoid generalizations that may not hold true in all contexts and processes. The concept of defining the social characteristics of a community is certainly not new and several social vulnerability metrics have been developed both within the context of natural hazards as well as other fields of study that demonstrate the formalized exploration of this topic. However, effort was taken in this work to develop a new social metric, which very intentionally is not described as a social vulnerability metric. This decision to develop a new metric and assign a new term to it is made because of the weaknesses of the term as articulated in (Tierney, 2018):

> First, even as evidence continues to mount, showing that patterns of inequity exist across communities and across hazards, by and large studies do not provide a full explanation as to why this is the case. Many studies of vulnerability are place-specific and tend to involve analyses that center on particular variables, such as locational information on particular toxic facilities or on natural hazards, race, and ethnicity. In consequence, a general theory of vulnerability has yet to emerge. Second and relatedly, as can be seen in efforts to develop vulnerability indicators by using census data and other sources, the tendency in the literature is, with a few notable exceptions, to describe vulnerability as an attribute of particular segments of the population —in other words, as a state —when for theoretical purposes vulnerability is more appropriately conceptualized as a process in which different groups are affected by changes in the broader political and economic environment that either reduce or increase their propensity for loss.

These weaknesses in the logic and application of social vulnerability metrics seem to derive from a larger desire for generalizability that manifests in three dimensions: space, components, and time. The first weakness described in the excerpt reflects a desire to derive a theory and/or metric based on one geographic area that can be used everywhere, spatial generalizability. The second weakness logically follows from a desire to derive an understanding of a process from a single or a small set of aggregate variable(s). To convert a dynamic process into a state (static) variable, one must freeze the process in time and solve for a finite set of components which will become the aggregate state variable (i.e. vulnerability as a state). Thus, characterizing vulnerability using a state variable that was derived in this way inherently assumes that the temporal context and outcomes of interest will remain constant. Assuming the process remains constant over a given dimension infers that any value can be taken in that dimension without impacting the process. Put simply, assuming time and components remain constant in order to define a state variable implies that they are generalizable traits. Thus, a state variable is decontextualized from the time and components with which it was initially solved in order to create a generalizable, single metric. Use of a single quantitative metric to convey social vulnerability has been warned against in the literature (Spielman et al., 2020). Trying to achieve generalizability in all three of these dimensions simultaneously predicates that the developer of such a metric incorporate assumptions that risk rendering the tool meaningless only because the situations in which it is meaningful, and thus recommended for implementation, have not been clearly outlined. The need for this clear delineation of when and how a metric like this should be used is stressed in Bakkensen et al. (2017).

In order to address these concerns, the term social susceptibility is used in this work rather than social vulnerability. The definition of social susceptibility is based upon the Wisner et al. (1994) definition of vulnerability, "the characteristics of a person or group and their situation that influences their capacity to anticipate, cope with, resist, and recover from the impact of a natural hazard." This could be edited and augmented to arrive at a concept of social susceptibility that stresses two traits: it is a static feature that operates as an input to a dynamic process, namely vulnerability, and it is specific to a set of negative outcomes. The definition of social susceptibility thus could be "the level at which a set of characteristics that have been linked to a specific set of negative outcomes following a natural hazard are present in a community." And in use, it would not be said that a group is susceptible to some negative outcome but rather that they have high susceptibility to a particular negative outcome following a hazard event. For example, in this work the term social susceptibility is used independent of context at times to be mindful of verbosity, but wherever it appears in the text it should be taken to mean "social susceptibility to long-term negative affects pertaining to decreases in population, housing stock, population, income, and educational attainment." This definition of a new term that is meant to be applied when relating specific inputs of the vulnerability process to specific outputs serves to reduce the number of generalizations inherent in this work and give a clearer path forward in the understanding of vulnerability as a process. Of the three generalizations pertaining to space, components, and time, susceptibility metrics are generalizable in space if they are developed using spatially diverse data; they do not allow for generalization via component; and they may still present weaknesses in time generalization due to the static nature of the metric. Hence, social susceptibility solves part but not all the problems of generalizability. Hopefully it is a meaningful step to shifting the conceptualization of vulnerability from a state to a process, so that when working to reduce vulnerability, policy makers can try to change the various mechanisms in the vulnerability process that oppress people rather than trying to change the people that these mechanisms oppress.

As a final definition and clarification that is necessary to understand the rest of this work, this work has adopted Tierney's understanding of recovery, which "consists of all the activities that encompass efforts on the part of those affected by disasters to overcome disaster disruption and continue to thrive" (Tierney, 2018). This is far too weighty a definition and understanding to presume that a single metric or output could capture it, as discussed above. Although population outmigration, housing availability, and restoration of economic function are important factors in community recovery (Peacock et al., 2007), the metrics used here cannot be assumed to accurately portray more nuanced behavior in the community such as climate gentrification, heterogeneity of recovery patterns across the community, and the sum total of recovery activities undertaken by local organizers. By the very nature of the monitoring factors being summary statistics, it is evident that this data has been stripped of two key dimensions in determining community recovery. First, the data have been spatially homogenized such that nothing can be known of the ways these features change over the landscape of the community. Second, the data are state variables, meaning they cannot speak to the process by which this state has been reached. Consequently, to not make any reductionist leaps, this work speaks only to the trend of the monitoring factors, describing it a stable or declining, and allows the readership to decide how this information can be used based on their specific research needs and equity considerations.

2.3. Social Susceptibility Metric Development

The concept of defining the social characteristics of a community is certainly not new and several social vulnerability metrics have been developed both within the context of natural hazards as well as other fields of study that demonstrate the formalized exploration of this topic. However, effort was taken to develop a new social metric, which very intentionally is not described as a social vulnerability metric, for this case study because previous research has documented the need for more empirically validated metrics (Bakkensen et al., 2017). The work from Bakkensen et al. (2017) also goes on to state the importance of choosing the right metric for a given application, or conversely, when developing a new metric, clearly delineating its intended use. Within the context of this mandate, no metric found in the literature had demonstrable predictive power for the long-term outcome trajectories for implementation in selecting communities for data collection after geographically large events, making it worthwhile to explore a new method for social metric evaluation.

In establishing this metric, several requirements had to be met. First, the metric had to explain long-term community behavior. Second, it had to be validated by both theory and empirical data. Third, it had to have far less data demand than previous examples because it would need to be rapidly implementable following a hazard event. And finally, it needed to be developed independent of natural hazard type or region to ensure its applicability for any climatic hazard capable of significant structural damage anywhere in the United States. Non-climatic hazards were not included in this study as the dataset utilized for this investigation only tracks climatic hazards, but application of this tool for these hazards is any area for future research. These requirements ensure that the metric developed will be useful, meaningful, and actionable for researchers, meaning that the community selection and prioritization process for future geographically large events will not only take into account the extent of damage but also the predicted rate of recovery (or lack thereof) due to social factors. The procedure for determining these factors is set forth in the following set of steps.

Step 1: Determine factors that will be used to monitor the community's response to an event. These will be referred to as "monitoring factors" and will be used as a proxy for measuring long-term community recovery.

Step 2: Generate a list of communities that experienced climatic hazard events capable of producing significant structural damage from the NOAA Storm Event Database that had event costs of more than $50 million in the years 2011-2013.

Step 3: Tabulate the increase/decrease of the selected monitoring factors in the selected communities in the years following the event.

Step 4: Select predictive factors to analyze for accuracy against the monitoring factors.

Step 5: Determine the aggregate list of predictive factors that has the largest correlation with the summed value of the monitoring factors.

Step 6: Bin each continuous predictive value into quintiles in order to more easily characterize the attributes of a community.

Step 7: Sum the normalized mean of the quintiles for each factor to get a final predictive sum.

Step 8: Further characterize this predicted outcome by dividing the sums of both monitoring factors and predictive factors into four outcome groups.

To decide on the factors that should be represented in the limited set of parameters, an evaluation of previous events was conducted. Because the 5-year ACS (American Community Survey) data is only available starting in 2010, the communities analyzed were those which experienced a climatic hazard event which exceeded $50 million in damage in the years 2011-2013 according to NOAA's Storm Events Database, excluding hail and drought as they were deemed outside the scope of most rapid response field studies. This meant that tornadoes, floods, storm surges, wildfires, flash floods, coastal floods, and high wind events were the only event types considered in the development of this tool (*NCDC Storm Events Database*, 1996; United States Census Bureau, 2020). It is acknowledged here that the threshold of $50 million is somewhat arbitrary and was primarily done in order to purge low consequence events from the dataset. Applying this criterion may have eliminated high consequence events from the dataset for which the dollar amount estimate of damage was artificially low. This represents a potential error source in the analysis.

The benchmark data was collected for each community the year before the event, then at 5 years after the event, and finally 2019. The third sampling year was 2019 as it was the last available year at the time of this metric development, meaning that for 2011 events the 2019 data was 9 years from the benchmark year, for 2012 event it was 8 years from the benchmark data, and for 2013 events it was 7 years for the benchmark data. These decreasing gaps are not ideal, and undoubtedly it will be of interest in the coming years to analyze later years of data to obtain no overlap in sampling dates, as well as to include additional event years to grow the sample size. However, as it stands at present this sampling resulted in a list of 139 communities to be analyzed for best fit. This list of 139 communities consists of both counties and county subdivisions. The granularity of the region chosen was dependent on the geographic footprint of the event and the level of data accessibility for each region.

The monitoring factors selected for the development of this tool were: Median Household Income, Population, Educational Attainment, Number of Households, and Number of Housing Units. As evidenced by this list, emphasis was placed on the shrinking of communities (population outmigration) and their housing infrastructure in the years following the event. This emphasis was placed because changes in population indicate and predicate changes in the social fabric of the community. Additionally, decreases in the housing supply at the community level can hinder the community's ability to recover and exacerbate social vulnerabilities (Prevatt et al., 2012). It is also important to note two features of this data. First, there is no control sample set of communities to determine if negative trends in the monitoring data are predicated by a disaster occurring, the decision to omit this was based on the tenant that "it is not the physical event itself that produces... losses. The event is merely a trigger; the losses are the result of processes that are internal to the social order" (Tierney, 2018). Put another way, there may indeed be negative trends for these monitoring factors in the absence of a hazard event in communities that also score high on the metric proposed here. If such a trend exists, this would not suggest the tool behaves poorly, but rather it would simply prove the argument made in the quote above. Second, this data was analyzed using a simple binary, in which a decrease was represented as a value of one and an increase or constant value was represented by a zero. This was done based on the requirement that this metric should be applicable in all size communities, which experience the same percentage decrease in population for instance in very different ways. The trends were established for: the year prior to the event to the first post-event sampled year; the year prior to the event to the second post-event sampled year; and the year prior to the event to the first post-event sampled year then the first post-event sampled year to the second post-event sampled year. These were intended to detect, respectively, rapid recovery and/or stability, prolonged recovery, and persistent recovery. All 15 binary results (5 variables and 3 trends per variable) were then summed to get a community outcome trend.

Once data had been collected and the trends had been determined for the 139 communities, the proposed social factors were compared against these results using a correlation analysis. The proposed social factors represented a combination of social vulnerability values established in the literature and new values that sought to represent previously documented, but difficult to quantify, concepts of community resilience. Several parameters commonly occurring in the literature were considered, including median household income, median age, educational attainment, race, ethnicity, tenure status, occupancy rate, population density, and rurality (Beccari, 2016; Cutter et al., 2003; Fatemi et al., 2017). This list was then expanded to include the parameters from FEMA's Community Resilience Indicator Analysis that could be readily called using the Census API for ACS 5-year data (FEMA, 2022). The summary of parameters added from this source can be found in italics in Table 1. Then several exploratory parameters were also included for consideration as proxies for various forms of community resilience and cultural capital referenced in the literature (Norris et al., 2008; Ritchie & Gill, 2007; Yosso, 2005). These included first a community self-perception parameter that attempted to capture the "self-categorization" explained in Ntontis et al. (2020) by aggregating community reviews via a common community review platform called Niche (*Niche*, 2021). Second, the role of family ties in building social capital by creating a combined factor from census information that accounted for the percent of households that were families and the number of children those families had. Third, a factor for navigational capital was proposed that used educational attainment in a slightly different way by only accounting for those individuals with college degrees or greater. Lastly, an exploratory term was defined for racial affinity groups to account for the power of collective action that minority racial affinity groups may experience, also called resistance capital (Yosso, 2005). The final parameter considered was the NOAA event cost in order to account for event intensity. The potential predictive parameters and the monitoring factors are summarized in Table 1.

Table 1:Predictive factors considered, and monitoring factors for assessing predictive power of resultant metric

Factors considered for inclusion in the predictive metric	**Monitoring factors which the linear fit model was being trained to predict**
1 - Educational Attainment	Median Household Income
2 - Median Income	Population
3 - Race	Number of Households
4 - Ethnicity	Number of Housing Units
5 - Rental Tenure	Educational Attainment
6 - Occupancy Rate	
7 - Median Age	
8 - Population Density	
9 - Rurality	
10 - Event Cost	
11 – Tiered/Scaled Value for Event Cost	
12 - Community Self-Perception	
13 - Racial Affinity Groups	
14 - Family Ties	
15 - %Less than HS	
*16 - %Population65+**	
*17 - %Limited English Households**	
18 - %Population Below Poverty Line	
*19 - %Unemployed in Labor Force**	
20 - %Female Unemployed	
*21 - %Single Parent Households**	
*22 - Mobile Homes as % of HUs**	
23 - %owner occupied	
24 - % Households w/o vehicle	
25 - GINI Index	

26 - Number of Healthcare professionals per 1,000 residents

**An asterisk denotes the selected features following a set of linear regression analyses.*

At this step, the predictive factors were binned into quintiles to delineate a community's areas of vulnerability, and these quintiles were given the descriptive categories of Low, Moderately Low, Moderate, Moderately High, and High. These categories represented less than the 20^{th} percentile, 20^{th} to 40^{th} percentile, 40^{th} to 60^{th} percentile, and greater than the 80^{th} percentile, respectively. Communities that fell into each of these bins were then given the median value of each bin, meaning the 90^{th} percentile, 70^{th} percentile, 50^{th} percentile, 30^{th} percentile, and 10^{th} percentile, respectively. These binned values were then min-max normalized and scaled from zero to one to ensure optimal results for the linear regression to be conducted in the next step.

A linear regression was conducted for all possible combinations of up to ten of the variables detailed above for which each parameter had a p-value of less than 0.05. Only combinations of up to nine factors were examined because the attempt to find satisfactory combinations yielded no results that satisfied the single-parameter statistical significance requirement. The maximum adjusted R-squared value for these various combinations was 0.3621 and resulted from the combination of six factors. The linear regression resulting in this adjusted R-squared value utilized median income, percent of population age 65 and older, percent of limited English households, percent of the workforce that is unemployed, mobile homes as a percentage of housing units, and the GINI Index. The marginal return of adding more factors was explored further by investigating the maximum R-squared value for every number of parameters. These are summarized in Table 2. Based on these values and to ensure that the metric had a small enough number of factors to be rapidly implementable, the combination of five terms was selected as the six-term combination saw only a marginal increase in the adjusted R-value of less than 1% when compared to the five-term combination. The equation for this combination is shown in Table 2 and a summary of statistical features of the model is included in Table 3.

Table 2: Combination of predictive factors resulting in highest adjusted R-squared value

Number of parameters	Adjusted R-squared value	% change in adjusted R-squared value	Parameters (coded values from Table 1)								
1	0.2356		16								
2	0.2790	0.0434	16	24							
3	0.3086	10.609	16	24	25						
4	0.3509	13.707	16	17	21	22					
5	**0.3606**	**2.764**	**16**	**17**	**21**	**24**	**25**				
6	0.3621	0.416	2	16	17	19	22	25			
7	0.3422	-5.496	9	12	16	17	20	23	25		
8	0.2465	-27.966	2	5	14	17	18	20	21	22	
9	-	-	-	-	-	-	-	-	-	-	-

Table 3: Characteristics of the selected model

Variable	Estimate	SE	t stat	p value
Intercept	1.371408521	0.59483	2.30556	0.022682862
Population 65+	4.671738751	0.68049	6.86524	2.30E-10
Limited English	-1.966429156	0.69931	-2.812	0.005670075
Single Parent Household	2.245928841	0.80608	2.78622	0.006112765
Households w/o a vehicle	2.350506492	0.83222	2.82438	0.005466983
GINI	-1.854795882	0.73262	-2.5317	0.012513357
Summary				
Number of observations: 139, Error degrees of freedom: 133				
Root Mean Squared Error: 2.56				
R-squared: 0.384, Adjusted R-Squared: 0.361				
F-statistic vs. constant model: 16.6, p-value = 1.08e-12				

$$SSM = 1.3714 + 4.6717po - 1.9664le + 2.2459sp + 2.3505hv - 1.8548gn \quad (1)$$

where *SSM* is the social susceptibility metric, *po* accounts for the population age 65 and over, *le* accounts for the percent of household with limited English, *sp* accounts for percent of single parent households, *hv* accounts for households without a vehicle, and *gn* accounts for the GINI index. This result is termed the social susceptibility metric for simplicity, but it should be noted that throughout this work, this term more specifically means social susceptibility to long-term negative community outcomes, especially population outmigration.

These individual factors are arrived at by first finding the local community value for percent of the population 65 and over, percent of households with limited English, the percent of single parent households, the percent of household without a vehicle, and the GINI index. These values are then divided by their corresponding national average for the year being evaluated.

As described previously, these parameters are then binned into quintiles. The median values of these quintiles are then min-max normalized to standardize the weighting of each parameter. The values were given a standard 0 to 1 scaling, which is shown in (2), where the X values indicate the various unscaled quintile values and x_i indicates the scaled value for each quintile. The results of this scaling are the values which ought to be used in (1).

$$x_i = 1 - \frac{X_i - X_{min}}{X_{max} - X_{min}} \tag{2}$$

Because the resultant social susceptibility metric is intended to be a lightweight tool in order to fill a unique need in the field, this optimized parameter selection was deemed sufficient for further investigation. The predictive metric results were grouped into two categories (high and low susceptibility), and each of those categories were broken into two subcategories, which form the general indicator groups. These subcategories were as follows: very high susceptibility, high susceptibility, low susceptibility, and very low susceptibility. The thresholds for these categories were determined iteratively based on fit with the monitoring factors' sums, by which very high susceptibility corresponded to marked decline, high susceptibility corresponded to decline, low susceptibility corresponded to recovery, and very low susceptibility corresponded to marked recovery and stability. In Figure 8, the gray line plots the aggregate susceptibility values for the

139 impacted communities mentioned previously against the corresponding cumulative probability for this dataset. Figure 9 shows a corresponding representation of the monitoring factor sums for the 139 communities. The subcategories described above are shown as shaded regions and the categories of decline and recovery are divided by the bold line. These figures show that the threshold chosen corresponded to approximately equivalent probabilities of each value. Descriptions of the equivalent categories are formalized in Table 4.

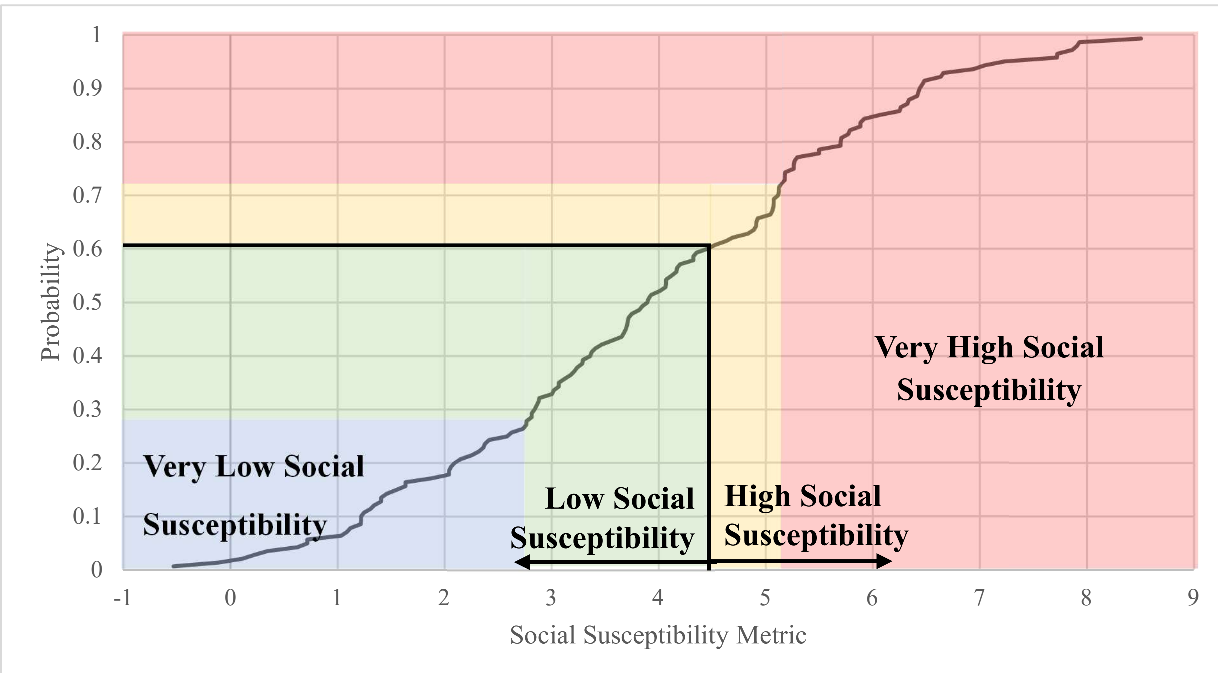

Figure 8: Social susceptibility metric - thresholds and categories

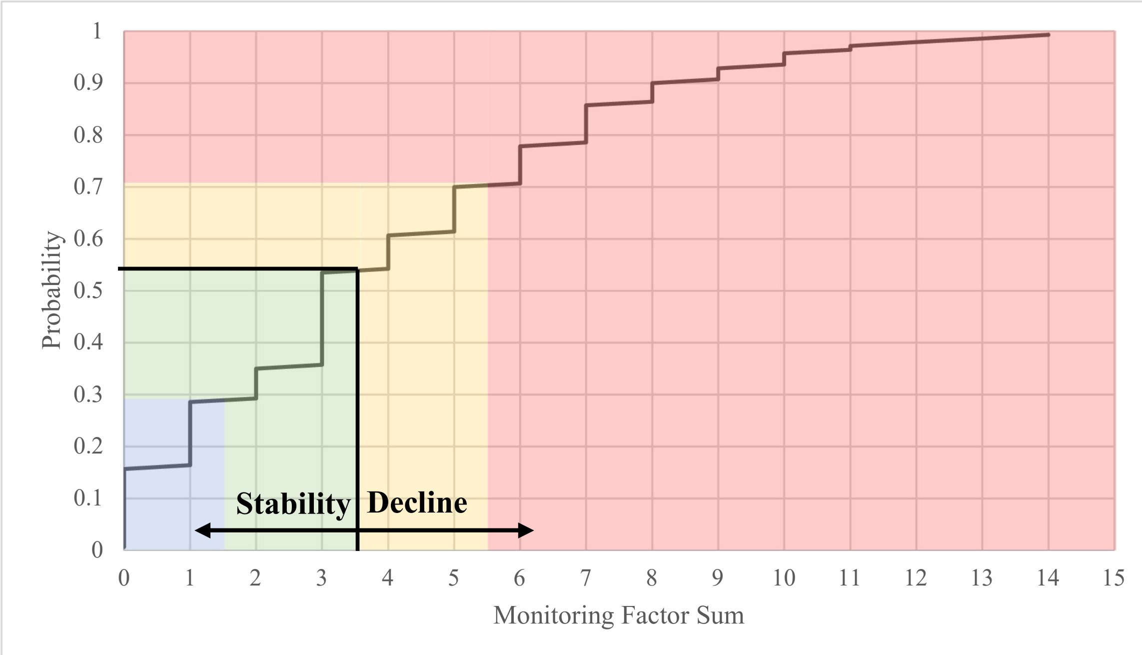

Figure 9: Monitoring factors sum - thresholds and categories

Table 4: Characteristics of Social Susceptibility and Outcome Tiers

		Score Range		Meaning of Outcome Tiers in the Years Following the Event	
Social Susceptibility Tiers	**Outcome Tiers**	**Predictive Factors**	**Monitoring Factors**	**Post-Event Monitoring Factors with Partial Negative Trends (Not Sustained)**	**Post-Event Monitoring Factors with Sustained Negative Trends**
Very Low	Marked Stability	<=2.75	<=1.5	At most 1 factor	None
Low	Stability	<=4.4	<=3.5	At most 3 factors	At most 1 factor
High	Decline	>4.4	>3.5	More than 3 factors possible	More than 1 factor possible
Very High	Marked Decline	>5.15	>5.5	All 5 factors possible	At least 1 factor, More than 2 factors possible

With these thresholds, i.e. definitions for factor stability, factor decline, etc., in Table 4, the percentage of communities predicted to see stability in their monitoring factors that actually did was 73.2%, and the percentage of communities expected to see declines in their monitoring factors that actually did was 73.7%. Thus, the metric accurately predicted stability and decline of a community's monitoring factors roughly three out of four times. Of the communities predicted to see marked stability in their monitoring factors after the event 70.3% did. On the other end of the spectrum, 60.0% of the communities predicted to see marked decline in their factors, did experience this. Visualizations of the number of communities in each category and the accuracy of prediction are shown in Figure 10. These results indicate a sufficient level of accuracy as a rapidly implementable tool for aiding in rapid post-event decision-making; however, the plot in Figure 11 shows a level of dispersion that indicates wider application beyond rapid post-event decision-making is not appropriate.

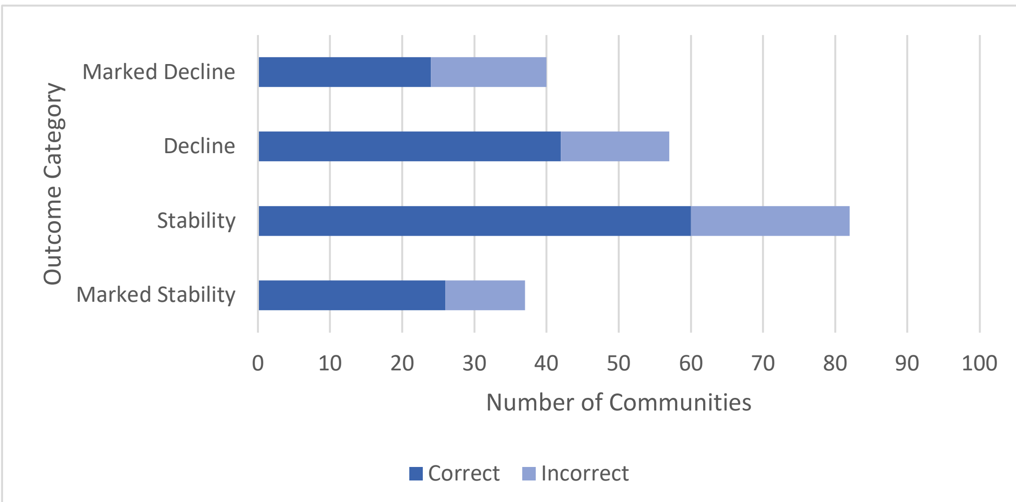

Figure 10: Communities categorized by predicted outcome state assessed for accuracy

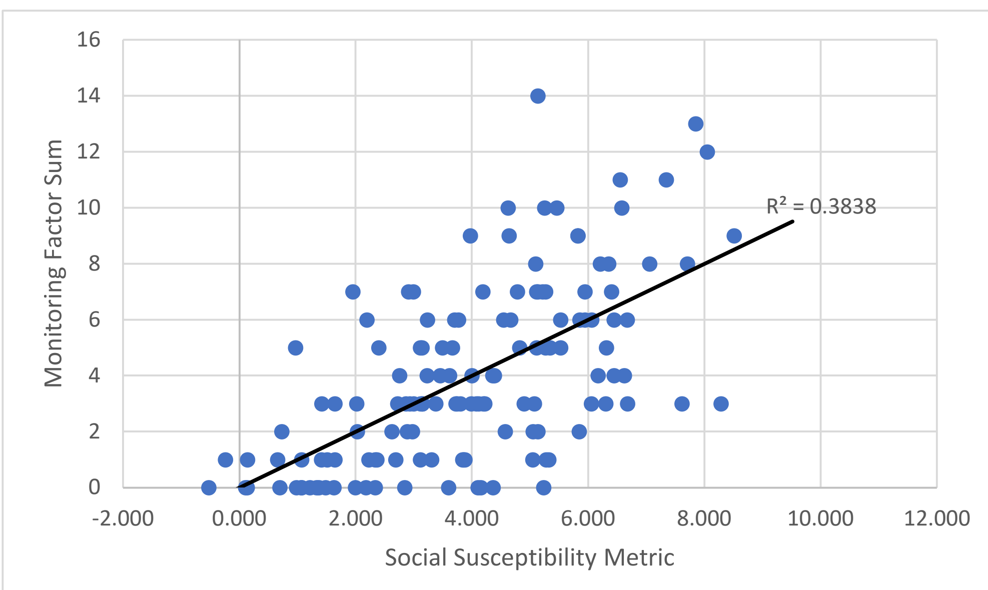

Figure 11: Plot of Monitoring Factor Sum vs. Social Susceptibility Metric

In order to express the validity of this metric in a landscape of similar social metrics, a comparison of the predictive power of this tool versus three commonly used and referenced tools was undertaken. It is important to note that predictive power is specifically referring to the alignment that the given metric has with the observed monitoring factors after the event and the results of these comparisons are not intended to cast doubt on any existing metric but only to suggest that more specificity may be required for how and when overarching metrics are used. The first comparison was to the CDC SVI tool. Because this work provides community scores at the county and county subdivision level, the SVI for the counties was able to be compared directly, while the SVI for census tracts were aggregated to the county subdivision level for comparison. The results are summarized in Figure 12 and Figure 13.

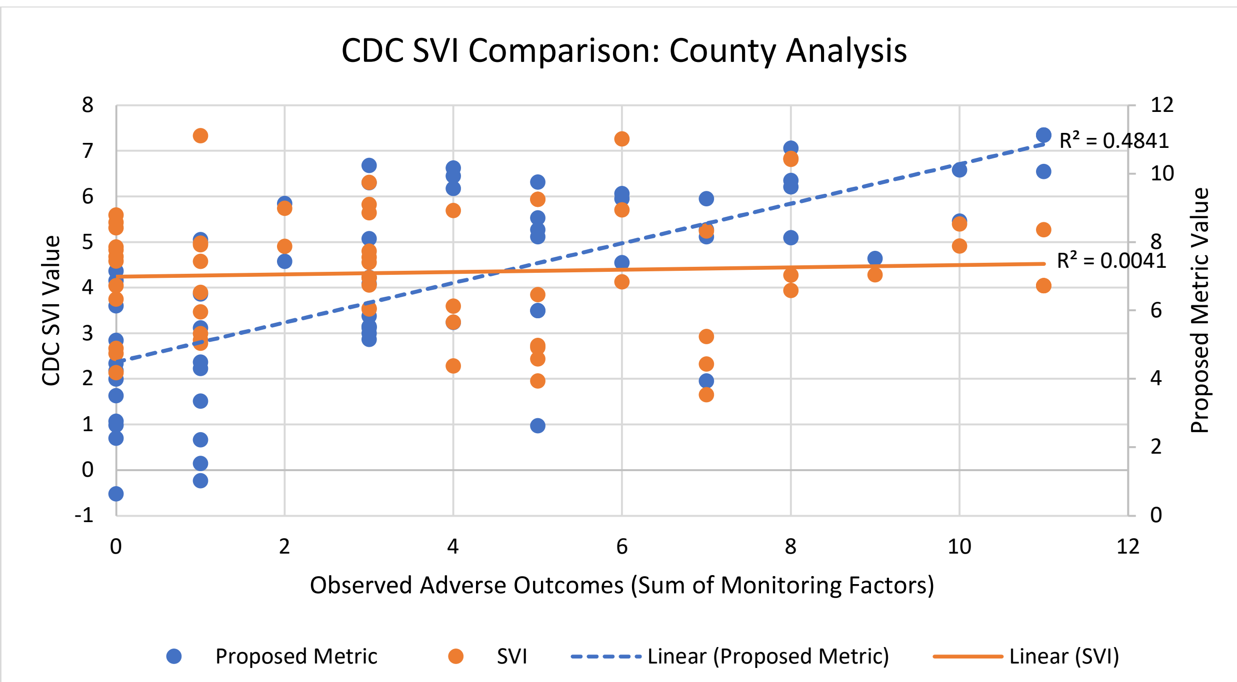

Figure 12: Comparison Predictive Power for Selected Monitoring Factors (CDC SVI County)

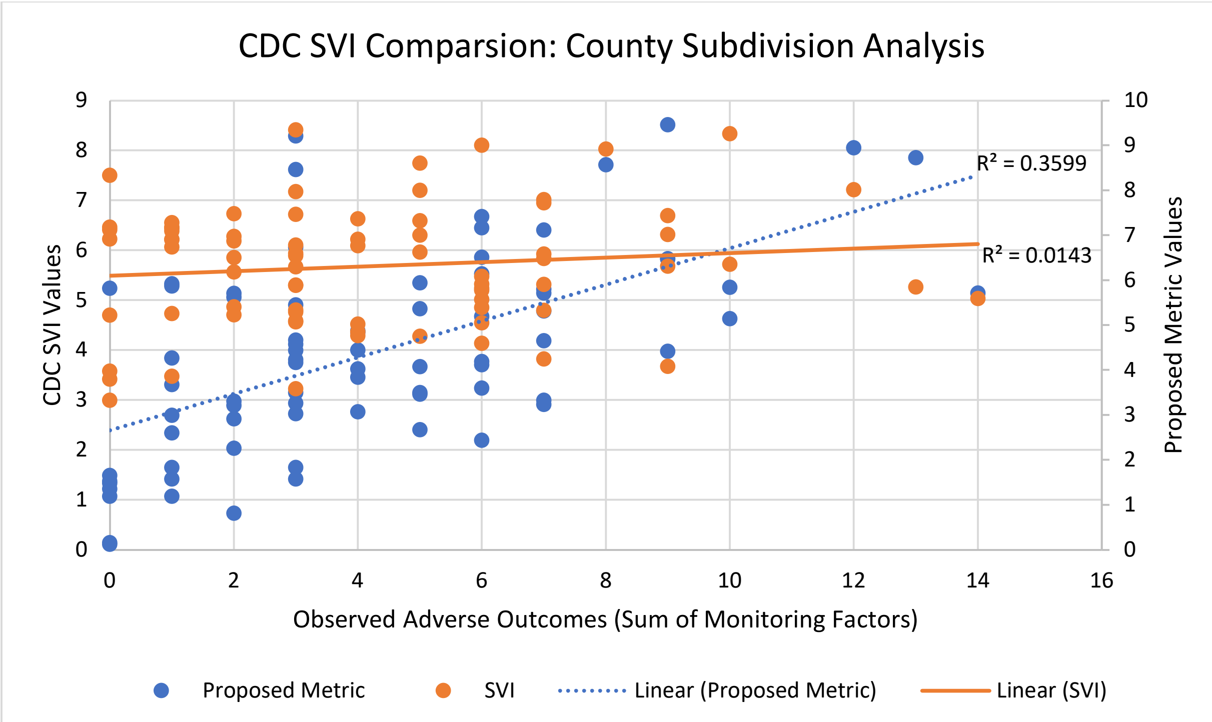

Figure 13: Comparison of Predictive Power for Selected Monitoring Factors (CDC SVI County Subdivision)

Here we can see that the CDC SVI values have almost no relationship with the long-term outcomes of interest for this work, having R-squared values of 0.0143 for county subdivisions and 0.0041 for counties. This furthers a key tenant of this work, namely that specialized metrics should be developed for specialized applications. It is also worth noting along with Figure 12 and Figure 13 that although the metric proposed in this work remains predictive at both granularities, the county analysis shows heightened levels of correlation with resulting monitoring factors when compared against the county subdivision analysis. For completeness, a similar analysis was conducted for the BRIC and SoVI metrics, which are summarized below. It is important to note that the 2010 datasets for both metrics were used, and only county-level data was available. All county subdivisions in the 139 communities were just coded to adopt their county's BRIC or SoVI metric to address this granularity issue. The results are shown in Figure 14, Figure 15, and Figure 16.

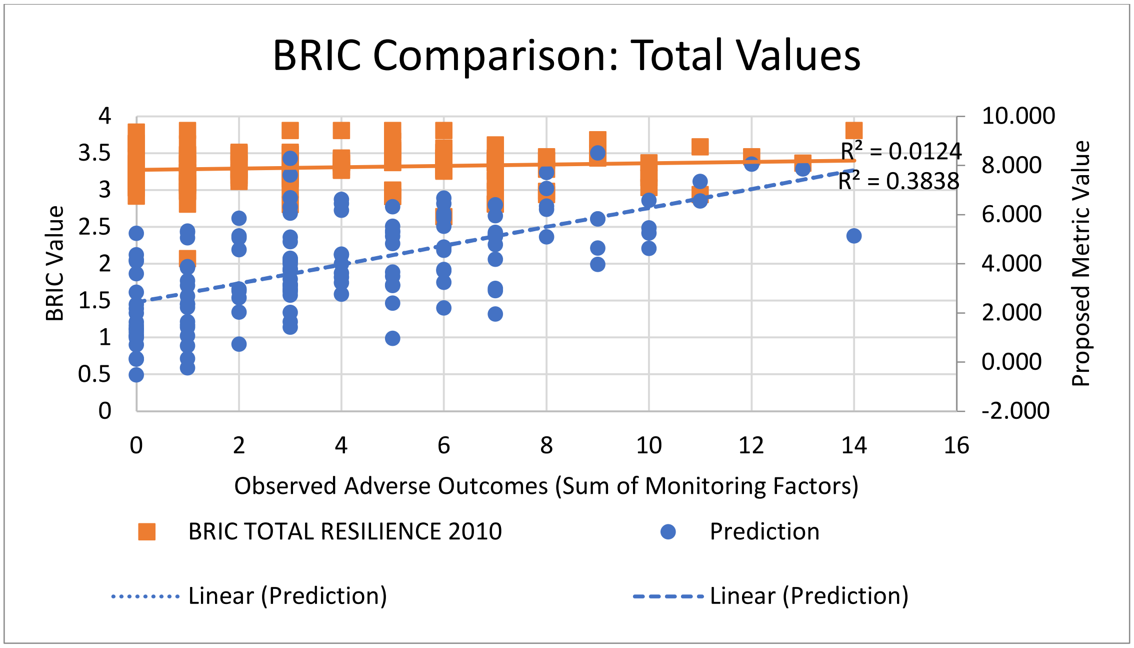

Figure 14: Comparison Predictive Power for Selected Monitoring Factors (BRIC Total)

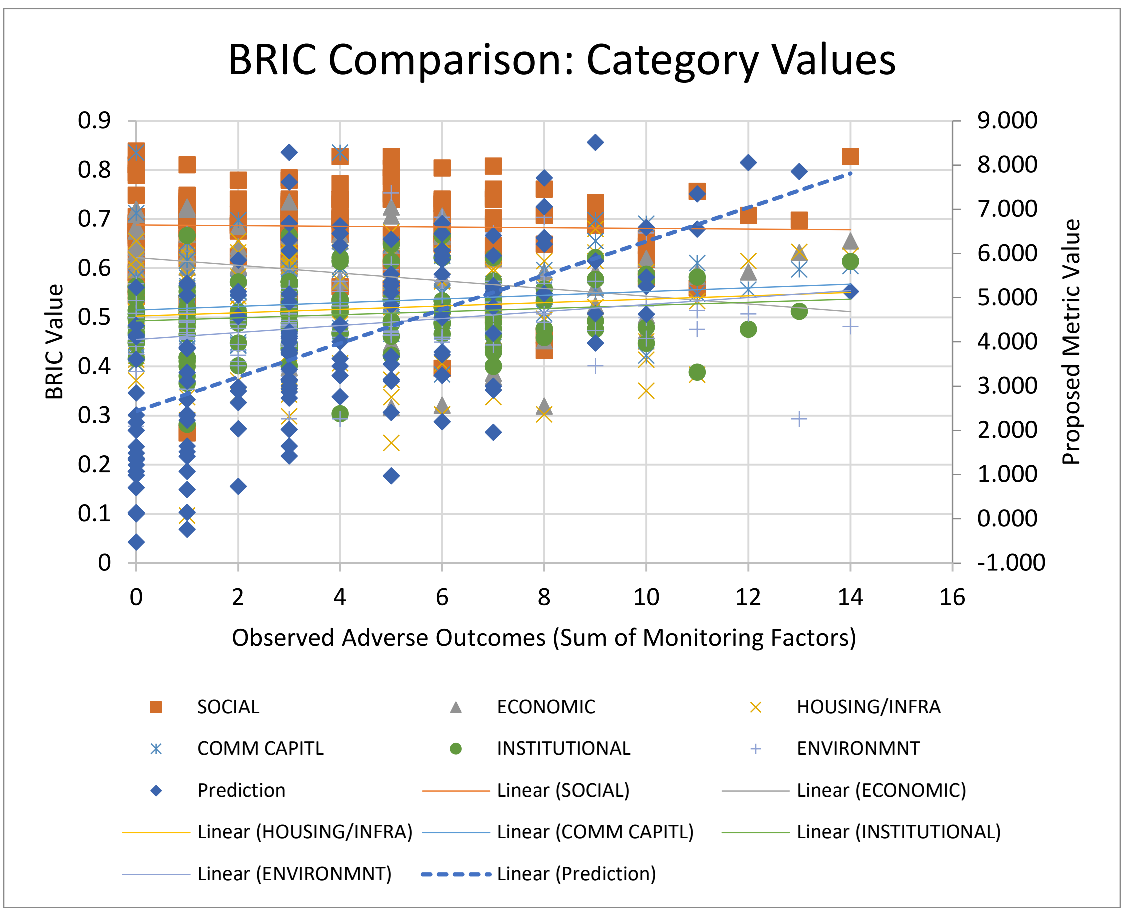

Figure 15: Comparison Predictive Power for Selected Monitoring Factors (BRIC Categories)

From this set of figures, it is evident that neither BRIC nor SoVI have the predictive power demonstrated by the proposed metric for the monitoring factors of interest. It is important to reiterate here the point of this comparison is not to invalidate or cast doubt on the usefulness of these metrics. They are backed by great caches of data and theoretical knowledge of the field. Instead, this comparison is included as a means of illustrating that the social factors best suited predict one outcome are not necessarily the best suited to predict any outcome. This brings the argument back to the central theme and driving force of this work, namely that communities should not be seen as vulnerable or not vulnerable but rather as a complex, dynamic mixture of strengths and vulnerabilities that will lead each community to have a tailored list of adverse outcomes to which it is susceptible.

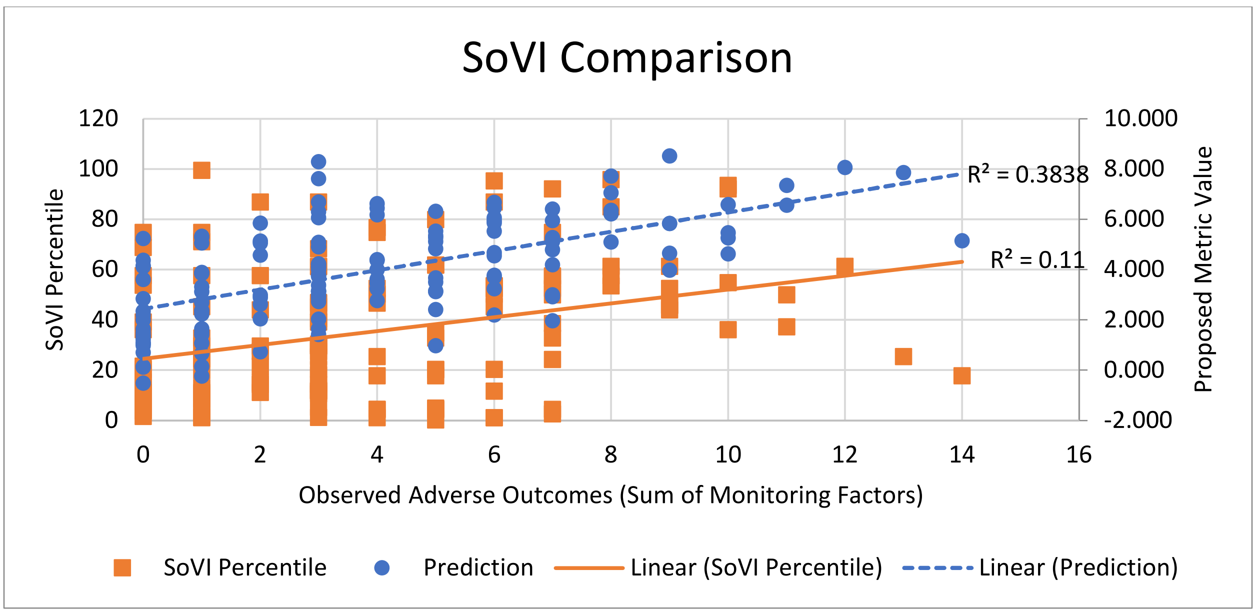

Figure 16: Comparison Predictive Power for Selected Monitoring Factors (SoVI)

2.4. Companion Tool for Limiting Community Selection Bias

Providing this SSM without any accompanying implementation tool would only partially address the issue of social characteristic consideration in post-event reconnaissance. As such, a simple and easy to use companion tool has been developed to streamline and standardize the decision-making process of community selection. The tool presented in Figure 17 prioritizes the selection of communities based on rarity of their social susceptibility status within the available pool of communities, the extent of damage (anticipated and/or reported), and the efficiency of traveling to each community. This framework is offered as the general guideline to be followed at every instance where the sample size needs to be reduced due to resource constraints while ensuring extent of damage and social susceptibility are systematically integrated in the decision-making process. A case study on the implementation of this metric and the companion decision-making tool can be found in the following chapter beginning at page 64.

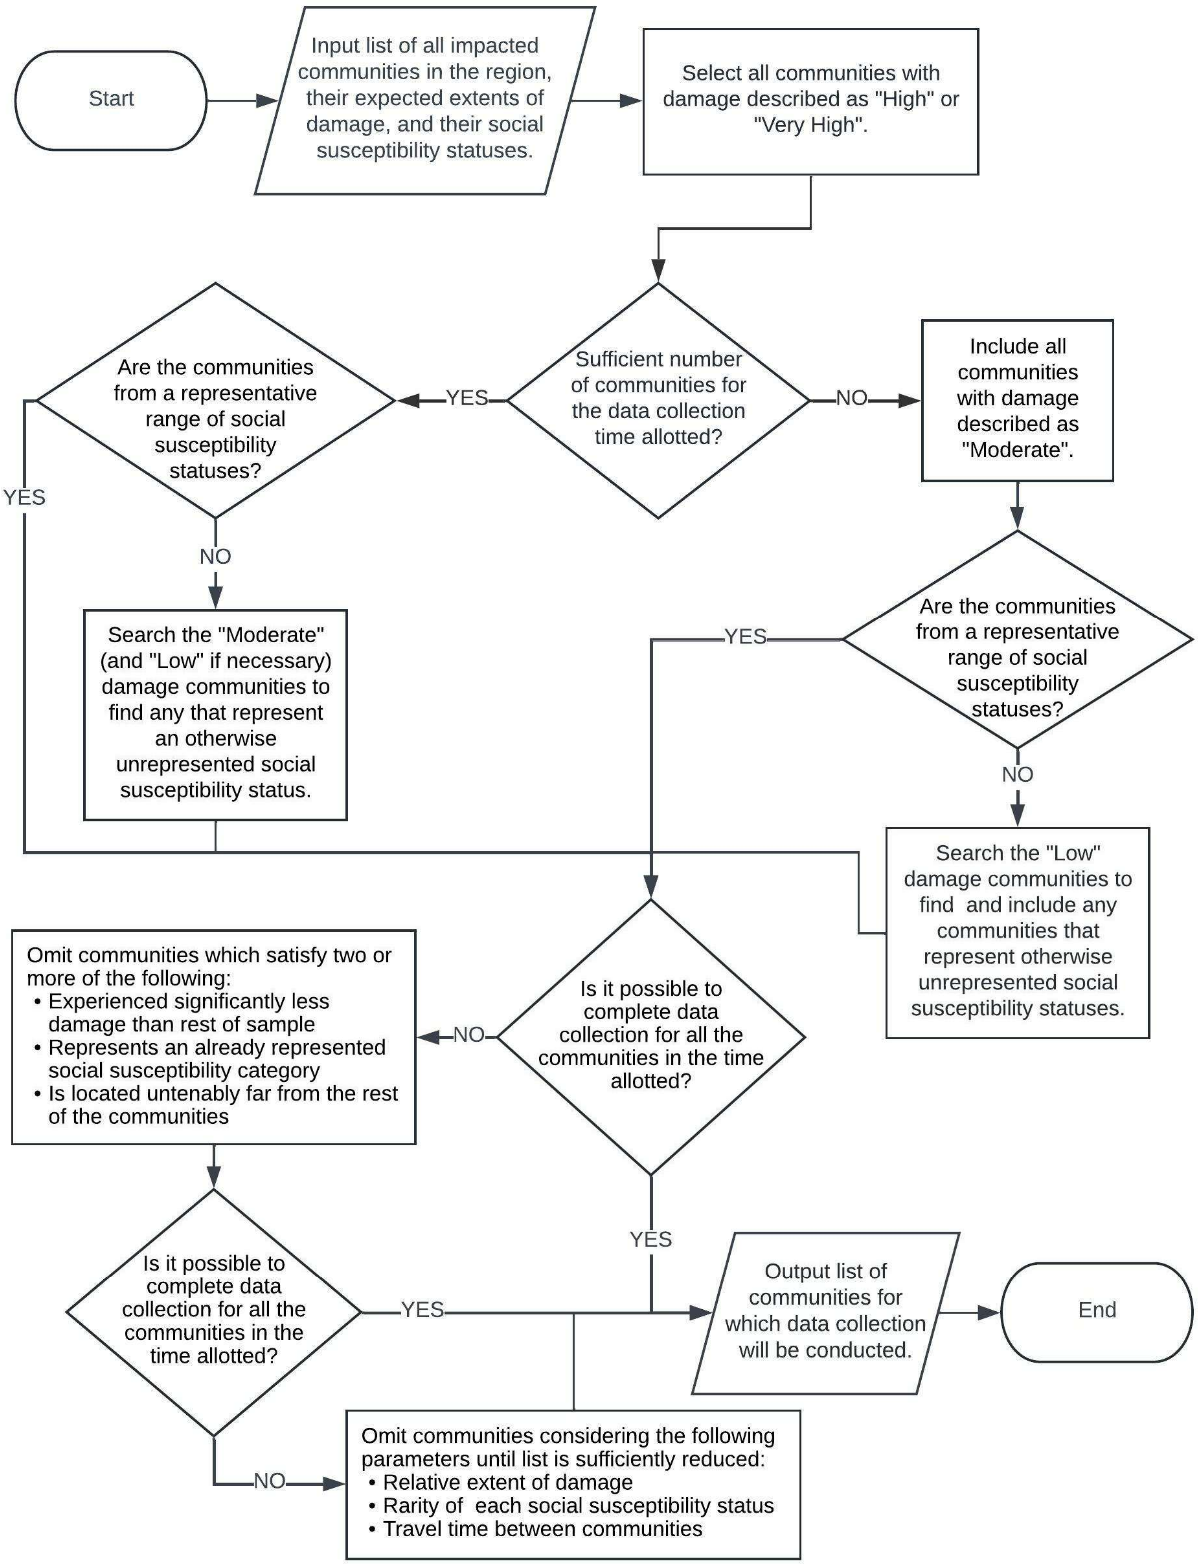

Figure 17: Protocol for Community Selection

2.5. Discussion And Next Steps

In reviewing the development and implementation of this metric and the associated decision-making tool, it is important to note an asset that has not yet been expressly mentioned, its generalizability. Some explanation has already been offered as to why the metric focusses on only whether fluctuations occur and not the extent of those fluctuations. In addition to previously referenced interest in providing a concise tool that is easily understandable and implementable, there is a second reason that speaks to the researchers' emphasis on the universal application of this tool regardless of community size or location, barring international use. While extent of impact is intentionally blurred in our monitoring factors, the duration of the impact is just as intentionally preserved. In selecting the parameters to be used in this metric as well as the combination of these parameters, an emphasis was placed on predicting the duration instead of the extent of negative trends as timelines are more easily comparable across dissimilar communities. As an example, it is more complex, and almost certainly inaccurate, to presume that a community of 100,000 and a community of 1,000 experience and perceive a 10% population decrease in the same way. Whereas 10 years of sustained population decrease can be assumed to have a more universal and reliable meaning, namely that the community is in decline regardless of its present size. Thus, the formulation of the monitoring factor measurements is not only intended to preserve the easily understandable metric schema but also to create a more generalizable, duration-based understanding of the community impact. The universal applicability of such a metric across not only different size communities but also different event types is an asset in disaster research within which the lack of generalizability of models and methods remains a concern (Johnston & van de Lindt, 2022). Thus, the monitoring technique put forth in this work is also viewed as a method that could be investigated and implemented separate from the rest of the framework in cases where comparisons across different types of hazard events can be difficult. Thus, this work has three deliverables that are used here in tandem, but each has the potential to be individually implemented. That is to say, the social susceptibility metric could be considered for implementation in applications that require a metric for predicting the long-term outcomes monitored in this work; the conceptualization of the community selection framework could be implemented with a different social vulnerability metric; and the monitoring protocol that emphasizes duration of impact could be used independently to compare across hazard types.

In discussing the results of the illustrative example presented above, it is apparent that the implementation of this metric for the decision-making tool maintains a certain level of inherent subjectivity. The final community selections made for the illustrative example show the options that must be weighed in this decision-making process, and they provide evidence of the fact that even with consideration given to the social susceptibility of each community, there are times when the sample set predicates that some social susceptibility tiers will be more thoroughly represented than others. This will inevitably be the case when the range of communities impacted is skewed. Nevertheless, an effort should be made to select communities that represent the full range of social susceptibility tiers present in the sample of considerably damaged communities. The results of this illustrative community selection process indicate that this tool certainly maintains space for interpretation and best professional judgement; however, it begins to put forth a methodology by which natural hazards research can incorporate a range of disciplinary perspectives from the earliest stages of planning and data collection.

With the scores for predicting recovery or decline hovering around 70%, this metric can be improved upon as more historical data becomes available. Increases in both number of communities and length of time for which those communities are monitored will aid in identifying emergent characteristics, further validating the parameterization, and where necessary modifying the parameterization or the weighting used. Even when data availability allows for patterns of recovery and decline to be more accurately known, such a tool is not intended to be implemented purely based on calculated values. This tool serves as somewhat of a boundary object, a place at which various fields overlap and can contribute. It is not intended to be implemented without input from an interdisciplinary team of researchers. In fact, it is meant to encourage discussion and deliberation within that team to understand the goals and perspectives of each team member. More clearly delineated means of orchestrating interdisciplinary collaboration are imperative for providing communities with holistic solutions in the preparation for and wake of hazard events. This integration needs to occur at every stage of data collection and analysis in order to create a pluralistic epistemology (Kendra & Nigg, 2014). The implementation of this metric and tool demonstrates that the rapidity of deployment in disaster research field studies does not preclude researchers from still engaging in interdisciplinary decision-making.

2.6. Conclusions and Future Work

It is important to note that there are many more lessons still to be learned in the development and implementation of this tool and its accompanying metric. The case study in this work is interesting in that it demonstrated implementation on a historic tornado outbreak especially when individual, short-track tornados likely would not impact enough structures to rationalize use of this tool. The case study is interesting to explore as we grow ever more concerned about the increasing impact of severe storms (A. B. Smith, 2020). However, there are limitations with using this event as the case study. Most notably, geographically larger events such as hurricanes and floods could have created a larger sample of communities to select from, leading to a more demographically representative group of communities being selected for a longitudinal study. Implementing this tool in a similar longitudinal study for a more geographically large event could inform the vulnerability process occurring as a result of the susceptibility metric described herein. Another area for continued ethical discussion is how recovery is quantified in this work and elsewhere. Future iterations of the metric may look to retrain the model with additional monitoring factors pertaining to diversity, equity, inclusion, and justice.

Future work may also include the exploration of the metric's implementation in events not tested and validated here as well as other spatial granularities. The events may include climatic events that do not yield structural damage such as droughts, non-physical hazards such as pandemics, and non-climatic hazard events such as earthquakes. Other spatial granularities may include census tract to align with other similar resources such as SVI, though Niche data may not be available at every census tract. The grounding questions provided by Wong-Parodi & Small (2021) for ensuring a decision-making tool is meaningful for intended users is an important resource to return to in the future work with this tool. Thus, this source will be used as a guiding document when refinements and expansions of the tool are undertaken.

As a final point on this topic, Bakkensen et al. (2017) not only spoke of the importance of empirical validation as mentioned earlier but also of the necessity for a metric to have a clear purpose. Speaking with greater clarity regarding the validation methods, limitations, and appropriate applications of a given social metric is an important step towards more actionable work in natural hazards research and beyond. To this end, the metadata template in Table 5 is offered as a brief semi-structured description of the metric provided herein and as a possible template for other work to provide clearer contextualization for development and implementation of such metrics to prevent their use in inappropriate and unvalidated applications. This table is inspired by the ODD (Overview Description Details) protocol in agent-based modeling (ABM), which originated from a similar need to standardize the characterization of models in that area of study (Railsback & Grimm, 2019, pp. 36-37). Indeed, many lessons may be learned in the work agent-based modeling has been able to accomplish in bridging many different disciplines with real-world overlap, and in the past few years, it has begun to be adopted as a boundary object in hazards research as well (Reilly et al., 2021). This first implementation of such a protocol in the context of social vulnerability metrics only explores implementation of the Overview section of the ODD as the other sections have details that are less applicable outside the ABM context.

Table 5: Social Susceptibility Metric Metadata

Purpose	Intended Audience:	Researchers
	Intended Use:	Rapid community selection for longitudinal data collection following geographically large events
	Stage of Disruption:	Recovery (long-term)
	Scale/Event Type:	Developed with data from various hazard events across the U.S. at the county and county subdivision level
Development	Parameter Consideration Method:	Theory Prior Social Metrics
	Parameter Selection Method:	Linear Regression (All Possible Combinations)
	Validation:	Empirical (long-term trends in census data for: Median Household Income, Population, Educational Attainment, Number of Households, and Number of Housing Units)
Data Requirements	Data Sources:	ACS 5-year Data Niche.com
	Inputs:	Median Household Income Population Area Niche User Rating Niche Overall Score Median Age National Averages of Median Age & Median Household Income
Results	Equation:	$y = 2.7902 - 3.2028in + 4.4495ag - 1.7423ru + 1.8580cs$
	Accuracy/Fit:	$R^2 = 0.303$ $p = 2.74 \times 10^{-9}$ High Vulnerability Accuracy: 70.31% Low Vulnerability Accuracy: 74.67%

The social susceptibility metric put forth in this method is intended to formalize the incorporation of social factors at every stage of data collection and analysis. It is suggested that the inclusion of interdisciplinary factors ensures that the research team will have a more holistic understanding of the challenges and assets a community possesses on the path to recovery following a hazard event. In providing an illustrative example as to how to utilize this metric in a standardized framework, this work aims to show that implementing interdisciplinary decision-making tools is still complex and nuanced, and regardless of the amount of guidance provided, there is still a great deal of professional judgment required to investigate the multi-faceted impacts of such events. Thus, it remains important to include an interdisciplinary and diverse team of researchers to achieve a multi-faceted understanding of communities and the hazard events they experience. This tool initializes the pathways to that shared understanding amid rapid field study preparations.

3. MIDWEST FIELD STUDY: DATA CREATION

3.1. Motivations and Overview

With the increased modelling capabilities of recent years, it is now possible to model not only community-level damage due to a hazard event but also the recovery trajectory of the building stock in the months that follow. The Interdependent Networked Community Resilience Modelling Environment (IN-CORE) allows users to perform damage and recovery analyses for a series of hazards. The development of this platform has been undertaken by an extensive team of researchers involved in the Center of Excellence for Risk-Based Community Resilience Planning (CoE) who work to ensure their models produce understandable and actionable outputs to support decision-making. As an element of actionability, the CoE must ensure that this suite of models has also been thoroughly validated. In the case of recovery modelling, this commonly requires that a model is empirically validated using hindcasting techniques based on field study data collected after a hazard event. This ensures that models provided within IN-CORE accurately characterize recovery patterns, allowing real communities to have reliable information on the likely bounds of predicted recovery following a hazard event.

As a continuation of this empirical validation of hazard models within the CoE, a longitudinal field study was launched following a tornado outbreak in December 2021 to document initial building damage as well as track the damaged buildings' recovery over a three-year period. Due to the geographically large nature of tornado outbreaks, it was necessary to implement techniques to rapidly decide which communities would be surveyed. Because it is well-documented in the literature that recovery depends in part upon social factors, it was determined that the research team would not only consider anticipated damage but also social susceptibility factors in selecting the communities to survey. Hence, the social susceptibility metric discussed in the previous chapter was utilized in this work for community selection. By choosing a group of communities for the longitudinal study that had different degrees of social susceptibility to natural hazards, the recovery models validated with this data could be shown to have predictive strength in a variety of contexts. The data processing and model validation process is still underway for this project as the data collection does not conclude until December 2024; however, results from the first six waves of data have been finalized and are presented in this chapter. The results are compared with the SSM by a rank order and social susceptibility tiers presented at the end of this chapter.

The larger goals of this work were to provide a robust, publicly available dataset to the field of hazards research, provide empirical validation to models within IN-CORE, and provide a set of in-depth investigations as part of this field study investigation to move the field forward. As part of the overlap of IN-CORE and the aims of this dissertation, a field study dataset has been developed that (a) builds a repository of reliable data and (b) determines the SSM's efficacy in predicting which communities would be able to recover from the profound impact of this event in the weeks, months, and years after the event due to not only the level of damage in each community but also the social and economic resources to which each community had access. In establishing a longitudinal study of an event that simultaneously impacted multiple communities, it was proposed that this work could provide greater insight into the mechanisms that aid in the recovery of communities after such events and conversely the attributes that diminish a community's ability to recover.

3.1.1. Event Description

On December 10-11, 2021, a violent outbreak of 66 confirmed tornadoes in nine states across the U.S. caused $4.1 billion of damage, 87 deaths, and injuries in the thousands (National Weather Service, 2021, 2022; Paducah National Weather Service, 2021). This deadly tornado outbreak began the evening of December 10, 2021 and continued into the early morning hours of the following day. In total, this event consisted of 66 confirmed tornadoes of which two were EF4 and six were EF3 (National Weather Service, 2021). The full geographic scope of this event is shown in Figure 18. The most destructive EF4, which had a staggering path length of 165.7 miles, resulted in 57 deaths and over 500 injuries. This long-track tornado caused extensive damage to several communities. The most significant damage from this tornado occurred in Mayfield, Kentucky. It caused the damage or destruction of 3,778 residential buildings, 183 commercial buildings, and 103 other buildings. This damage included the loss of Mayfield's historic downtown area, the county courthouse, and a candle factory to the southwest of Mayfield (Paducah National Weather Service, 2021). As such, Mayfield was the focus of this longitudinal field study and community selections were made under the assumption that Mayfield would be a cornerstone of the investigation.

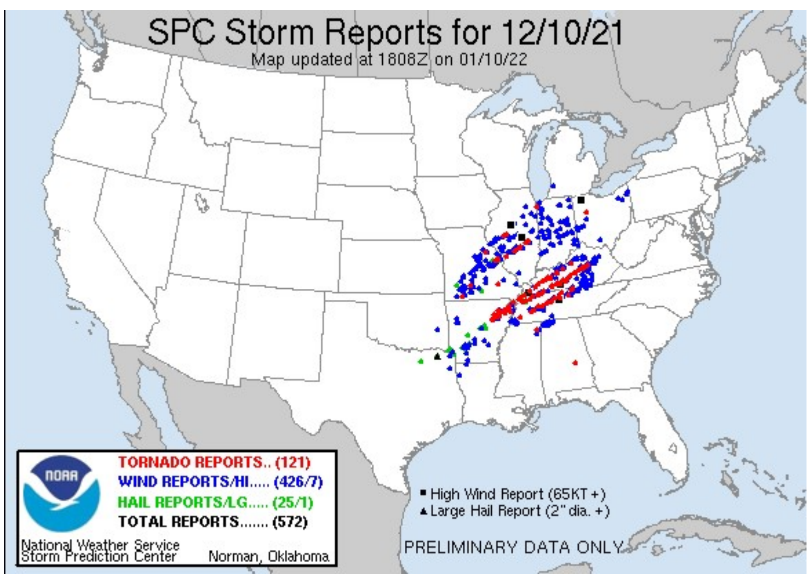

Figure 18: Tornado Outbreak on December 10-11, 2021 (National Weather Service, 2022)

3.2. Methodology

3.2.1. Social Susceptibility Metric (SSM) Implementation

The SSM framework provided in the previous chapter as a means of community selection in geographically large events was first implemented in this field study. In the days immediately following the event, news and aerial data were surveyed to establish which communities had been affected, and where possible, qualitative descriptions of anticipated damage were also recorded. All the communities noted in the preliminary damage search were then evaluated for their overall susceptibility score. These susceptibility scores were put into the categories described above, namely Markedly High Susceptibility, High Susceptibility, Low Susceptibility, and Markedly Low Susceptibility. Due to the low data needs of the metric, this data was collected rapidly from the U.S. Census website. Once each community had been placed into each corresponding bin, the communities were then placed into tiers of importance for data collection in the field study. These tiers were established by weighing the suspected extent of damage, the rarity of the community's susceptibility category within the sample of communities, and the efficiency of resource use as outlined in Figure 17. The rarity of the community's susceptibility category was prioritized in this case to ensure that communities that represent unique categories in the sample set will be monitored throughout the longitudinal study. Prioritizing rarity was the prerogative chosen for this work in order to answer research questions outside the scope of this work, but the research focus could just as easily have been placed on communities with very high social susceptibility scores exclusively rather than maintaining a representative sample. The communities considered and the tiers demarcated are included in Table 6 and sorted by priority. The field study team utilized the proposed tiers of importance to collect data in the form of 360-degree geotagged videos on

December 22-23, 2021. The resultant list of communities for which data was collected in this field study is shown in Table 7.

Table 6: Selected communities grouped by priority of data collection

Town of Interest	County/ County Subdivision	Social Susceptibility Metric	Social Susceptibility Tier	Priority Tier
Leachville, AR	Neal township, Mississippi County	0.89	Very Low Social Susceptibility	1
Bowling Green, KY	Bowling Green CCD, Warren County	1.37	Very Low Social Susceptibility	2
Edwardsville, IL	Edwardsville township, Madison County	1.59	Very Low Social Susceptibility	1
Samburg, TN	District 5, Obion County	3.47	Low Social Susceptibility	2
Mayfield, KY	Graves County	5.10	High Social Susceptibility	1
Bremen, KY	Bremen CCD, Muhlenburg County	5.39	Very High Social Susceptibility	2
Mayfield, KY	Mayfield CCD, Graves County	5.85	Very High Social Susceptibility	1
Monette, AR	Buffalo township, Craighead County	6.08	Very High Social Susceptibility	1
Trumman,AR	Willis township, Poinsett County	6.32	Very High Social Susceptibility	2
Centertown, KY	Centertown CCD, Ohio County	6.96	Very High Social Susceptibility	3
Hartford, KY	Hartford CCD, Ohio County	8.27	Very High Social Susceptibility	2
Dawson Springs, KY	Dawson Springs CCD, Hopkins County	8.45	Very High Social Susceptibility	2
Princeton, KY	Princeton CCD, Caldwell County	8.78	Very High Social Susceptibility	3

Table 7: List of Communities for which data was collected during field study

Town of Interest	Social Susceptibility Metric	Social Susceptibility Tier	Predicted Outcome for Monitoring Factors Based on Social Susceptibility
Leachville, AR	0.89	Very Low Social Susceptibility	Marked Stability
Edwardsville, IL	1.59	Very Low Social Susceptibility	Marked Stability
Samburg, TN	3.47	Low Social Susceptibility	Stability
Mayfield, KY (Graves County)	5.10	High Social Susceptibility	Decline
Bremen, KY	5.39	Very High Social Susceptibility	Markedly Decline
Mayfield, KY	5.85	Very High Social Susceptibility	Markedly Decline
Monette, AR	6.08	Very High Social Susceptibility	Markedly Decline
Centertown, KY	6.96	Very High Social Susceptibility	Markedly Decline

Hartford, KY	8.27	Very High Social Susceptibility	Markedly Decline
Dawson Springs, KY	8.45	Very High Social Susceptibility	Markedly Decline

For sustainability of the longitudinal study, it was decided that the list of communities to be analyzed in the later data collection waves should be further reduced. This was accomplished by first asking all participating researchers to provide a brief description of the extent of damage seen in each community for wave one. This description was requested as a means of incorporating professional judgement and getting a broader sense of the damage for which the video data was still in the preliminary stages of processing. Once these assessments were provided, the communities for the longitudinal study were selected in accordance with Figure 17 from the previous chapter.

The results of this framework being applied are shown in Figure 19 with color coding based on social susceptibility tiers. A map of all surveyed communities is shown in the left image of Figure 19 with very high social susceptibility shown in red, high social susceptibility shown in orange, low social susceptibility shown in green, and very low social susceptibility shown in blue. The selected set of communities is shown in the right image of Figure 19 and listed in Table 8. Mayfield in the final community selection has been denoted by dark orange to denote that although Mayfield has high social susceptibility, its surrounding county has only high social susceptibility. When comparing Figure 19 to the idealized community selection process shown at the beginning of this chapter, it can be observed that this sampling is notably more skewed to the very high socially susceptible communities. This is because communities more directly in the path of the tornadoes in this outbreak happened to be in regions of greater social susceptibility. The grouping of communities used in the longitudinal study is thus representative of this high social susceptibility, with only one very low social susceptibility instance in Leachville and one low social susceptibility instance in Samburg. It is worth noting an interesting contextualizing detail here as well. Although Mayfield ranked in the very high social susceptibility category, as seen in Table 8, Graves County, the county in which Mayfield resides and for which Mayfield is the county seat, fell within the bounds of the high social susceptibility category. Thus, the inclusion of Mayfield could provide greater insight as to the role the greater geographic context plays in recovery. Hence, Mayfield has been denoted as dark orange rather than just red in the right image of Figure 19.

Table 8: Communities selected for longitudinal study

Town of Interest	Social Susceptibility Metric	Social Susceptibility Tier	Community Monitoring Factor Summary Statistics Tier
Leachville, AR	0.895	Very Low Social Susceptibility	Marked Stability
Samburg, TN	3.470	Low Social Susceptibility	Stability
Mayfield, KY (Graves County)	5.101	High Social Susceptibility	Decline
Bremen, KY	5.395	Very High Social Susceptibility	Marked Decline
Mayfield, KY	5.847	Very High Social Susceptibility	Marked Decline
Monette, AR	6.081	Very High Social Susceptibility	Marked Decline
Dawson Springs, KY	8.453	Very High Social Susceptibility	Marked Decline

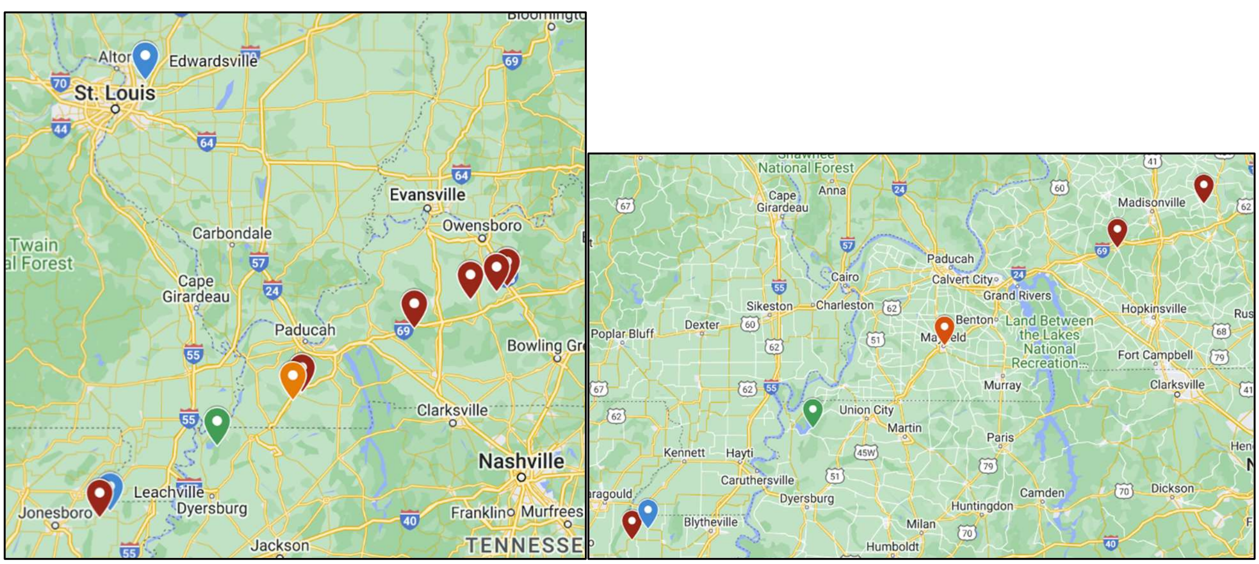

Figure 19: Left: Communities initially surveyed, Right: Communities selected for longitudinal data collection

3.2.2. Data Collection

3.2.2.1. Wave 1

The author has contributed to the collection, processing, and evaluation of various waves of data for this field study as part of a larger team of researchers. The full documentation of the study is outlined in van de Lindt et al. (in review). The initial wave of data collection was conducted by dividing eleven researchers into four teams, each of which had a vehicle. Each team was assigned a community or multiple communities to survey on the first day of data collection with the plan to adjust as necessary after the first day due to some unpredictability in both the road conditions and the equipment performance. Three teams started from Memphis, Tennessee while the fourth team started from St. Louis, Missouri to collect data in Edwardsville, Illinois. The teams convened in Murray, KY after the first day of data collection in order to decide on a plan that optimized the data collection process for the next day to survey as many prioritized communities (see Table 6) as possible. This plan was selected in accordance with the protocol outlined in Figure 17. The teams met in Louisville, KY the third night with all teams able to complete the plans set out for them the night before. This concluded the first wave of data collection for this field study.

This relatively small team was able to perform extensive data collection over just two days of surveying due to the technology employed. Whereas previous work involved on-site damage assessment, taking single photos, and surveying the damage on foot, this field study used 360-degree GoPro cameras mounted to the vehicle roofs to collect data as the team drove through each community. This process adhered to the methodology put forth in Crawford et al. (2017). For areas where the road was obstructed due to debris or barricades, the team supplemented video data with still photos using the QuickCapture application to generate images. A GPS unit was also used to track the GoPro video footage; however, later analysis found that the GPS tracks generated by the

GoPro cameras were sufficiently accurate for location tracking. An example of the 360-degree video generated by the GoPro camera in Mayfield, Kentucky is shown in Figure 20.

Figure 20. Example of 360-degree video coverage

3.2.2.2. WAVES 2-6

For the longitudinal study, it was determined that data would be collected over the course of six waves and two years. Upon review of the recovery data, this timeline was extended to include a seventh wave of data collection for the communities which required additional time to document recovery, the results of which will not be included in this dissertation due to timing conflicts. Subsequent waves of data collection were streamlined further by removing three communities that saw minimal damage in the first wave of data collection. These omitted communities are Centertown, Edwardsville, and Hartford. The communities appearing in every wave of data collection are shown in Table 8. This reduction in scope along with the team becoming more accustomed to the data collection process allowed most later waves to be accomplished by a team of only four people in two cars over the course of two days. The general routes of these later waves are shown in Figure 21 and Figure 22. For some waves, more researchers were able to join and/or less daylight was available due to the time of year and, in these situations, Car 1's tasks were divided such that Car 1A surveyed everything South of Mayfield and to have Car 1B survey everything North of Mayfield. After each wave of data collection, the data was catalogued in a shared google drive before undergoing further processing, which will be detailed in a later section.

Table 9: Timeline of Data Collection Waves

Wave	Time Since Event	Date
1	3 weeks	December 2021
2	3 months	March 2022
3	6 months	June 2022
4	1 year	December 2022
5	18 months	June 2023
6	2 years	Jan 2024
7	3 years	December 2024

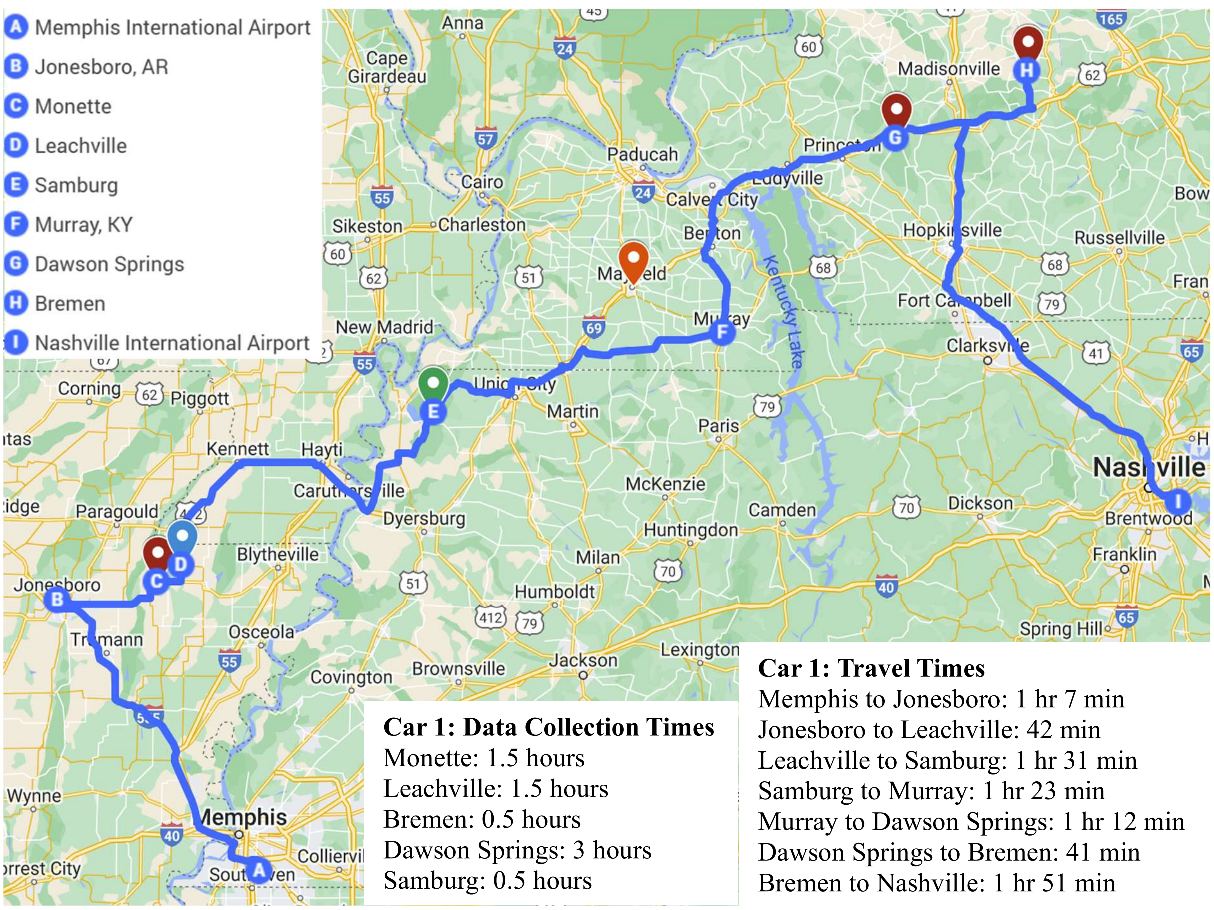

Figure 21:Standard Route for Car 1 in Waves 2-6 to reach all communities aside from Mayfield

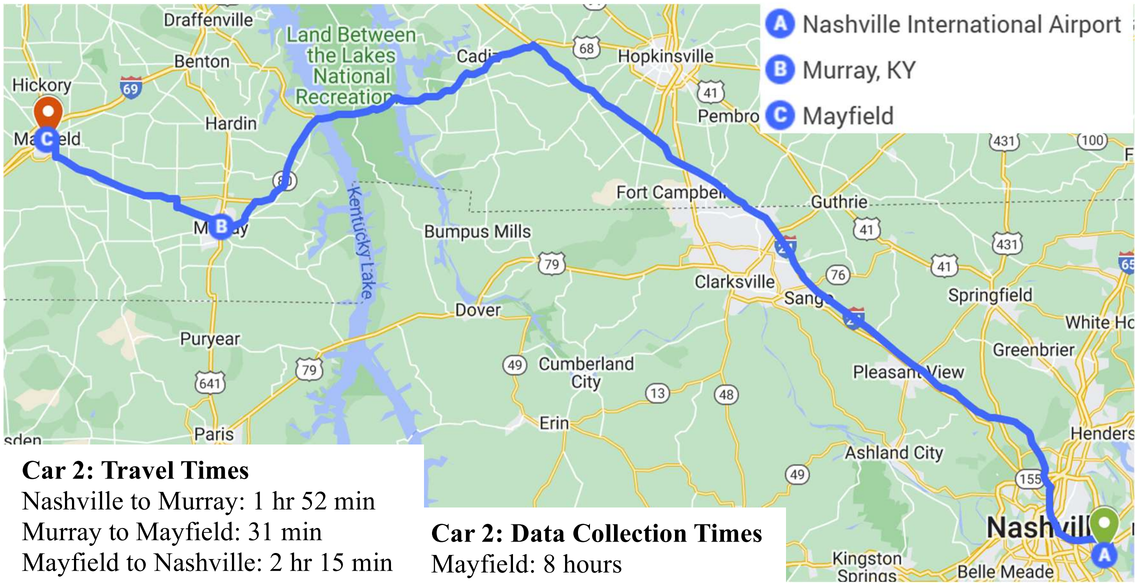

Figure 22: Standard Route for Car 2 in Waves 2-6 to focus on the largest single data collection effort of Mayfield

3.2.2.3. Wave 7

Upon review of the recovery trends for the first six waves, it was decided that Wave 7 would only survey the two slowest communities to recover, Dawson Springs and Mayfield. More information on these communities and their building stock trends will be shown at the end of this chapter, indicating that by two years after the event these communities still had not seen a return of 90% of their building stock. Thus, they were the primary focus of the seventh wave, which will not be detailed here but will be made available via data publication in the coming months.

3.2.3. Data Processing

3.2.3.1. Prerequisite Video Data Processing

As part of the tradeoff inherent in reducing the time required in the field for data collection, the post processing required in creating meaningful datasets from the raw video footage has been extensive. The first task was to convert the video format from proprietary GoPro files to the standard non-proprietary mp4 format to adhere to dataset management bast practices and to promote inter-platform operability in subsequent steps. Then there was the unique method by which GoPro cameras chapter their videos approximately every 8 minutes in order to prevent the loss of an entire video due to corruption of a small section of footage. For ease of data annotation at later steps and to streamline the handling of these video files, the 8-minute videos were concatenated to have a length of approximately 32 minutes (4 videos concatenated). This had a fair amount of fluctuation as data collectors would restart or pause video recording intermittently for any number of reasons. The resultant files were named following a human readable convention that indicated the town and the chronological order of the video files. For instance, "Leachville4.mp4" is the fourth concatenated video captured in the town of Leachville, Arkansas for a given wave. Folder hierarchies distinguish the wave in which a video was collected so this has not been included in the naming convention to avoid confusion between wave and video numbers.

As the 360-video data was converted and concatenated, so to the accompanying LRV data had to been converted to the non-proprietary csv format and stitched together to give video timestamps and coordinates of the camera at each instant. With all pre-requisite elements generated, the georeferenced video files then had to be converted into a map of dynamic survey points that could be annotated for damage, recovery, and a host of other attributes. This work also needed to be able to be done simultaneously and remotely by a team of remote student workers. To meet these requirements, the CoE established a collaboration with the Center for Advanced Public Safety (CAPS) at The University of Alabama. This group already had a webtool platform called Extreme Events that has been used by other researchers to visualize, interact with, and augment geospatial data. Although this CAPS platform had been developed previously, this project's need for customized functionalities as well as some deprecations in the platform since previous implementations, required troubleshooting and weekly meetings with the CAPS team to ensure the undergraduate student data annotators could continuously clean, annotate, and revise the dataset as easily as possible. A screenshot of the CAPS interface is shown in Figure 23. The building centroids as well as the extracted images could be tagged in a variety of customizable ways to meet the investigatory needs of the researchers involved in the project. The blue pentagons represent buildings, and the purple points represent survey locations at which still images were extracted from the 360-degree video. By toggling the Wave drop down, these purple points would filter to show just the survey locations for the selected wave.

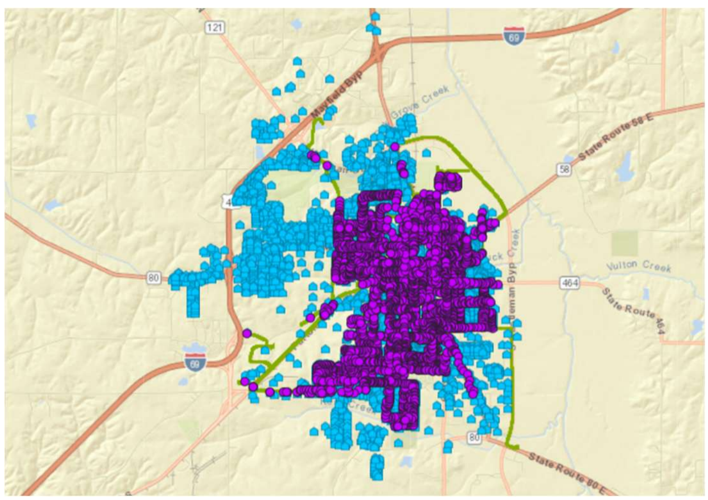

Figure 23: Screenshot of Mayfield wave 1 data from the Extreme Events website

3.2.3.2. Student Training and Mentorship

To manually annotate such an expansive dataset of over 5,000 buildings and 30,000 corresponding survey points, a team of over a dozen undergraduate students was brought on to the project alluded to previously. These students required training in damage assessment protocols as well as building components that are commonly referenced in damage assessment protocols but are absent from the vernacular, e.g. “parapet walls.” Thus, they received a set of virtual, asynchronous training sessions followed by checks for understanding. This was of course just the beginning of the training, however. In the subsequent months, the students were supervised for the first 150+ data annotation sessions to ensure that all questions were quickly addressed and that the students were correctly interpreting the description in the IN-CORE and EF Scale documentation. To further build on the validity of these damage assessments, students tagged in ever-changing teams of no less than three to ensure that all assessments provided were reached by consensus and not just the opinion of a single individual or even the potential echoing chamber if students had been assigned to fixed teams and hours.

As an intentional though supplemental benefit to this technique of tagging by randomized consensus, this field study could provide busy students with a flexible and engaging work environment. Firstly to this end, the research team prioritized their agency by allowing them to select their own hours, making it clear that confirmed annotation sessions that they had to cancel were hours that their fellow taggers couldn't work. By being more accountable to their peers rather than a boss or manager, students only canceled when absolutely necessary and often found replacements for themselves when they did have to cancel at the last minute. Second, students were encouraged to bring their whole selves to the research work through the way annotation sessions were managed, and relationships were fostered. The results of this engagement meant that students felt comfortable and even compelled to ask any and all questions, recommend website interactivity modifications, and suggest the addition of new tags for emergent characteristics in the dataset. This not only made the research better, but it also allowed students to socialize virtually with their peers in a freshly post-COVID world where loneliness was having a noticeable impact on student well-being (Richardson et al., 2024). Although just a by-product of this work, it has been meaningful to see these students form impactful friendships, have movie nights, and even race in triathlons together. I could not be more grateful to this group of students who worked so tirelessly to ensure the results produced were as accurate as possible and did so while cultivating so much joy and camaraderie.

3.2.3.3. Data Annotation

To begin this work, the team started with a core set of data tags which the team was focused on tracking over time. These tags were: Damage State from IN-CORE, Building Archetype from IN-CORE, Damage Indicators from the EF Scale, Degree of Damage from the EF Scale, and Recovery State. Over the course of data tagging, other tags were added to increase the complexity and richness of the dataset based on requests from fellow researchers or the student data annotators. These secondary tags were helpful in creating a more nuanced narrative around a given building, but the primary tags listed previously remained the focus of the work and were thus cleaned and validated more systematically than the secondary tags.

Each data annotation session involved a team of three or four students using the Extreme Events website provided by CAPS to navigate the communities and investigate each building for which there was sufficient, high-quality video footage to assess damage and recovery. Where necessary the students used resources such as Google Earth, Google Streetview, and other waves of data to clarify any inconsistencies and uncertainties. Some such anomalies included discerning post-event from pre-event damage, verifying total structure destruction in high damage areas, and determining building archetypes and damage indicators for high damage structures. These supplemental datasets were invaluable in cleaning the data for particularly perplexing structures, and the students were industrious in their ability to use this amalgam of resources to create as clear of a picture of reality as possible.

As each wave of data was completed, the team received data exports of the results from the CAPS team. These results were reviewed for missing damage or recovery tags, and the points were flagged for the students on an interactive and editable ArcGIS Web Experience to allow them to note points of concern, review those points, resolve any errors, and then remove the error flag from the ArcGIS web layer. This process was completed for inconsistencies within each wave of data as well as across all the waves of data to create the most complete and refined dataset possible. Throughout this process the team received additional exports of the CAPS data across the waves, and when necessary, the data cleaning process was repeated until any data errors constituted less than 1% of each wave's dataset. As of the writing of this dissertation, the students have annotated, reviewed, and satisfactorily cleaned all data for the first six waves. Wave 7 is currently being annotated and cleaned.

3.3. Results

3.3.1. Overview

Although the potential research questions answered by this study are expansive and aim to propel the field forward more generally, the driving goal of this study as a whole has been to provide empirical validation to IN-CORE models. The full task of model validation is outside the scope of this study; however, this work is under way and this field study investigation offers a meaningful contribution to the task of ground-truthing IN-CORE models (W. "Lisa" Wang et al., 2024). The most integral tags that data annotators have assigned to buildings for initial damage assessment were IN-CORE Building Archetype, IN-CORE Damage State, Damage Indicator, and Degree of Damage. The assignment of the Damage Indicator and Degree of Damage tags allows an EF rating to be assigned for each region so that a model tornado used in IN-CORE can be validated. Then the simulated damage can be validated by comparing the model outputs Damage State assignment with the observed Damage State assignment for each building. These results demonstrate the actionability of IN-CORE and its suite of analyses. As this larger model validation work continues, the primary deliverables for this work within the scope of this dissertation are to first provide this data in an accessible way and second to provide a case study to test the results of the social susceptibility metric generated earlier. To the first task, this data is being made available on the platform DesignSafe-CI and will have a doi forthcoming. To the second point, the first two years of damage and recovery data have been collated, reviewed, and presented here to discuss the predictive strength of this SSM and the lessons learned in its first application.

3.3.2. Discussion of General Trends

A brief review of the data and some key findings are summarized here to rationalize the contribution of this study. To summarize the collected data and building stock recovery trends across the six communities for the first two years of the study, Figure 24 has been included below. This figure shows that of the approximately 5,000 buildings surveyed across all waves, nearly 3,500 were undamaged after the event according to the thresholds outlined in the IN-CORE Damage State documentation for tornadoes. These undamaged structures are truncated out of the figure to focus on building stock damage and recovery trends. Of the 1500 buildings impacted by the event, approximately 400 reached a Damage State 4 or were immediately cleared, suggesting that damage was extreme and irreparable. For the remaining approximately 1100 buildings at varied levels of damage, over 200 reached a Damage State 3, approximately 400 were at a Damage State 2, and more than 500 were noted as a Damage State 1.

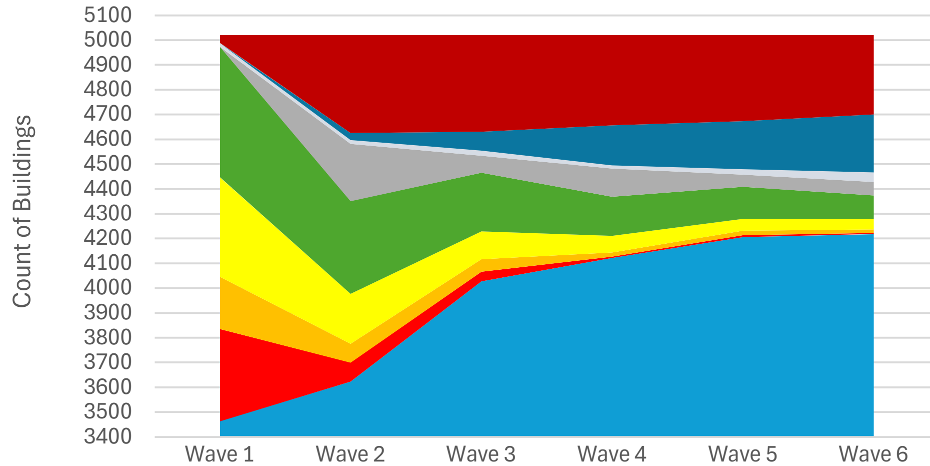

	Wave 1	Wave 2	Wave 3	Wave 4	Wave 5	Wave 6
Cleared & Not Replaced	33	395	391	364	348	320
Cleared & Replaced	0	29	75	162	194	234
Null	14	15	21	13	22	39
Not Surveyed	0	231	69	114	49	54
DS 1 (Slight)	527	374	236	157	129	96
DS 2 (Moderate)	401	201	113	68	47	41
DS 3 (Extensive)	211	76	50	16	19	14
DS 4 (Complete)	371	77	39	6	7	5
No Damage	3464	3623	4027	4121	4206	4218

Figure 24: Summary of all tracked buildings across six waves and six communities

In subsequent waves, the post-event pathways seen in the dataset are broken into four primary categories: repair, clearing and reconstruction, clearing only, and inaction. The first two of these methods contribute to positive trends for each community's building stock, while the latter two do not. Certainly, the clearing away of a severely damaged building is better than inaction if for no other reason than human safety, and it readies the site for possible future construction. However, the data has shown that it does not indicate rapid reconstruction is to follow. Indeed, relative to the 424 buildings cleared by 3 months after the event, only 235 new structures had been built two years after the event to recoup these building stock losses. This reconstruction can be slower due to several factors from funding to permitting, as such clearing is only a preferred option for restoration when repairs are impossible or highly impractical due to the extent of the damage. As for repairs, the portion of owners that chose this path versus clearing or inaction are shown in

Figure 25. A consistent band of about 150 buildings were in the process of being repaired in each wave except for Wave 6, growing the number of fully repaired buildings with each passing wave. Wave 6 indicates the beginning of this repair cycle slowing as those buildings which are going to be repaired are dwindling and the number of buildings for which it seems no action is going to be taken steadies and even grows slightly between Wave 5 and Wave 6 as property owners either abandon their building or determine they do not have the capacity, economic or otherwise, to perform repairs. The timeline comparison between repaired and reconstructed buildings indicates that while repairs were underway right after the event and begin to plateau 2 years after the event, reconstructions only start in earnest at the 6-month mark and continue relatively consistently after that point. For the sake of concision, it can be expressly stated here that Waves 2-6 saw 29, 67, 120, 74, and 59 newly completed or under-construction buildings, chronologically. These figures are partially included in the "Cleared & Replaced" element in Figure 24, where completed constructions are summed wave after wave and the under-construction structures are only counted in the wave within which they were noted. This trend of slower reconstruction versus repair is not new, but the clarity with which this data proves this point is useful.

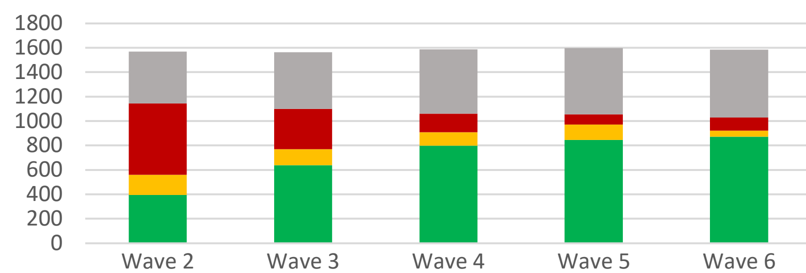

	Wave 2	Wave 3	Wave 4	Wave 5	Wave 6
Cleared Buildings	424	466	526	542	554
RS3: No Repairs Underway	583	329	151	85	108
RS2: Repairs Underway	164	131	110	126	50
RS1: Fully Repaired	396	638	799	844	871

RS1: Fully Repaired RS2: Repairs Underway RS3: No Repairs Underway Cleared Buildings

Figure 25: Repair States for all six communities across waves 2-6

The restoration pathways are also broken down by initial damage state in Figure 26. This figure shows the timeline for when different damage state buildings were cleared as well as the repair state of uncleared buildings in Wave 6. Comparing the stacked sum for each category to the totals shown in the initial damage state column of Figure 24, the values are demonstrated to be identical. This first and foremost demonstrates the reliability of the data set while also reflecting the strength of longitudinal field studies in their ability to track through time the status of nearly every building in a study area to see their recovery path post-event and their more long-term outcome. This figure shows a separation of restoration pathways with lower-level damage often being rectified via repairs and higher-level damage being addressed with reconstruction. This aligns with the conversation above as well. It can also be noted that the vast majority of structural clearing occurred in Wave 2, which is reflected in other figures as well but is perhaps most clearly demonstrated in this representation.

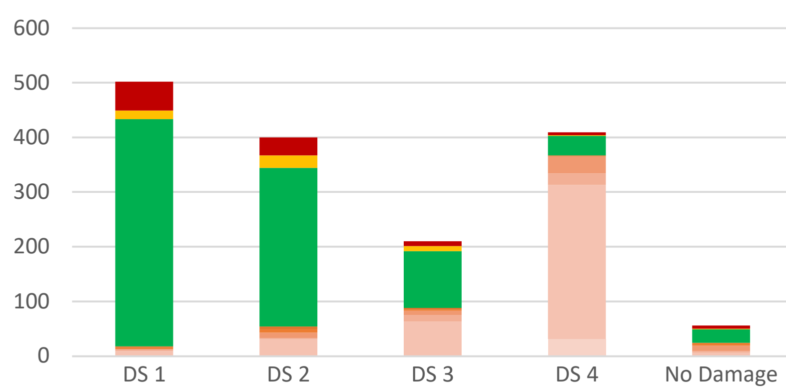

	DS 1	DS 2	DS 3	DS 4	No Damage
Wave 6 No Repairs	53	33	9	5	6
Wave 6 Repairs Underway	16	23	9	1	1
Wave 6 Repaired	415	290	104	36	25
Cleared in Wave 6	2	5	3	2	0
Cleared in Wave 5	3	6	2	0	5
Cleared in Wave 4	1	10	8	31	10
Cleared in Wave 3	3	2	12	21	3
Cleared in Wave 2	9	31	63	282	6
Cleared in Wave 1	0	0	0	31	0

Figure 26: Status of buildings grouped by Wave 1 Damage State

This representation of the data also highlights a few key features that might be worth discussing or at least clarifying. First, some buildings with a damage assessment of No Damage were repaired. This is simply because the damage states defined in IN-CORE do not capture all possible permutations of damage. Thus, there were some instances where a structure could not be characterized as damaged but was indeed in need of minor repairs. Second and more impactfully, it is evident that some buildings which had little to no damage were ultimately cleared away. There are three primary reasons that this might occur based on lessons learned in this study. First, although industrious, the student annotators cannot assess what they cannot see, so it is possible that there was more severe damage not captured by any of the accessible resources and footage. This may have led to artificially low damage assessments for some structures. Second, there appear to be some instances in the data where a building with minimal damage was cleared because a large swath of its neighboring buildings was cleared due to much higher damage. And in razing the entire block, the owners of these remaining structures allowed their home to be removed as well. The exact reasons for this can only be speculated upon. And finally, for this event, there was also a large winter storm that came through much of this region less than a month after the tornados, leading to frigid temperatures and 4-6 inches of snow for some of the impacted communities including Mayfield, Dawson Springs, and Bremen (Paducah National Weather Service, 2022). For homes whose envelope had been breached by the tornadic wind and had not been sealed off from the elements yet through repairs or tarping, this winter storm may have led to significant water intrusion. Alternatively, buildings whose heating source was non-functional at the time of the storm may have faced the damage associated with burst pipes. Thus, some of the lower damage cleared buildings may have withstood additional damage after Wave 1 but before the team returned for Wave 2 that led the structure to be demolished. To increase the rapidity of building stock recovery, these low damage clearings should be avoided wherever possible. As weather patterns become more volatile, this may require more rapid deployment of resources to impacted communities to shore up buildings before a second event follows in relatively rapid succession as was the case here. Deferred repairs that are due to delays in resource allocation are avoidable losses and should be prevented wherever possible.

Aside from an interest in general trends across the area surveyed, there was also an interest, as stated previously, in examining the building stock recovery trajectories for each community. These trajectories will be presented in the following section in and compared to the SSM values that were initially used to select the communities surveyed in this study. This represented the first implementation of the SSM for community selection and the goal was to ensure that the results of this work captured a representative sample of communities so that the opportunity is there for future work to correlate the relationship between disparate recovery trends and policy decisions, resource allocations, etc.

3.3.3. Comparing Results to SSM Predictions at Two Years Post-Event

To start, it is important to note that the original SSM was trained on and selected against the long-term outcomes of several factors: median household income, population, educational attainment, number of households, and number of housing units. This field study by comparison only tracked this final parameter regarding housing units as well as other building's status. As such, the results here are only speaking to the decline of a single factor and consequently may not perfectly align as social data on the community was not simultaneously collected and mapped across the six waves. In fact, it was during this first application of the tool that a secondary task was pursued to determine if predictions could be made as to the extent of recovery or decline for each long-term outcome rather than the aggregate set of monitoring factors. This curiosity led to the work on population outmigration prediction presented in Chapter 5 of this dissertation. Nevertheless, this data was analyzed with the expectation that the SSM scores would still meet their intended goal of predicting the relative recovery of communities in a field study and guide the research team to select communities that would experience disparate recovery trends due to the same event.

Figure 27 represents the building stock recovery trends for all six communities tracked in the longitudinal field study. The initial damage shows the disparate impact that this event had on communities across the study region with Mayfield, Dawson Springs, Bremen, and Samburg all having more than 80% of the surveyed building stock impacted by the event to some degree. It is worth noting here that the surveyed building stock is not and should not be interpreted as the complete building stock for each community. In some cases, such as Bremen, the surveyed area expanded well beyond the town's official boundaries to capture damage to unincorporated neighborhoods. Meanwhile in larger communities like Mayfield and Dawson Springs, the research teams could not reasonably survey every building in the community. As such, the building stock damage percentages would likely be lower for these communities if there was not an effort made to focus on damaged areas.

To account for discrepancies in recovery percentages based on this more ad hoc method for determining which areas of each community would be surveyed, Figure 28 has also been provided to show the recovery trends as a portion of only Wave 1 damaged buildings. Thus, at Wave 1 all recovery has been set to 0% and then rebounds in the subsequent waves according to relative recovery until it returns to 100% or more, indicating that all damaged buildings have been repaired or replaced to return the building stock to the pre-event levels. The drawback of this visualization is that it can look noisier for low damage communities like Leachville and Monette. These communities even appear to have a regression in recovery at different points according to this graph. However, this is actually a result of the initial damage being so low (around 10% of buildings surveyed) that subsequent waves with any error or missed survey points may represent a meaningful portion of the buildings for which recovery has not been verified. This is the first indication that Leachville and Monette did not receive sufficient damage to inform recovery trajectory trends more broadly.

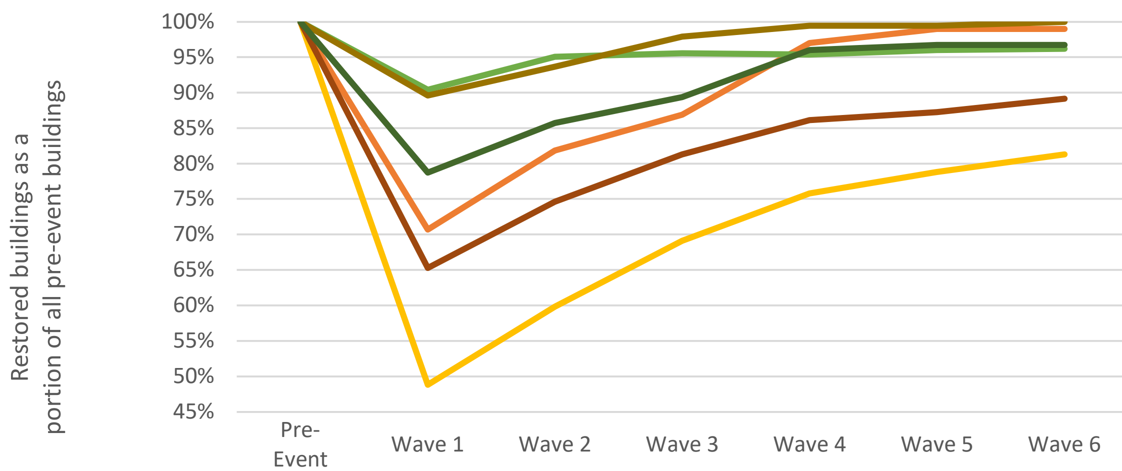

	Pre-Event	Wave 1	Wave 2	Wave 3	Wave 4	Wave 5	Wave 6
Bremen	100%	71%	82%	87%	97%	99%	99%
Dawson Springs	100%	49%	60%	69%	76%	79%	81%
Leachville	100%	90.4%	95%	96%	95%	96%	96%
Mayfield	100%	65%	75%	81%	86%	87%	89%
Monette	100%	89.6%	94%	98%	99%	99%	100%
Samburg	100%	79%	86%	89%	96%	97%	97%

Figure 27: Building stock restoration as a portion of all pre-event buildings for each community over six waves

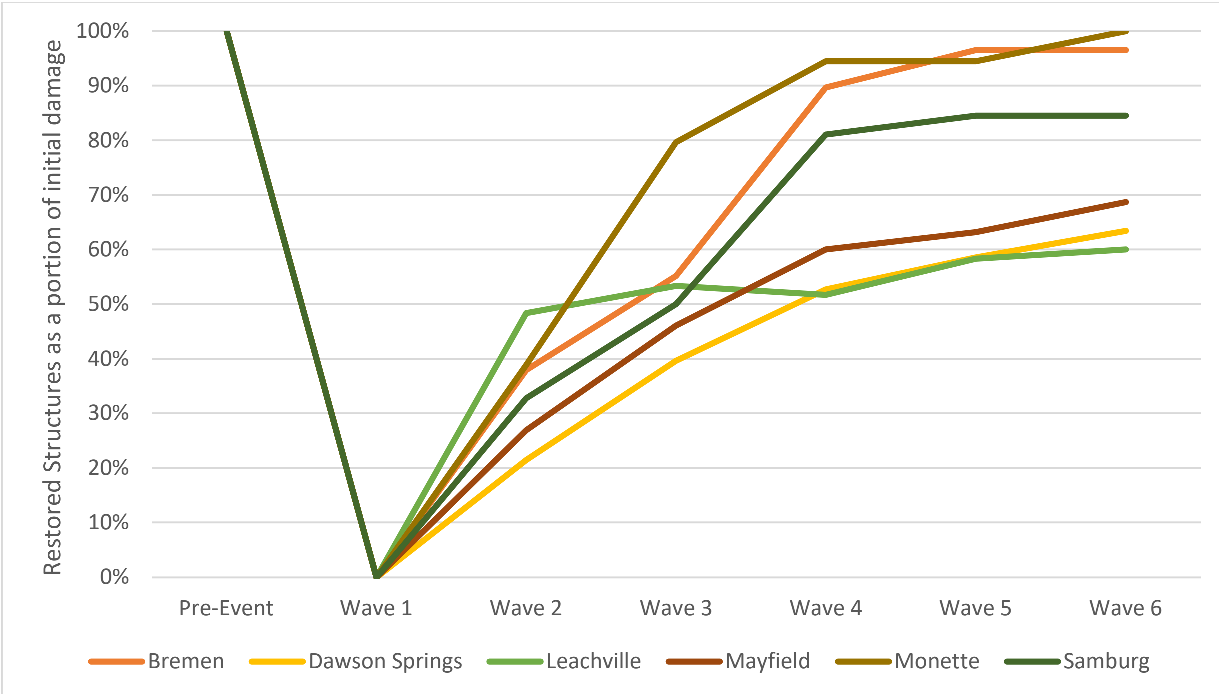

Figure 28: Building stock restoration as a portion of Wave 1 damaged structures for each community over six waves

In comparing these outcomes to the Social Susceptibility Metric, it was decided that the best method of comparison would be rank order lists. This was done because, as mentioned before, not all monitoring factors were monitored for this test. So, the resultant predictions from the SSM could not be applied to a single monitoring factor. For instance, it cannot be assumed that if Bremen will certainly see a decline specifically in housing unit recovery because its overall SSM projects marked decline. However, it seems fair to assume that marked decline of several other monitoring factors would have some relationship with either initial damage or long-term recovery. Table 10 has been generated to reflect this relationship between SSM and building stock outcomes due to the event. These outcomes include the initial damage post-event (Wave 1 values in Figure 27), the relative recovery at the 2-year mark (Wave 6 values in Figure 28), and absolute recovery at the 2-year mark (Wave 6 values in Figure 27). For ease of interpretation, all values that are +/- 1 off the predicted ranking are shown in light blue, while perfect matches are shown in dark blue. The light orange indicates poor fit. This poor fit only occurred in the two communities already flagged above, namely Leachville, Arkansas and Monette, Arkansas. It seems evident from the previous discussion and the results in Table 10 that these two Arkansas communities are likely bad candidates for implementation of the SSM in community selection. This has two likely causes. First, these communities saw significantly less damage as a portion of their total building stock, only being impacted by the tornado on their northern edges. Second, unlike the other communities that were quite far apart from each other and relatively independent from their neighboring communities, Monette and Leachville are only about 10 minutes apart by car, meaning that their recovery trends are almost certainly intertwined with one another. Because the granularity of the SSM model is at the county subdivision level for this study, the SSM may not reflect interplay between these county subdivisions as recovery is underway. Seeing as how recovery is oftentimes a regional endeavor, it is not recommended that adjacent communities be analyzed separately but rather should have their scores combined to reflect their mutual reliance. Alternatively and preferably, if the geographic range of the event allows, it may be preferable to avoid selection of adjacent communities unless the focus of the study is specifically the regionality of recovery.

Table 10: Rank order of several dimensions for all communities

	SSM	Predicted Order	Absolute Restoration Order	Relative Restoration Order	Initial Damage Order
Bremen	5.4	3	2	2	4
Dawson Springs	8.5	6	6	5	6
Leachville	0.9	1	4	6	1
Mayfield	5.8	4	5	4	5
Monette	6.1	5	1	1	2
Samburg	3.5	2	3	3	3

If for instance in this study, the two Arkansas communities are retroactively removed from consideration, the study would yield the results shown in Table 11 with more than half of the cells indicating a perfect prediction match and no instances of poor fit. The only communities with fluctuating order are Bremen and Samburg whose recovery patterns have been closely mirroring each other with small margins of separation (see Figure 27). Although this is only a single application of the SSM, this improved predictive performance would seem to indicate that the SSM performs better when low damage and/or adjacent communities are omitted from the sample. Additional applications of the SSM would allow for more meaningful and systematic review of the effectiveness of this metric.

Table 11: Rank order of several dimensions for communities aside from Arkansas communities

	SSM	Predicted Order	Restoration Percentage at Wave 6	Absolute Restoration Order	Relative Restoration Order	Initial Damage Order
Bremen	5.4	2	99%	1	1	2
Dawson Springs	8.5	4	81%	4	4	4
Mayfield	5.8	3	89%	3	3	3
Samburg	3.5	1	97%	2	2	1

3.4. Summary and Conclusions

Due to the nature of the events studied in disaster research, large and reliable datasets are difficult to find and even more difficult to compare against one another without being able to control for any of the changing variables. Establishing this dataset for a series of communities that were simultaneously impacted and did not experience another event of similar magnitude over the course of the study period allows for a more controlled study. The researchers involved in this work are continuing to use this data to investigate a range of topics pertinent to hazards research today. Within the scope of this dissertation, this work has taken strides to better understand how the same event can have varied long-term impacts for communities based on social characteristics, which can in turn inform modeling efforts made later in this dissertation.

In this chapter, the data that has been created and collated here has been compared against the SSM predictive tiers. When paring this list down to only those communities that saw significant damage, the SSM is found to be quite predictive in determining relative recovery rates as well as other parameters. For instance, when reviewing the results summarized in Table 11, the alignment of the initial damage state rankings with the predicted rankings is an unexpected but possibly elucidating revelation. Exploration into the alignment of initial damages and SSM values would not only be an interesting and helpful by-product of the metric but could also mean that wider use of the SSM outside of just recovery contexts may be appropriate. This is however a small dataset and would need a much larger body of corroborating data to be deemed appropriate to use in other contexts. Further validation of this wider implementation constitutes a possible area of future investigation.

This dataset demonstrates the way in which recovery trajectories can vary for a set of communities impacted to different extents by the same hazard event. This type of controlled investigation provides a window into the mechanisms that may help or hinder building stock recovery. More investigation of this dataset as well as development of similar datasets following other types of hazard events that span multiple communities will certainly be fruitful in uncovering knowledge about these recovery mechanisms and how they can be leveraged to benefit communities.

4. TESTING POTENTIAL COMMUNITY TORNADO SHELTER LOCATIONS WITH SHORTEST PATH MODEL

4.1. Background and Motivations for a Sheltering Model

Tornadoes affect many communities throughout the southeastern and central United States each year. In this region, in-home tornado shelters are rare, especially amongst lower income households. Thus, when a tornado warning occurs, many people are left with very few safe options. FEMA reports that tornado warnings offer between ten to fifteen minutes of time to seek shelter. This time could be spent navigating to a nearby, well-documented shelter with which the household is already familiar. This model uses a shortest path approach to test the viability of placing future community shelters on the same site as dollar stores across the region.

Many low-income residents in tornado-prone regions do not have adequate access to tornado shelters and live in homes that are likely unable to offer life safety even in low-level (e.g. EF1, EF2) tornadic events. As such, community tornado shelters are needed to bridge this gap in life safety protection, but determining the ideal location for these shelters is difficult and not frequently determined in a scientific way. Furthermore, these shelters are often only made possible with grant money from FEMA programs that require some level of matching. If the prospect of writing the grant application doesn't dissuade small, under-resourced, and possibly even unincorporated communities, then the cost sharing model very well may (Riley, 2024). Hence many of the residents of these communities do not have access to community shelters, and there seems to be few paths forward for public-only funding models to bridge this gap, i.e. often only available after a tornado has impacted an area through FEMA Hazard Mitigation Grant Funds. In order to put forth one possible partial solution to this community tornado shelter dilemma, the use case presented here uses variety (or dollar) stores as potentially well-positioned sites for tornado shelter spaces. Anyone who has been to the rural southeastern and central United States can speak to the ubiquity of their stores, and in fact, the annual report from the largest of these chains notes that they have stores within 5 miles of 75% of all U.S. residents (Dollar General, 2021). If these stores could share their sites with community tornado shelters and possibly aid in funding the acquisition or construction of such shelters, then community tornado shelters could roll out rapidly to many more populations that need these life-saving resources.

As for the motivations fueling a dollar store chain's potential interest in this type of work, dollar store chains have faced scrutiny in the past few years. As many say their business model takes advantage of rural communities by locating food deserts, building new stores in these locations, and then providing shelf-stable, sub-optimal nutrition options to their customers who have few other options. These scathing reviews have created a public relations problem for dollar stores, which the largest dollar store chain in America, Dollar General, is trying to combat with a combination of rebranding efforts and philanthropic work. Regarding rebranding, they have introduced two new store concepts. The first is a store all its own called pOpshelf, which provides more non-food products, "such as seasonal and home décor, health and beauty, home cleaning supplies, and party and entertainment goods." The second concept aims to address the primary critique head on: DG Fresh is an initiative to bring (more) perishable items into their stores, which purportedly will "reduce product costs, enhance item assortment, improve our in-stock position, and enhance sales" (Dollar General, 2021). These efforts at diversification and rebranding are mentioned here because they speak to the fact that Dollar General does focus on its public relations challenges after negative reports and could be open to new means of brand improvement. This might suggest that they are perfectly oriented at present to take on a new venture if that venture aims to improve the way their relationship with rural communities is portrayed. Furthermore, sharing their site with a community tornado shelter and possibly aiding in the funding of such a project would not be their first venture into community aid programs, or even disaster aid, as noted by the community page of their website (*Community*, 2024). All this seems to frame well the fact that a public-private partnership between dollar store chains and FEMA, or some other federal agency, could save lives, decrease local economic burden, and provide a desired public relations boost for dollar store chains across many rural and under-resourced communities that they are currently located within.

4.1.1. Potential Shelter Locations Used in the Model

Variety stores, also known as dollar stores, are used in the case study implementation of this tool as a means of demonstrating not only the viability of this model at scale but also the case to be made for collocating community tornado shelters with variety stores. Because variety stores have a business model that stresses customer convenience (i.e. low travel times), designs for small footprint and dispersed locations, and serves many low-income, rural populations. These populations are an ideal audience for these new community tornado shelters as they are frequently outside the travel area for a traditional community tornado shelter and are more likely to live in poorly constructed buildings. Additionally, if a household has frequently seen a tornado shelter during their grocery shopping visits, it seems to follow that they would be more likely to consider this as a possible option when a tornado warning is issued. A final benefit of this collocation with variety stores is the fact that variety store locations are already well-documented and accessible online, meaning that residents and those already on the road could quickly search, or better yet be alerted about, their nearest potential shelter location in a standardized manner via their mobile device.

The idea to test the viability of variety stores as a sort of proxy for shelter location optimization was first arrived at organically during data collection for the first wave of the Midwest Field Study discussed in the previous chapter. This was simply because there were variety stores seemingly everywhere in the field study region. Indeed, the town of Bremen, Kentucky alone has three variety stores within the present 9-minute tornado sheltering travel radius from the center of town, while the population within the town limits was only 172 in the 2020 census. Refer to Figure 29 to see the prevalence of these stores across Dixie Alley. It was hypothesized that in these very small towns and rural areas with a nearby dollar store, or several, a pre-built shelter placed somewhere on the (frequently oversized) lot of a dollar store could potentially accommodate many households within a 10-minute radius of the store while in other instances a smaller sheltering radius or larger shelter may be necessary. And though much can be said on the topic of chain stores, and dollar store chains in particular, becoming increasingly common across rural America, the availability of a life-saving resource on the site of a dollar store chain location would streamline the process of helping some of the most at-risk members of a community have access to a shelter in many areas where these locations either don't exist or are not well documented.

For this investigation, the area of interest was limited to a state which has been subject to several deadly tornados in recent memory and for which a single chain made up the largest portion of variety stores to avoid having to collect data for all the major retailers in this market. As such, Alabama and the chain with the largest market share shown in Figure 29 were selected for analysis. A clear caveat to the analysis provided in this chapter is the fact that sheltering times and required capacity of each shelter location would decrease if other brands were added to the model, making this idea only more viable than presented here.

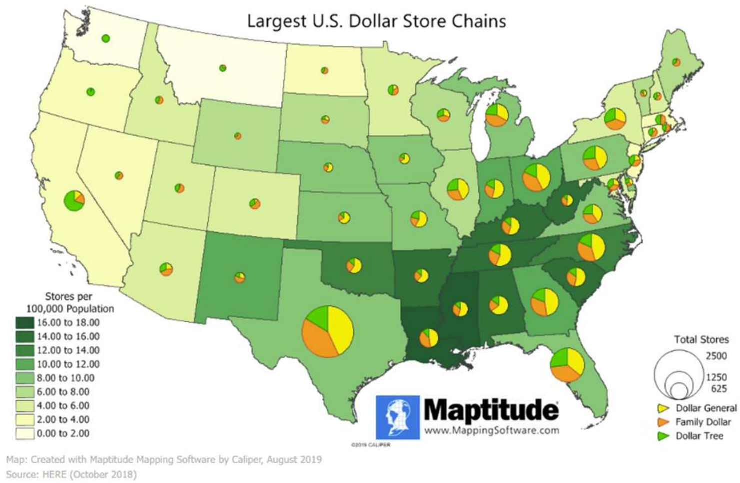

Figure 29: Map of Dollar Stores by State (Largest U.S. Dollar Store Chains, 2019)

4.2. Approach

This model has in fact been developed as two standalone yet chainable models. The first model identifies the closest shelter for each residence from the National Structure Inventory in a continuous space, colloquially referred to as "as the crow flies" (US Army Corps of Engineers, 2022). Then, the second creates the fastest OpenStreetMap route between each residence and the closest shelter that was identified in the first model. This choice to bisect the tasks proved to be necessary when running this model on a large scale. The example implementation of this model shown below primarily uses one-degree by one-degree cells (except in the most populated areas) for analysis that are then overlapped in a patchwork to cover the entire state of Alabama. If only the second model was used and households within this model were allowed to select their own closest shelter given the full range of possible routes from their residence to every dollar store within a one-degree by one-degree cell, this introduces too many permutations and drastically hampers the efficiency of the model while also risking that the nearest shelter is actually just over the border of an adjacent cell. As such, these models are divided into these two primary tasks for each residence: (1) picking the closest shelter, and (2) routing to that location.

4.2.1. Continuous Model

The continuous model has households determine their nearest variety store and then record this store location to be used in the subsequent model. The efficiency of this model allows it to run for a five-degree by five-degree cell relatively quickly. For instance, the model ran for the entire state of Alabama in under 3.5 hours using 2500 CPU cores. The only necessary input is a table of initial locations of the households (their housing units) and another table of candidate shelter locations. From these tables, the two necessary entities are formed. The first is the shelter locations, referred to in the code as "DS" or dollar stores in alignment with the use case the prompted the code's development. These are static entities that calculate the distance between themselves and every household in the model's space and then pass this information back to the model. The second entity type is those that are seeking shelter, referred to in the code as "households." These households receive the list of possible shelter locations and their corresponding Euclidean distances from their start location and then select the closest shelter location with which to pair. This paired list of households and their selected shelter location is fed into a csv, and this serves as the primary output of this sub model. There are several ways to calculate a closest available resource, but this model was developed in order to make it fully chainable with the subsequent model such that, if desired, a user can add both codes to a single file and have the whole sheltering process run without requiring the intermediate step of outputting to a .csv file. This is not done in the use case below because the region of analysis in the continuous model had to be divided and ran in several sub models due to the region's size and the computational cost of the OpenStreetMap ABM.

For this use case's implementation of the code, the candidate shelter locations were the locations of the state's largest variety store brand, and the households' initial locations were all the residential structures from the National Structures Inventory, assuming for simplicity that every residential structure has one household. Thus, this application of the sheltering model and its companion OpenStreetMap model look to simulate how many individuals could navigate to a nearby variety store from their home in the event of a tornado warning being issued. It does not account for those individuals who are already on the road at the time of the warning or who are in commercial structures nearby. It is also assumed that each household will travel in a single car. Households are taken to all be the mean value of 2.5 people for the state of Alabama based on 2020 Census data for the sake of efficiency; however, creating a normal random distribution centered around this value could be done by adding a new attribute to each household and randomly assigning household size by sampling from a normal distribution. Although more would need to be known of the standard deviation for household size as well. The resultant map of where each household is planning to shelter is shown in Figure 30, along with a zoomed snippet to show what this looks like locally. In Figure 30, the variety store locations are displayed by black-outlined, oversized gray dots and the households that will potentially shelter at each gray dot are shown in varying colors to indicate the different variety stores to which they are closest. Some colors are necessarily repeated due to how many variety stores are present in the dataset.

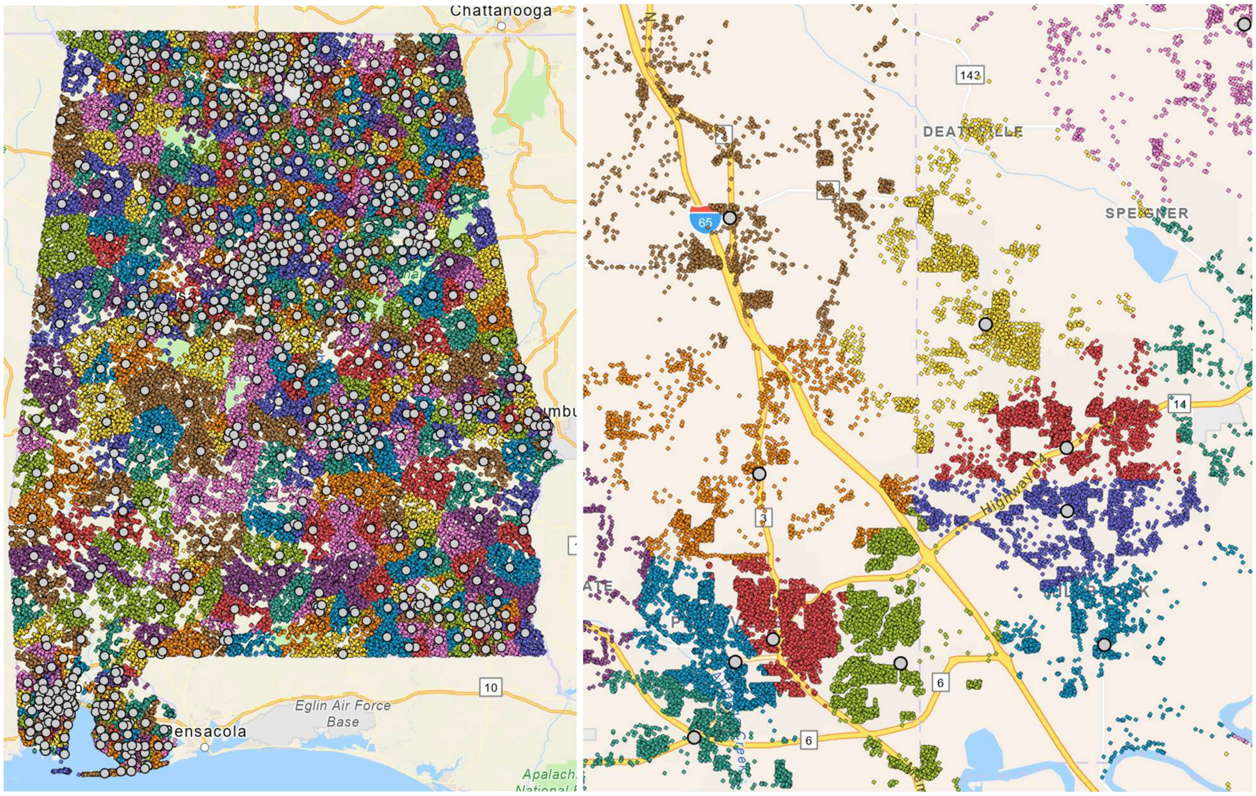

Figure 30: Continuous model results where each color represents households navigating to a different dollar store

4.2.2. Open Street Map Model

Once every household has been paired with a shelter, the results are passed into the OpenStreetMap space to allow each household to route to their respective shelter on the street network downloaded from and available via OpenStreetMap (OpenStreetMap Contributors, 2024). The households will be assigned a random amount of time to leave their home and reach their nearest roadway in their vehicle based on a bounded normal distribution centered at two minutes and standard deviation of three minutes with a lower bound of one minute and an upper bound of nine minutes, see (3). Thus, if a home is randomly assigned a home departure time of more than nine minutes, then the home will be reassigned a home departure time of nine minutes, though this will affect a very small number of households due to the characteristics of the curve. Conversely, if the household is assigned a home departure time of less than one minute, then it is realistically set to one minute.

This adjustment will affect many households. This normal distribution is characterized by using logical reasoning as to the time required to suddenly leave one's home and certainly an area of the model that can be adjusted in other applications or when more warning time is provided.

$$Departure\ Delay \sim N\left(\frac{2}{60},\left(\frac{3}{60}\right)^2\right);\ LB = 1;\ UB = 9 \tag{3}$$

The first step is for every household to determine the shortest route to the closest shelter found in the continuous space model. If the household's route is greater than the set maximum sheltering time, then the household will not attempt to shelter. This limit is placed on the households both to avoid households still being on the road when the warning time elapses and to realistically simulate the fact that households who know they live far from a shelter are unlikely to see community sheltering as a viable option. In this exploratory model, it is assumed that every household wishes to shelter at the beginning of the model run unless their route is unrealistically long. This assumption is not made because it is believed that all households will wish to shelter following a tornado warning but rather because this model does not allow households to impact the commute times of other households thus it is as if this model is running for each household in its own space devoid of other households. Thus, this model looks to capture how quickly each household could reach their nearest shelter in an idealized scenario. It is then recommended that this sheltering time is scaled by a factor of safety to account for other drivers in a real-world scenario. For tornado warnings that lead to minimal action from the public, this value may be closer to one, and in situations where there are widespread sheltering efforts, a higher value is more appropriate. This second case portraying widespread sheltering has not been historically seen following a tornado warning, and as such, this model allows households to shelter without regard for their fellow households to prevent extraordinary traffic situations that are unrealistic for instances of tornado sheltering. It is of course hoped that as more innovation occurs in the space of community sheltering, such as that proposed here, households will feel more compelled to seek shelter when tornado warnings are issued, which will complicate the matter of congestion and traffic impacts.

The maps used in this model are downloaded from the crowd-sourced mapping service OpenStreetMap using the Julia package LightOSM (Deloitte Digital, 2020). The full list of nodes and edges that make up the road network are downloaded via a json file and this file informs the creation of a continuous space model that makes all regions that do not fall along a road network impassible. Thus, the road network functions in this model not unlike the tunnels in a mine with the primary difference being that information regarding the types, speed limits, lanes, etc. of roadways are also stored when this information is available on OpenStreetMap. For efficiency, slight road curvatures are neglected. A final important note on the LightOSM Julia package is that it uses dimensions of kilometers and hours. These units are reflected in the parameter provision throughout the model.

At the interface of the LightOSM and Agents packages, agent speed can be defined in three primary ways according to Agents.jl's dependent package LightOSM.jl: speed, time, and lane efficiency. The first is a straightforward provision of a speed which the agent will travel at for the duration of the model run unless adjusted at an intermediate step. By setting the map's unit to time instead, each household's speed will be set according to the max speed allowed on a given road according to the OpenStreetMap file. The third and final option is to take the prior time calculation and scale it according to the number of lanes on the road, such that less lanes will lead to slower speeds. This routing type is called "lane efficiency," and its lane scaling factors are as follows: a one-lane road equals 0.7, a two-lane road equals 0.8, a three-lane road equals 0.9, and a four-lane road = 1.0. This last modification helps to address the fact that the households in the model navigate to their shelter irrespective of the other drivers on the road; however, the logic behind these reductions does not hold merit in a sheltering scenario and thus this option was not used in the present study. Based on these values, the upper bound of sheltering time would be the time divided by 0.7, assuming all roads have one lane each way. This upper bound is helpful to bear in mind when considering congestion and will be extended when considering sheltering congestion later in this chapter. The randomized speed option also has some strengths in that it allows more household choice; however, leaving this speed constant for the duration of the model run is unacceptably unrealistic and resetting the speed at every second is a possible permutation of the model that was not explored here. For this case study, the decisions made at every step assumed rational and obedient actors. As such, the middle option was chosen by which the speed limit defines the speed at which every household is traveling on a given road. When no speed limit is defined in the dataset, the default is assumed. Default speed limit values are assigned according to the road tag: "motorway" equals 100, "trunk" equals 100, "primary" equals 100, "secondary" equals 100, "tertiary" equals 50, "unclassified" equals 50, "residential" equals 50, and "other" equals 50 with units of kilometers per hour.

As alluded to above, this model is operating in truly continuous space meaning that households can occupy the same exact place and locations are true floats, allowing for households to occupy any location along a road. This is unrealistic and does not sufficiently characterize the impact of traffic. This decision is considered acceptable for the use case of tornadoes due to two reasons. First and primarily, tornado warnings do not enjoy universal belief let alone universal knowledge of public shelter availability (Brotzge & Donner, 2013). This confidence in warning systems is further diminished depending on time of day with midnight to 4am being the lowest point in this collective confidence (Krocak et al., 2021). Thus, the assumption that every household will shelter and lead to intense congestion is highly unlikely unless public sentiment to warnings in tornado prone regions shifts dramatically from the present situation. Second, it is recommended that the factor of safety for sheltering time be set to two. This would mean that if the anticipated warning time was set to a lower bound of ten minutes, then only those households within 5 minutes of their local variety store would be advised to shelter. Making the necessary modifications to account for traffic concerns would also require more elaborate consideration and quantification of each household's probability to shelter in order to simulate a realistic number of drivers on the road. These extensions and elaborations are recommended for other permutations of this sheltering model to handle different hazards, warning times, and traffic considerations.

The model thus iterates through the complete list of households. Each household generates the route from their home to their shelter. Households then determine if their sheltering route is too long to plausibly shelter in time, in which case they stay home. Then the households that can shelter in time prepare to leave their home according to their randomly assigned preparation/home departure time. Then the household moves along their shortest route at the roadway's speed limit without regard for fellow households for however many discrete timesteps (minutes) it takes to arrive at the shelter location. The time is then recorded in the sheltering timing during the following timestep to allow one minute for households to exit their vehicle and enter the on-site tornado shelter. At the final timestep of the model, 20 minutes after the warning was issued, the shelters will tally how many households they have on their site, and the model will terminate.

For this use case, the model is run for overlapping grid cells of varying size as shown in Figure 31 with grid cells containing more households being shrunk slightly to reduce the difference in computational requirements. Still images taken from the video output of one iteration of grid cell number 19's model is shown in Figure 32, covering the area to the northwest of Montgomery, Alabama. In this use case and the code by default, only ten iterations of the model are run due to the computationally expensive nature of running an entire state's worth of data. However, when the region of interest is reduced, the model can easily be run for 100 or 1,000 iterations to effectively characterize the output probabilities. Furthermore, the model can be made more stochastic by modifying the model to run based on randomized speed and randomized warning times. This permutation of the model is available at https://.github.com/blythejohnston.

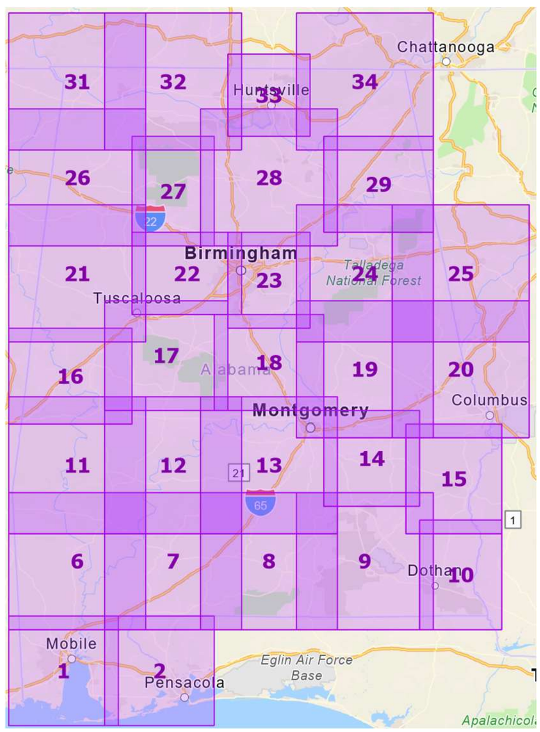

Figure 31: Analysis Cells for the OpenStreetMap model

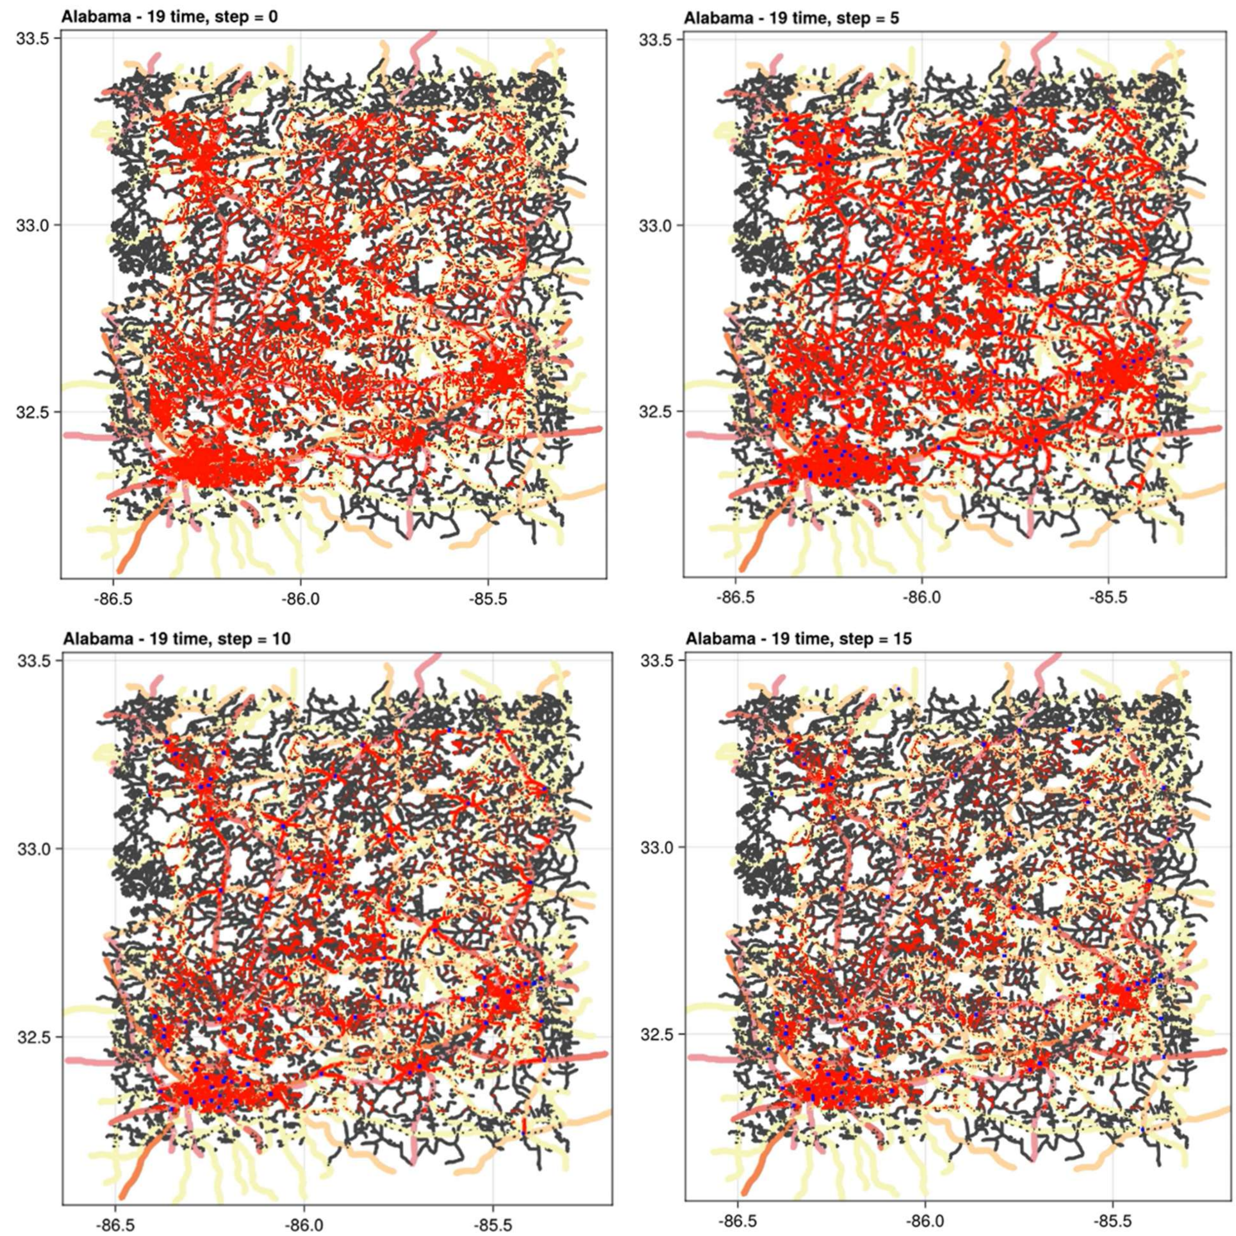

Figure 32: Still images from the video file output for grid cell 19 in which red points represented non-sheltered agents and blue points represent sheltered agents (at their shelter)

4.3. Results

The results of the OpenStreetMap were output for every grid cell into corresponding csv files. For households and shelters where cells overlap, the households' sheltering times and chance of sheltering, as well as the shelters' number of households, are averaged. This chance of sheltering is arrived at by running ten parallel realizations of the model to capture variability of the data. Due to the low model run count, the values presented here may not adequately capture variability and are thus only shown here as an illustrative and exploratory use case of this sheltering model and the proposed collocation of variety stores and tornado shelters across an entire state. More in-depth analyses of a particular area would reduce the geographic range of a particular model run and increase the model run count.

Returning to this example, the grid cell output files were modified and combined to reflect the summary analysis across all grid cells, which is shown in the rest of the results presented here. An example of the processed results has been shown for grid cell 19, which is shown in Figure 33. The upper left image shows the continuous model results matching households to shelters with the households living in homes within each shelter's influence area denoted by a varying color to reflect to which shelter they will navigate. The upper right image layers the OpenStreetMap shelter times on top of the upper left map. These are displayed on a continuous scale from two minutes to fifteen minutes. These sheltering times are the average times across ten model runs and incorporate the departure delay time outlined above. Those households for which sheltering is not possible within the defined maximum warning time of 20 minutes are not displayed on the continuous scale, and thus you can see the underlying layer for these points. No sheltering is projected to occur for households living in those smaller multicolor housing unit points shown. Lower left shows this map of shelter times simplified to focus on the predicted sheltering times with all other layers turned off. The lower right image shows this map broken into tiers of sheltering time to delineate those households for which sheltering ranges from "unlikely to shelter in time" to "recommended to shelter." This final category accounts for those households who can reach their shelter in five minutes or less and are thus recommended to shelter for a warning time of ten minutes and factor of safety for road congestion of two. An identical analysis was performed for all other grids, which were then concatenated to create a statewide analysis.

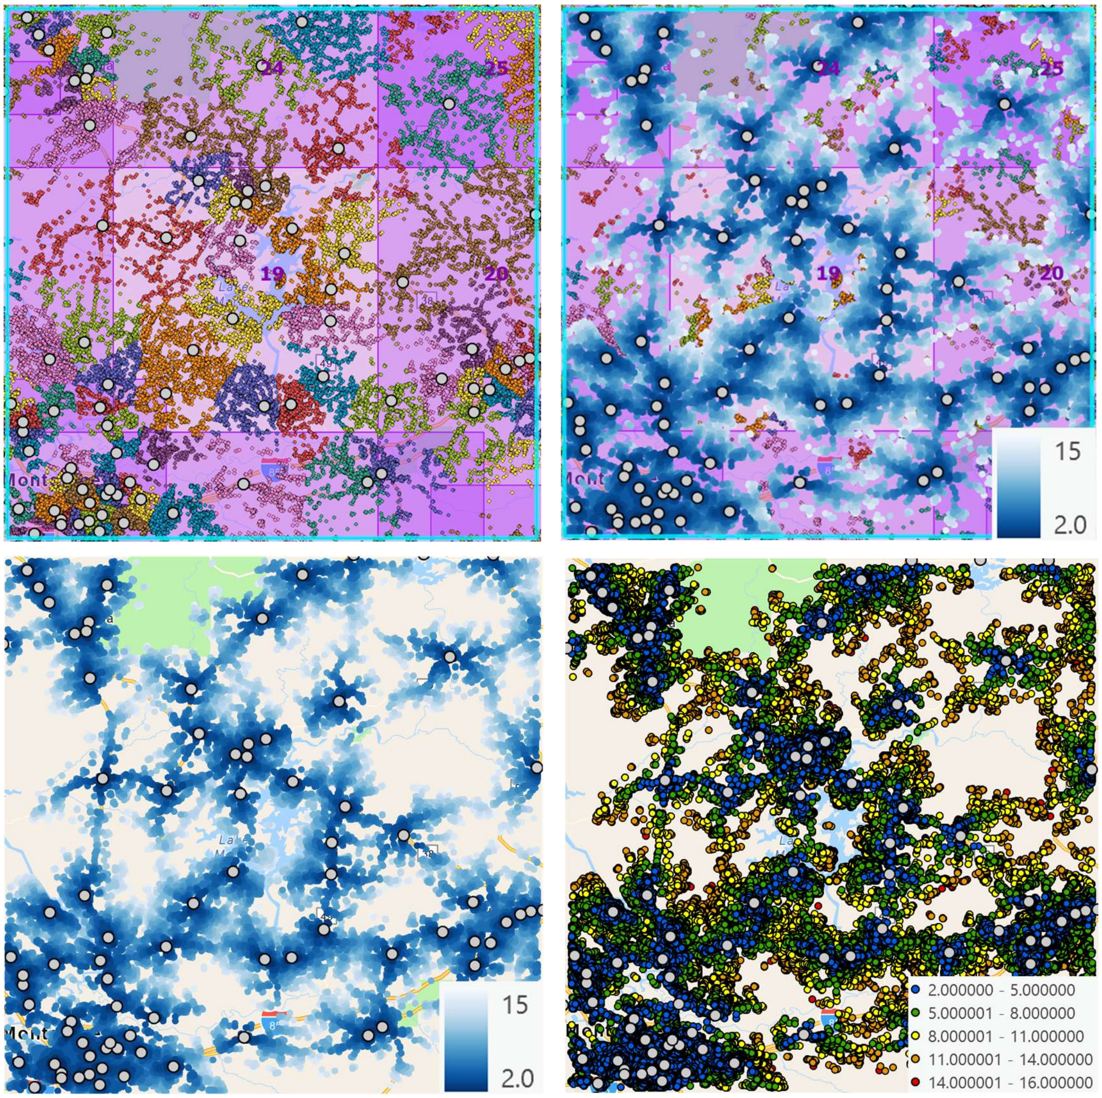

Figure 33: The results from the sheltering model for grid cell 19.

In conducting this analysis for the entire state, it was found that approximately 1,770,000 of the 1,920,000 residential structures in Alabama are within 16 minutes of their nearest variety store, including departure time. Of those 1,770,000 households, nearly 850,000 residential structures in the state of Alabama live within 5 minutes of a store location for the state's largest variety store brand. This number would increase if considering all variety store brands and without granting consideration of departure time. The visual summary of this complete analysis can be found in the left image of Figure 34, while the truncated analysis for just residences within five minutes of a variety store is shown in the right image of Figure 34. Considering that there are about 880 locations for the primary variety store brand in the state, this analysis would suggest that an average of 960 households with an average of 2.5 people will shelter at each variety store location if only those within five minutes choose to shelter. This means that 882 locations could provide shelter to over 40% of the population of Alabama if all households in this five-minute radius sheltered and proximity were the only limiting factor, meaning that parking space would have to be unlimited and shelter space would need to be plentiful. Indeed, the average required shelter size would be 14,250 square feet, taking the shelter space recommendation to be 6 square feet per person, and the average parking area would need space for 1,000 vehicles. Thus, it becomes clear that the limiting factor for this analysis is not the proximity of variety stores to enough households but rather the parking capacity at these locations, the traffic control at the shelter sites, and the required size of installed tornado shelters.

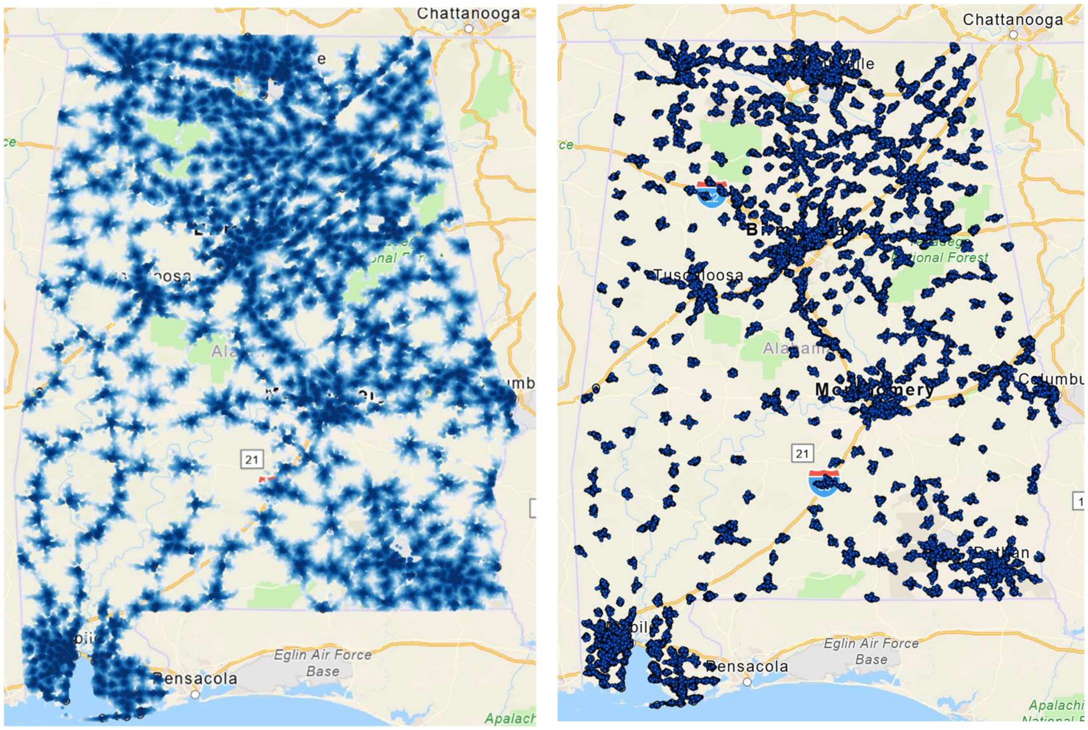

Figure 34: Whole State analysis for Alabama displaying the sheltering times for households across the state navigating to their nearest variety store (left: within 20 minutes; right: within 5 minutes)

As such, the data was re-analyzed to determine how these average values were dispersed across the state. This would allow each variety store location to be divided into a case of potential shelter location. The first case would be those whose projected sheltering rate would be low enough to merit use of the approach used in this model by which households travel by car to a location within a five-minute radius of their home. The second case would be suburban and urban locations for which navigating to the shelter by car would be too difficult to coordinate and would lead to an unmanageable number of sheltering households. In this case, a foot-traffic only sheltering approach is recommended by which households determine if they have sufficient time to walk to their nearest variety store, i.e. a five- to ten-minute walk at the individual user's pace. And finally, the third case represents those locations for which the siting of the variety store is largely rural, but too many households exist within a 5-minute boundary, so the sheltering region needs to be reduced further to ensure enough space is available for all residents.

This analysis yields the distribution curve shown in Figure 35 and the maps shown in Figure 37 for the total count of sheltered households at five minutes after the warning is issued (left) and the maximum time provided by the model (right). Figure 37 demonstrates that many but not all the shelters receiving over 1,000 households in the first 5 minutes are in urban areas, making this subset (urban with 1,000+ sheltering households) good candidates for case #2 delineated in the previous paragraph calling for foot-traffic only. The high sheltering estimates occurring in less urban or suburban areas are recommended to either narrow the sheltering region by decreasing the available sheltering time or possibly not recommend this location for a tornado shelter. This requires additional consideration of overflow alternatives, other variety store brands' locations nearby, and other community shelters that are already available. The technique of pairing down the sheltering region even further is shown in **Error! Reference source not found.** for the variety store location outside Eclectic, Alabama. As for case #1 in which households drive to their well-managed and smaller sheltering location, this is recommended to be considered on a case-by-case basis for the variety stores expected to see less than two-hundred households sheltering in the first five minutes. The map of the 68 potential store locations for this is shown in Figure 38.

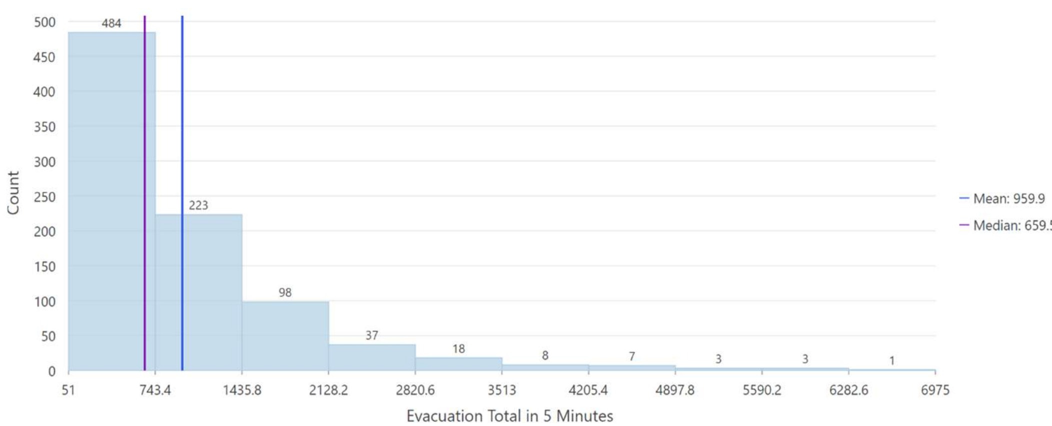

Figure 35: Distribution of the anticipated household sheltering count at each variety store location

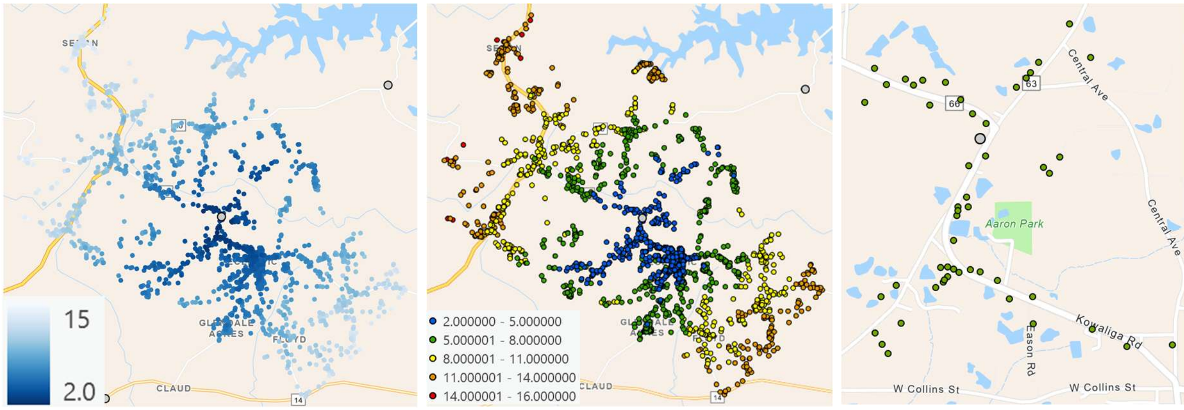

Figure 36: Sheltering times for all households allocated to store 1919052. Left: all sheltering households for the 20-minute model run (1369 households). Middle: tiered sheltering times to see the blue tier of 5 minutes or less (422 households). Right: the select households who could shelter in 3 minutes (51 households)

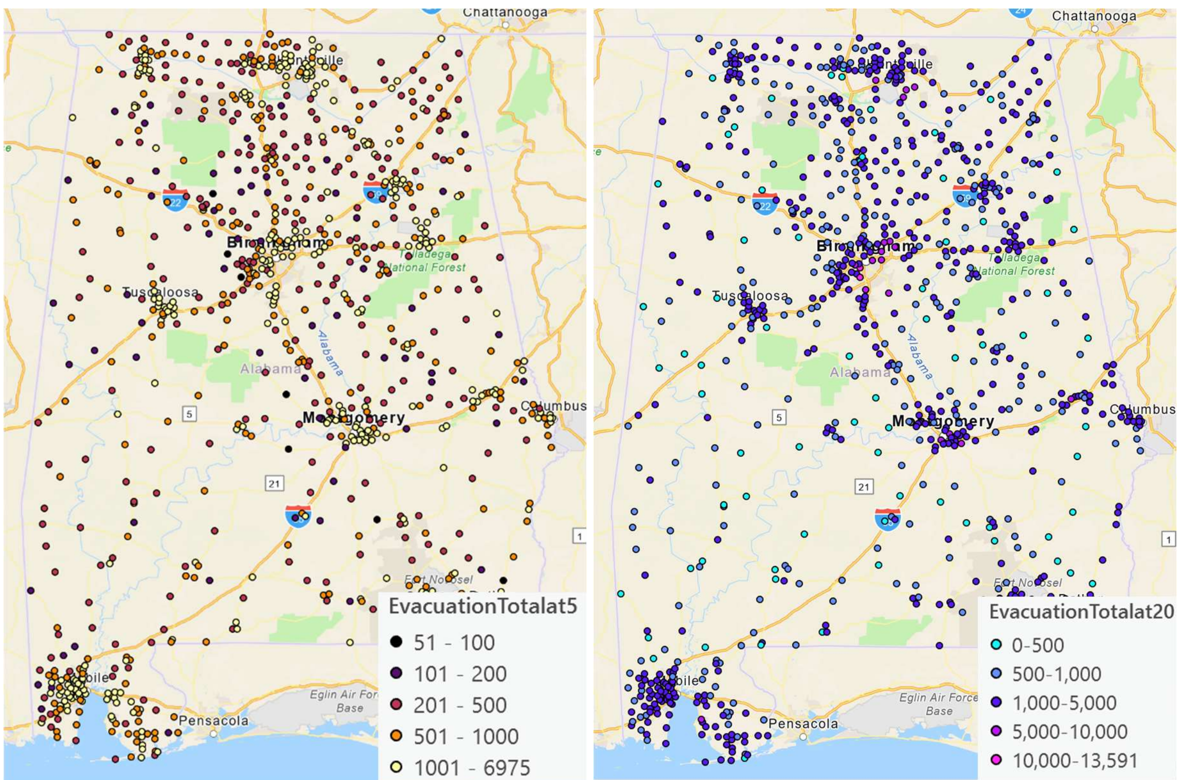

Figure 37: Variety store locations with sheltering totals at 5 minutes (left) and 20 minutes (right) after the warning is issued

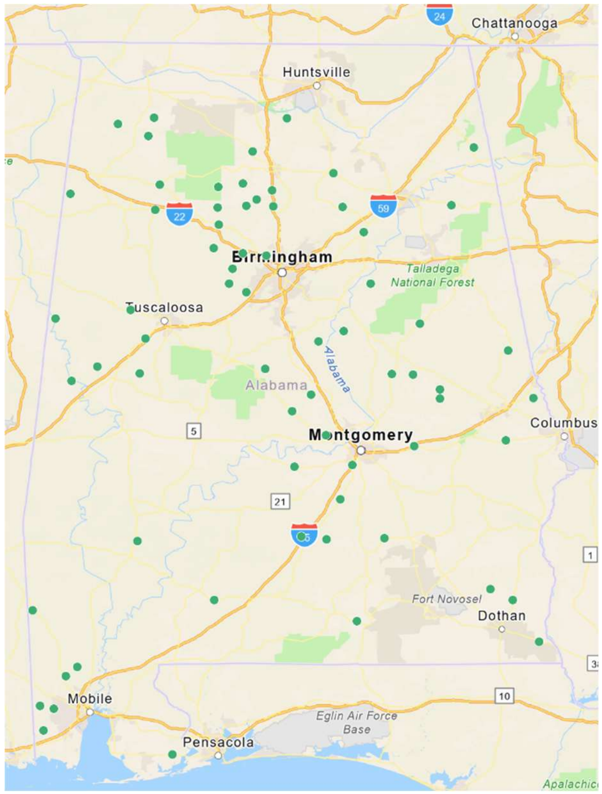

Figure 38: Variety stores for which the number of sheltering households 5 minutes after the warning is anticipated to be 200 or fewer households

4.4. Discussion and Areas for Future Work

This work demonstrates that this model can sufficiently characterize the potential for a given location to operate as the site of a tornado shelter by being sufficiently close to a critical mass of sheltering households. Similarly, this use case has demonstrated the ubiquity and strategic location of dollar stores across the state of Alabama, and the case for neighboring states is similar. Thus, this hypothetical of placing tornado shelters on the site of dollar stores across low-income and active tornado regions seems to represent a viable option for unincorporated and/or underserved populations. Optimization of tornado shelter locations is a vital area of investigation in the coming years, and it is hoped that this generalizable model can be adopted into this optimization protocol to allow for tornado shelters that will provide life safety for the most vulnerable members of a community during these events.

The most apparent drawback of this model is that "the crow flies" over mountainous terrain, water bodies, and other barriers on the earth's surface with ease, meaning that this model will have households adjacent to impassible boundaries travelling a greater distance to drive around a break in the road network to a shelter on the other side rather than choosing the more obvious option that is farther away in continuous space but closer according to the road network. A minor example of this is shown in Figure 39 for the use case presented here. This limitation is a minor issue in areas where the road network effectively crisscrosses the region, but if this is a concern two solutions are proposed, first the second model can be adjusted to find the nearest option along the road network or the output file from the continuous space model can be modified manually to send some households on peninsulas, riverbanks, etc. to a more suitable shelter. The first alternative will slow down the model immensely and thus should only be run for very small regions, making it not very scalable, and the second will rely on local knowledge and opinion, making it not generalizable. Thus, the bisected model described and demonstrated below is seen as the best alternative to this problem.

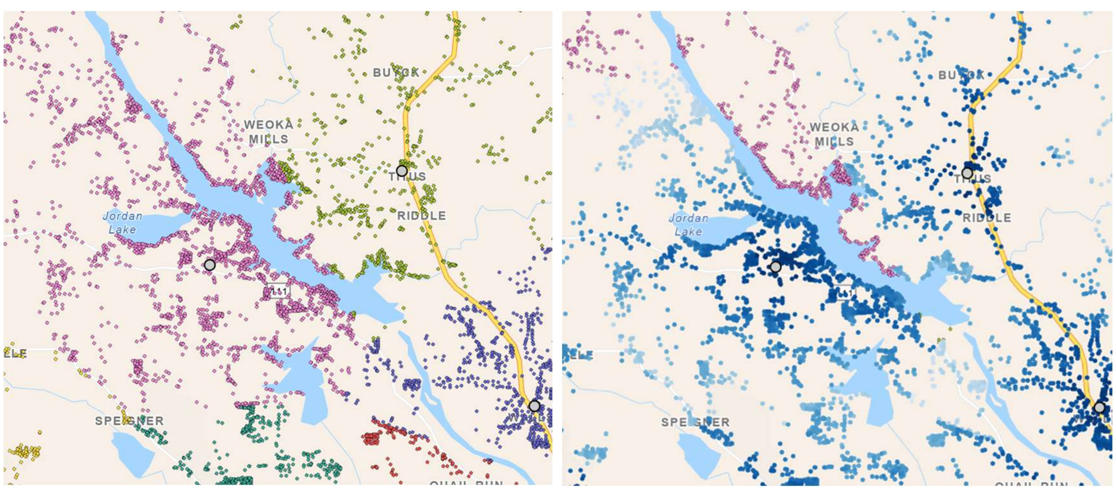

Figure 39:An instance of the continuous model overlooking a physical barrier in the road network (pink points in left image) leading to poorer performance of the subsequent model (non-sheltering on the NW riverbank in the right image)

As stated above, this model is most suitable for situations where the primary goal is to determine the viability and optimality of a potential shelter. As such the lack of traffic considerations for this model is not seen as a hinderance to addressing its primary goal; however, other models or future exploration would be warranted to create a similar fully generalizable model that can account for traffic conditions. This would be best suited as a third model in this chain in which the first is the continuous model, the second is the OpenStreetMap model to check the optimality of the potential shelter locations, and the third would be a similar OpenStreetMap model that requires households to provide sufficient space to nearby households in order to more accurately define the sheltering time needed beyond just the factor of safety applied here. Nevertheless, the methodology provided here serves to address a pressing question in tornado shelter literature at present which is determining idealized shelter locations in a scientific and unbiased manner.

Other elements that are always of concern when modeling sheltering patterns include the need for traffic control during this rapid onset event, the accessibility of such solutions for limited mobility individuals, and the households without a vehicle not being able to navigate to a shelter on foot. These concerns would need to be addressed as any sheltering theories move toward actualization, likely by members of the disaster preparedness field who have more experience in these topics. The case study examined here certainly shows promise, but in order to actualize this concept, several steps would need to be taken. First and foremost, interest in this concept would need to be expressed by large dollar store chains, and FEMA or another federal agency would need to establish a protocol by which a public-private partnership can be used to provide life-safety. A slow roll out of such a partnership would certainly be wise with annual targets for new tornado shelter constructions after viability of the model is shown. This project would be somewhat of a long road with many parties to convince along the way, but it is the belief of this author that this viability study has yielded results indicating that this work could result in material benefit to communities.

4.5. Summary and Conclusions

Although there have been several investigations regarding evacuation modeling, less work has been done in the space of systematic tornado sheltering. Furthermore, the provision of these shelters in unincorporated communities is difficult to coordinate and fund, making it less common. This leaves many rural populations living in structures without adequate sheltering spaces at risk of tornado damage and further risk to life safety. The generalizable tool outlined in this chapter could help to site tornado shelters across a region while only analyzing the user-input locations to allow greater community engagement with the location selection process. Meanwhile, the particular use case implemented in this first iteration of the model could help to alleviate the disproportionate risk that falls upon the rural and underhoused during these events. This analysis found that over 40% of Alabama residents live within 5 minutes of a variety store, including housing unit departure delay time. This seems to suggest that the concern would not even have to be one of travel time but rather of shelter and parking lot capacity. This percentage was admittedly higher than initially expected and truly demonstrates just how prevalent variety stores are across the nation and in the southeastern United States particularly, making them perfect candidates for a public-private partnership for community tornado shelters. The model's approach, allowable size, generalizability, and impactful use case presented here represent a meaningful contribution to the larger conversation of community tornado shelter siting. Although not without its fair share of simplification, it is hoped that this model can move the conversation toward more innovative approaches for community tornado shelter siting to fill current gaps in life safety provision for our most vulnerable communities.

5. EXPLORATION OF AN AGENT-BASED MODEL FOR POPULATION OUTMIGRATION

In the introductory literature review of this dissertation, agent-based models (ABMs) are put forth as one means by which disciplinary boundaries can be alleviated in the solution of real-world problems that are inherently interdisciplinary and internetworked. As is the case with all bottom-up approaches in modeling, emergent characteristics and the mechanisms that drive them are made more evident by ABMs. Although newer to the field of disaster research, ABMs have been used in fields such as ecology and anthropology for decades. These types of models are characterized by their approach of assigning attributes to individual agents to initialize the model and then have the agents interact with each other and their environment according to a set of defined rules and activities that mimic pertinent aspects of reality for the research question. Thus, ABMs lend themselves to all applications in which it is either more straightforward to identify the mechanisms compelling agents to act in a particular way or more fruitful to investigate these mechanisms to better extricate causation from correlation. Agents in these models can be everything from elk traversing a landscape to people trying to join the workforce. The ABM presented here have been developed in the Julia coding language with in-depth use of the “Agents” Julia package intended for ABM development (Datseris et al., 2024).

5.1. Motivations and Overview

After a disaster event, there is an increase in community instability due to losses in functionality within the physical, economic, and social infrastructure systems. This instability can lead households to temporarily dislocate, which is nearly unavoidable for some households whose housing sustained significant damage due to the event. However, when those households choose to outmigrate, which is defined as permanently dislocating to another community, then the original community begins to suffer more long-term and largely insurmountable hurdles such as losses in governmental funding, decreases in economic activity, and less tangible, although equally unnerving losses in human connection within the community. By a single name, the community loses some of its resilience. This creates a cycle of susceptibility to negative long-term outcomes, and with each hazard that occurs, they become less equipped to resist and recover. Thus, it is vital for the sake of our most susceptible communities to invest in improving community resilience. Yet community resilience is a difficult concept to definitively measure, and it is even more difficult to understand what community factors contribute, and in what amount, to this resilience. In some ways it might be easier to understand what resilience is not than what it is and then use a sort of process of elimination to determine what predicts community resilience.

One clear indication of a lack of community resilience is a high propensity for outmigration after a hazard event. If more can be understood about why outmigration occurs on a household level, a clearer aggregate understanding of outmigration may also be reached, and by inversion, a greater understanding of community resilience. This sort of household-level model would require a bottom-up modeling approach rather than a top-down modeling approach to understand and predict not only how many households will outmigrate, but importantly which households will outmigrate. In knowing more about the households that are outmigrating more actions can be taken to ensure that these households in particular are being provided with the appropriate resources and opportunities to stay in their community. Understanding more about this outmigrating population is the intended outcome of this modeling effort. In order to create a bottom-up outmigration model an agent-based modeling approach has been chosen in which the primary agents are households with secondary agents for businesses, government institutions, and schools. A top-down approach has been melded with this ABM in order to ensure that the values assigned to each agent are in alignment with actual hindcasted outmigration patterns. To accomplish this, a linear multi-regression analysis was developed to predict how many households would outmigrate and due to which factors, while the ABM was used to better understand who would outmigrate and by what mechanisms.

5.2. Methodology

One of the greatest challenges with a bottom-up modeling approach when choosing households to be the units of analysis, or "agents" in ABM, is that very little can be precisely known about each household, and their actions will not always adhere to rational behavior. To alleviate some of these challenges, the households in this model have been synthetically generated in accordance with not only community-wide census values but also correlation between these values. Second, a degree of randomness has been introduced to adequately capture the variation in a particular household's outcome. Due to the fidelity of this modelling approach, some assumptions have still been made, but these have been minimized as much as possible and are noted throughout the ABM discussion later in this chapter as they are introduced in the model. To more thoroughly roadmap the process by which this model was created, the general stages of analysis, the constituent steps, and the necessary interconnectivity of these steps is diagrammed in Figure 40.

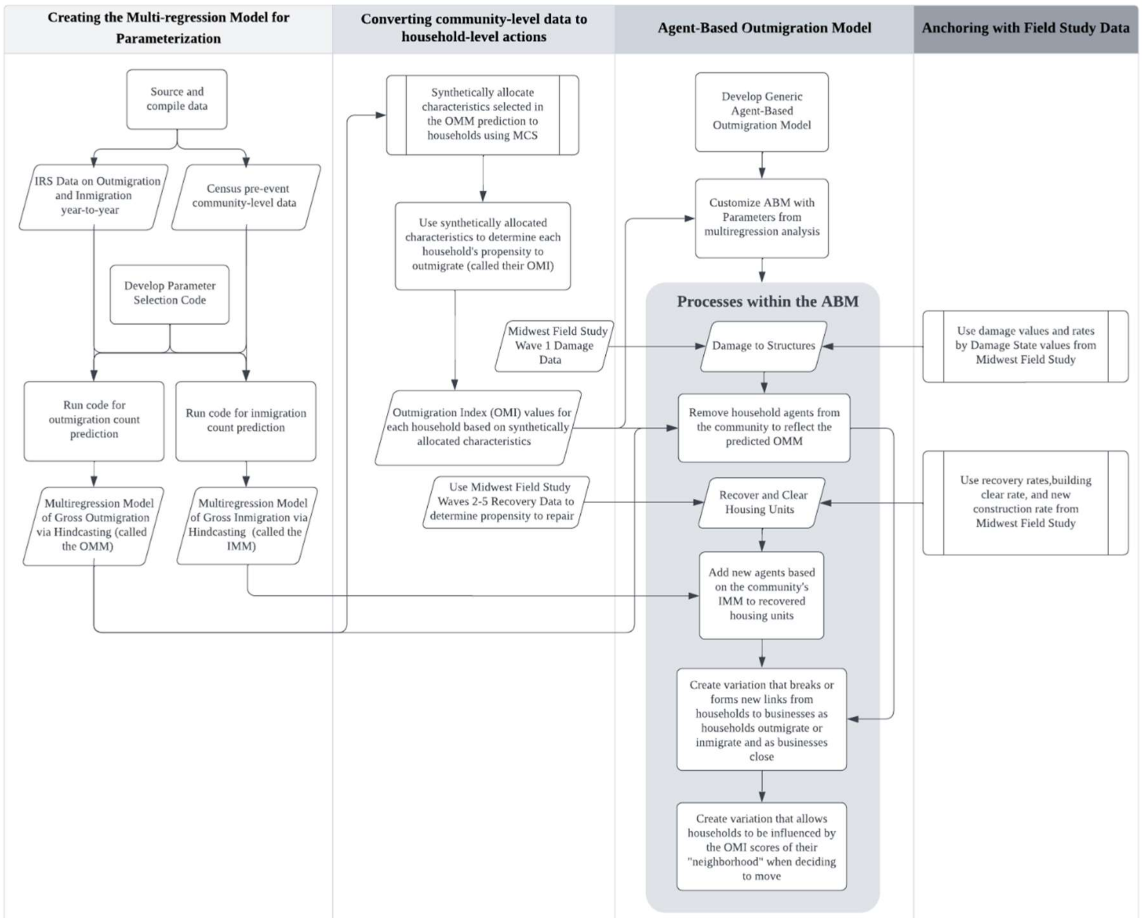

Figure 40: Agent-based outmigration model development process

5.2.1. The Multi-regression Model for Parameterization

5.2.1.1. Methodology for Multi-Regression Migration Analysis

Although the primary output of this effort is an ABM for outmigration prediction, this ABM needs to be informed by real-world data and nationally available metrics in order to ensure the results are more than just hypothetical scenarios. As such, this model integrates a parameterization of the propensity for outmigration based on a hindcasted training dataset. This is based on a multi-regression analysis analogous to the social susceptibility metric selection procedure described in Chapter 1 of this dissertation. In fact, it used the already compiled list of potential predictive factors shown in the left column of Table 1 in Chapter 2 for model training. However, this iteration of the multi-regression model had a few modifications to better suit this implementation.

First, the model was trained only to predict outmigration and inmigration patterns rather than a suite of long-term monitoring factors. In fact, these two factors were predicted in separate runs of the model to allow for further differentiated prediction of gross outmigration, gross inmigration, and by extension, net outmigration/inmigration. Although net outmigration is the primary concern when discussing outmigration more generally, it is in fact necessary for this ABM to know gross outmigration and gross inmigration predictions to simulate this flow of outgoing and incoming residents to then arrive at the net outmigration value as the model runs. This data on inmigration and outmigration of households was sourced from the yearly Statistics of Income national tax dataset (IRS Statistics of Income Division, 2020). Because this data was only available at the county level, the model only used counties, rather than counties and county subdivisions as seen in Chapter 1, to select parameters and train the model.

Second, because more accurate data was available via the SOI tax data, the outmigration predictions could be calibrated for each year rather than five and ten years post-event as with the ACS data used for the monitoring factors in the SSM development chapter. The bounds of this yearly analysis were bounded at four full years after the event. This meant that more years of hazard events could be compiled and used in the training dataset. Because post-event outmigration and inmigration data was tabulated through 2020, events occurring in 2012, 2013, 2014, 2015, 2016 had outmigration patterns tracked for four years after the event (plus Year 0, the year of the event). While more recent events such as those in 2019 were only tracked for one full year as well as the year of the event. See Table 12 for more information on this. To account for some variation year over year in national moving trends, the counties were compiled not by year of event, but by chronological analysis range of the event as shown in Table 12. For instance, all counties with an event occurring in years 2012-2020 were used as part of the training dataset for the “0 Years Since Event” outmigration and inmigration models, and so on for each model up to four years since the event. Beyond four years past the event, the sample dataset grew too small to yield generalizable results.

Table 12: The years in which hazards had to occur for a county to be included in the given model's training dataset

0 Years Since Event Model	**1 Year Since Event Model**	**2 Years Since Event Model**	**3 Years Since Event Model**	**4 Years Since Event Model**
2012	2012	2012	2012	2012
2013	2013	2013	2013	2013
2014	2014	2014	2014	2014
2015	2015	2015	2015	2015
2016	2016	2016	2016	2016
2017	2017	2017	2017	
2018	2018	2018		
2019	2019			
2020				

As in the SSM multi-regression model, every combination of factors was tested and the highest R-squared model for which individual p values remained under 0.1 was selected as the best model. The best linear model for each column in Table 12 was selected for both inmigration and outmigration. The variables chosen and the coefficients of these variables are shown in Table 13 for outmigration and Table 14 for inmigration. In the case of outmigration, there seemed to be significant noise in the parameters selected for each model, though a few factors consistently appeared. A factor is considered to be a consistently selected factor if it appears in at least a majority of the models, three or more. With this in mind, the consistent factors for outmigration were the percentage of limited English-speaking households, the percentage of unemployed females over the age of sixteen, the percentage of the population without health insurance, the percentage of families headed by a single parent, the percentage of mobile homes, the percentage of owner occupied homes, the percentage of households without a vehicle, and the population density of the community. As for inmigration, the resultant models selected a more stable group of parameters but only marginally more than the outmigration models. Setting the same standard for consistency as was used in selecting factors for outmigration, the analysis showed that in fact all the same factors selected from the outmigration models were selected for the inmigration models with the addition of median income and educational attainment. Because these consistently identified factors were chosen for the outmigration and inmigration prediction in the ABM and to avoid overparameterization and reconfiguring of the migration calculations for each year, this abbreviated list of parameters was reran as the best fit model for outmigration and inmigration to yield the results seen in Table 15. Though investigation of the changing strength and influence of different factors given different time horizons after an event would be an interesting area of further investigation.

Table 13: A summary of coefficients for the best fitting linear multi-regression outmigration model for each dataset along with the associated R-squared value

	0 Years Since	**1 Year Since**	**2 Years Since**	**3 Years Since**	**4 Years Since**
R-squared of best model	**0.46**	**0.42**	**0.44**	**0.48**	**0.54**
% with less than high school education	-1.92				
% 65 and over	-1.92			-5.75	
% with limited English proficiency		-1.45	-2.72	-3.93	
% below poverty line					-11.14
% unemployed civilian labor force	-3.44	-5.52			
% females 16+ unemployed	4.34	8.27	4.58	8.25	10.21
% population with no health insurance	1.90		6.96	6.57	
% noninstitutionalized pop w/ disability	2.36	4.78			
% single parent households	-5.84	-9.53	-15.54	-22.27	-32.58
% HUs that are mobile homes	0.38	0.72			1.80
% owner occupied HUs	-10.15	-31.66	-25.78	-33.56	-49.54
% households w/o a vehicle		-5.32	-9.66	-10.05	
GINI index					37.51
Number of health professionals per 1,000 people				-7.01	-10.76
% with associates degree or higher	-11.46	-19.20			
% with bachelor's degree or higher	2.95	6.39			
median household income	5.54	8.76			
% minority population			23.07		

% renters		-5.39			
median age		-6.97			-25.26
rurality (less than 200 people/mi^2)			2.40		
U.S. normalized population density	0.05	0.10	0.14	0.17	
size of largest minority race group (racial affinity groups)			-29.23		

Table 14: A summary of the best fitting model for inmigration for each dataset along with the R-squared associated with each model and its coefficient

	0 Years Since	**1 Year Since**	**2 Years Since**	**3 Years Since**	**4 Years Since**
R-squared of best model	***0.54***	**0.49**	**0.53**	**0.56**	**0.59**
% with limited English proficiency	-1.99	-2.49	-3.48	-4.56	-5.42
% unemployed civilian labor force					17.07
% females 16+ unemployed	3.06	5.48	7.30	10.55	
% population with no health insurance	3.50	5.20	12.44	17.21	18.48
% single parent households	-6.76	-11.44	-18.55	-25.53	-54.05
% HUs that are mobile homes		0.84	1.35	1.75	
% owner occupied HUs	-10.02	-19.40	-24.84	-30.58	-53.52
% households w/o a vehicle	-5.40	-10.81	-11.46	-14.27	-14.41
GINI index			-29.71	-35.33	
Number of health professionals per 1,000 people	-1.61	-	-	-	-12.09
% with associates degree or higher	-9.83	-13.86			
% with bachelor's degree or higher	3.18	5.81	7.10	9.45	
median household income	5.36	7.26			19.04
% minority population					5.86
U.S. normalized population density	0.10	0.16	0.22	0.28	

5.2.1.2. Results of Multi-Regression Analysis

The selected factors were isolated and re-run to identify their optimized coefficients. Also at this step, the coefficients generated were annualized by dividing each coefficient by the years since the event. As can be seen for the factors in Table 13 and Table 14

Table 14: A summary of the best fitting model for inmigration for each dataset along with the R-squared associated with each model and its coefficient

	0 Years Since	**1 Year Since**	**2 Years Since**	**3 Years Since**	**4 Years Since**
R-squared of best model	***0.54***	**0.49**	**0.53**	**0.56**	**0.59**
% with limited English proficiency	-1.99	-2.49	-3.48	-4.56	-5.42
% unemployed civilian labor force					17.07
% females 16+ unemployed	3.06	5.48	7.30	10.55	
% population with no health insurance	3.50	5.20	12.44	17.21	18.48
% single parent households	-6.76	-11.44	-18.55	-25.53	-54.05

% HUs that are mobile homes		0.84	1.35	1.75	
% owner occupied HUs	-10.02	-19.40	-24.84	-30.58	-53.52
% households w/o a vehicle	-5.40	-10.81	-11.46	-14.27	-14.41
GINI index			-29.71	-35.33	
Number of health professionals per 1,000 people	-1.61	-	-	-	-12.09
% with associates degree or higher	-9.83	-13.86			
% with bachelor's degree or higher	3.18	5.81	7.10	9.45	
median household income	5.36	7.26			19.04
% minority population					5.86
U.S. normalized population density	0.10	0.16	0.22	0.28	

, the coefficients grew with each year, which was simply a feature of the training dataset that summed the outmigration and inmigration percentages for each year to not lose any long-term trend information. Conversely these percentages could have been annualized, but it was simpler to annualize the results. These results are presented for outmigration, inmigration, and net migration (a simple difference taken between the inmigration and outmigration values) in Table 15. As one might imagine, the consolidation of factors has led to smaller R-squared values. This R-squared reduction is more pronounced for outmigration because outmigration had less consistency across the factors in its original best-fit models. This is however where the ABM can help fill in some missing information on building status and offer greater explanatory power as outmigration is not determined solely by social factors. Conversely, the strength of the inmigration linear models is beneficial because characterization of inmigration in the ABM is largely bound by the results of the linear model result. These coefficients shown in Table 15 are averaged for outmigration and inmigration to yield equations (4) and (5), which are subsequently incorporated into the ABM. The details of how these social factors are calculated for each household and then put into these equations are detailed in the ABM discussion later in this chapter.

Ahead of the ABM discussion though, it is interesting to scrutinize the results in Table 15. In the original formulation of this linear model, it was presupposed that some social factors long believed to be tied to poor resilience would have more clearly negative consequences on outmigration and/or inmigration projections. However, because many of the same factors have been selected for both inmigration and outmigration models, a different narrative emerges by which factors seem to be associated with either stabilization or destabilization of a community's population. Because all the factors that appear in both sets of models have the same sign for inmigration and outmigration, one can see that factors with negative coefficients lower the migration rates in the community (stabilization), while factors with positive coefficients increase the migration rates in the community (destabilization). The factors associated with stabilization and destabilization are summarized in Table 16. The mechanisms driving a given factor to belong to one column or the other in this table can certainly be hypothesized. At this point, this discussion will be tabled by acknowledging that lived experience and logic seems to explain the majority of these factors' association with either stabilization or destabilization of a community's population, and the ABM to follow will delve more into this. The net migration columns are also included Table 15 to draw attention to possible summary trends for each factor, but because the inmigration model has two additional factors which may have correlation with the other factors, the net migration linear model should not be used as anything more than an illustrative tool as to the general impact of the factors outlined here.

With empirically derived equations for outmigration and inmigration, the focus of this work now shifts to the ABM portion of the analysis. The aggregated gross outmigration prediction shown in (4), called the OMM (outmigration metric), will be synthetically downscaled to individual households in the following model. At the individual household level this value will become their Outmigration Index (OMI), and this will inform their likelihood of outmigrating along with other factors like in ability to find adequate housing or employment. The inmigration metric, or IMM, defined in (5) will be used at the community level as a primary guide for household agent "creation" and placement in the community space at each timestep to counteract the outmigration that is simultaneously occurring.

Table 15: Results of multi-regression analyses for only selected parameters – outmigration, inmigration, and net effect for each year after the event

	Variable for (4) & (5)	Annualized Outmigration (+ indicates outmigration)					Annualized Inmigration (+ indicates inmigration)					Net Population Impact (+ indicates net increase in population)				
Years Since Event Model		0	1	2	3	4	0	1	2	3	4	0	1	2	3	4
R-squared for Model		0.37	0.35	0.41	0.43	0.49	0.52	0.47	0.51	0.55	0.60					
Constant		17.93	19.32	20.94	20.47	23.43	15.63	15.59	16.11	15.01	16.79	-2.30	-3.72	-4.83	-5.46	-6.64
% Limited English	*l*	-0.66	-0.72	-0.66	-0.48	-0.31	-1.33	-1.05	-1.07	-1.00	-0.81	-0.67	-0.32	-0.41	-0.52	-0.50
% Females 16+ unemployed	*f*	1.50	1.58	1.67	1.86	1.91	2.22	2.23	2.35	2.71	3.07	0.71	0.65	0.68	0.85	1.16
% Population with no health insur.	*h*	2.89	2.27	1.87	1.79	1.55	3.83	3.07	3.42	3.60	3.64	0.94	0.80	1.55	1.81	2.09
% Single Parent HHs	*s*	-4.42	-4.30	-5.04	-4.98	-5.71	-6.32	-5.13	-6.24	-6.51	-7.85	-1.90	-0.83	-1.20	-1.53	-2.14
Mobile Homes as % of HUs	*m*	0.18	0.17	0.24	0.22	0.28	0.55	0.45	0.39	0.40	0.60	0.36	0.28	0.15	0.18	0.32
% Owner Occupied	*o*	-8.04	-8.80	-9.84	-9.79	-11.97	-9.03	-8.04	-8.05	-7.61	-10.48	-0.99	0.76	1.78	2.18	1.49
% HHs w/o vehicle	*v*	-3.55	-3.53	-2.97	-2.82	-2.52	-4.39	-5.02	-4.18	-3.84	-2.80	-0.83	-1.49	-1.21	-1.02	-0.29
% with Bachelors degree or greater	*e*						-0.52	0.68	0.68	0.74	0.96	-	-	-	-	-
Median Household Income	*i*						4.74	2.68	2.14	2.12	3.27	-	-	-	-	-
U.S. normalized population density	*p*	0.07	0.04	0.04	0.04	0.03	0.10	0.08	0.07	0.06	0.04	0.03	0.03	0.03	0.03	0.01

When averaged, the results for the five linear models for outmigration and inmigration shown above are as follows:

$$OMM = 20.42 - 0.57l + 1.71f + 2.07h - 4.89s + 0.22m - 9.68o - 3.08v + 0.04p \quad (4)$$

$$IMM = 15.83 - 1.05l + 2.52f + 3.51h - 6.41s + 0.48m - 8.64o - 4.05v + 0.51e + 2.99i + 0.07p \quad (5)$$

Table 16: The migration trend associated with each factor that appeared in both migration models

Associated with more migration stabilization	**Associated with more migration destabilization**
Limited English	Females 16+ unemployed
Single Parent HHs	Population with no health insurance
Owner occupied	Mobile Homes as % of Hus
HHs w/o vehicle	US Normalized pop density

5.2.2. The Agent-Based Models

5.2.2.1. THE MODELING SCENARIO

Due to the wealth of data that has been produced via the Midwest Field Study as described in Chapter 2, the city of Mayfield was selected as the initial community for which this outmigration model would be generated. While the social data and the linear model for outmigration and inmigration are pivotal to the parameterization of the ABM, it is also necessary to have information on the building stock in which households will live, work, and connect. The data that are necessary inputs for this model include the anticipated damage due to the hazard event; the projected repair rate and probability of building debris being cleared away based on that anticipated damage; as well as the rate of new construction starts. Because this initial model is being hindcast following an actual event, these variables are derived from real-world data to compare directly with actual outmigration patterns in the results section of this chapter, but in scenario analyses for future events, the code could also utilize the outputs of IN-CORE analyses since it has been written and developed to ingest and be informed by IN-CORE damage states.

5.2.2.2. THE SPACE

The agents in an ABM must be placed into a space in which they can perform actions, interact with other agents, and make decisions. These spaces can take the form of a road network as seen in the last chapter, a grid space, a completely continuous space, or any number of other spaces aimed at imitating a physical or non-physical "space." For this implementation, it is necessary for the agents to exist in a space that is not preconfigured by the Agents Julia package. This space can best be thought of as a collection of bins with each bin representing a building in the community and identified by its GUID field. Most bins can only store one agent with the exception of buildings that have been flagged as apartments, which are assigned a random number of slots between two and sixty within their bin to simulate each apartment unit. This value of sixty has been estimated for the community of Mayfield which has been analyzed to have smaller apartment buildings unlikely to house more than sixty families. This value could certainly be adjusted for larger cities with larger apartment buildings.

Although these bins do store information on where that particular structure is physically located, the space of this model is not an x-y plane in which these structures are placed at their x-y position. This was done to make the model more performant. Early iterations of this model allowed agents to operate in a x-y plane on which structures were located, but this led to poor performance because the model had to waste memory storing information about space that played no role in this model and was not inhabitable by agents, such as parks, roadways, and undeveloped land. A simplified visualization of this space consolidation and the resultant space can be found in Figure 41.

Figure 41: Schematic visualization of the conversion of building data from the x-y continuous map space (left) to discrete bin space (right) for a section of Mayfield. The legend here shows the archetypes of these structures (Memari et al., 2018)

Aside from the latitude and longitude of each building, each bin in the defined space also stores information about its archetype and the associated owner type for that archetype. This prevents agents from moving into inappropriate buildings, meaning no households will inhabit a school, no businesses will inhabit a single-family home, etc. Although there are situations where the latter happens in communities, this idealized model does not allow for this. Further exploration on this point and modification of this habitation rules could lead to discussions on how zoning changes affect recovery following a hazard, though that is outside the scope of this analysis. Other space variables include its initial damage state, occupancy status, school zones, the income of its original resident, recovery status, and flags for new constructions. How these parameters are used will be discussed later in the methodology discussion. To review the complete list of space factors and their meaning, see Table 17.

Table 17: Space-level Parameters

Model Parameter	Parameter Description
s	agent id(s) in the building
guid	guid of the building
archetype	IN-CORE building archetype. This model at present assumes the archetype to be based on IN-CORE's suite of tornado archetypes.
x	longitude of the building's centroid
y	latitude of the building's centroid
damagestate	initial damage state of the building due to the event
occupied	occupancy status of the building
elemschoolzone	elementary school zone that the building sits within
midschoolzone	middle school zone that the building sits within
highschoolzone	high school zone that the building sits within
n_householdsinstruct	number of households in the building
n_husinstruct	number of housing units in building
damage	Updated damage state after repairs have been completed
cleared	Boolean noting if a building has been cleared away or bulldozed
constructed	If the building is constructed during the model run period, this field indicates with a string value when that construction occurs. Otherwise, it will be listed that this building is an original structure.
constructtick	This integer field converts the "constructed" field into a time value in alignment with the units of the model's tick value.
apartment	Indicates if this building is an apartment.
ogincome	Stores the income of the original resident's income. Will be 0 if the no resident is assigned during initialization
ownertype	Deduces from the "archetype" and "apartment" fields which type of agent can occupy this building.
buildingfail	This field indicates whether the building will maintain functionality by 2 years after the event. Buildings with damage states of DS2, DS3, and DS4 either undergo repairs or fail, while DS1 and No Damage buildings can maintain functionality without repairs.

5.2.2.3. The Model Parameters

Primarily for the purpose of monitoring and bounding the model, there are also over-arching model parameters that are defined. These parameters mainly consist of basic values included in most models for housekeeping, places to store model counts, and user-defined input values. The most rudimentary of these parameters is the "tick" value, this value monitors the time that has passed since the event and defining the discrete amount of time that each "tick" represents is how the temporal scale of the model is defined. For this model, the temporal scale is set to a month, and the model is allowed to run for 48 months. This time limit is defined as the "simulationperiod" and is set in accordance with the time scale used in the linear migration models. The final housekeeping parameter included is the "progress_bar" which allows for the model runs to show their progress in an easily intelligible way in the terminal.

The model counts are stored for the model at n_structures, n_unoccupied_structures, n_occupied_structures, n_outmigrated, n_inmigrated, household_number_searching, n_unhoused_households, n_housed_households, and n_total_households. Where possible, these parameters are given initial values in alignment with census data values. These include n_housed_households and n_housingunits. The rest of the parameters are specific to this ABM and will vary for every community and need to be provided via an input parameter csv file. These parameters include the previously mentioned number of households and number of housing units, as well as the annualized inmigration and outmigration percentages for a given community, and the chosen definition of outmigration by time dislocated. As will be described in more depth later, outmigration can occur due to a relatively high OMI for a household, job loss, or because the household has been dislocated from their home for an untenable amount of time with no ability to find another suitable home. These parameters are listed in Table 18.

Table 18: Model-level Parameters

Model Parameter	Data Source	Parameter Description
tick	Model controlled	Time counter for the model (units = months)
simulationperiod	User defined	The length of time each simulation should run.
t_outmigration	User defined	The minimum time an agent needs to be dislocated before they can outmigrate due to dislocation alone.
progress_bar	Model controlled	An internal parameter for displaying model run progress.
n_structures	Model counted	The total number of buildings in the dataset.
annualoutmigration	User calculated	The input value for OMM from the linear regression calculation provided in (4).
annualinmigration	User calculated	The input value for IMM from the linear regression calculation provided in (5).
n_unoccupied_structures	Model counted	The number of empty buildings.
n_occupied_structures	Model counted	The number of occupied buildings
n_housingunits	From Census to initialize then model counted	The total number of housing units from the input building data set, including units that are projected to be constructed during the model run period.
n_schools	Model counted	The number of schools in the community.
n_outmigrated	Model counted	The number of households that have left the community permanently.
n_inmigrated	Model counted	The number of new households that have entered the community.
household_number_searching	Model counted	The number of households looking for a new place to live. In its present configuration, this is identical to dislocated.
n_unhoused_households	Model counted	The number of households who have not been able to find a place to live and are thus either dislocated or outmigrated.
n_housed_households	Model counted	The number of households in housing units within the community.
n_total_households	From Census to initialize then model counted	The total number of households including outmigrated and inmigrated households to capture the total number of households that have moved through the model space during the simulation period.

5.2.2.4. THE TYPES OF AGENTS

Households

The primary agents of interest in this model are the households. All other agents only exist in this model to interact with and influence the household agents to stay or leave the community. The households are agents that are assigned to housing units as long as sufficient housing units are available. The number of household agents created is in alignment with the household count provided by the U.S. Census for the year prior to the event. The parameters for households are included in Table 19. The derivation of all social factors shown in the table will be explained in the later discussion of model initialization.

Table 19: Household-level Parameters

Parameter	Description
id	Agent’s unique identifier in the ABM
pos	The building that the agent occupies which comes from “guid” in space dataset
pos_idx	This is the index of the agent’s position in the space dataset. It allows for crosswalk of information from agent to space and vice versa.
education	This parameter reflects whether anyone in the household has attained a bachelors degree or higher. 1 if yes, 0 if no.
income	This parameter reflects the estimated household income based on the value of the housing unit in which they reside.
limenglish	This parameter reflects whether anyone in the household can speak English very well. 1 if no, 0 if yes.
unemployedfemales	This parameter reflects if the household has an unemployed female over the age of 16 who is in the labor force. 1 if yes, 0 if no.
nohealthinsurance	This parameter reflects whether the household has health insurance coverage. 1 if no, 0 if yes.
singleparent	This parameter reflects whether this household is a family run by a single parent. 1 if yes, 0 if no.
inmobilehome	This parameter reflects whether this household resides in a mobile home. 1 if yes, 0 if no.
ownedhome	This parameter reflects whether this household owns their housing unit. 1 if yes, 0 if no.
novehicle	This parameter reflects whether this household has access to a vehicle. 1 if no, 0 if yes.
popdensity	This is a community-wide value used in the OMM and IMM calculations. This value is calculated as the people per square mile.
OMI	The OMM value discussed above but downscaled here to calculate the outmigration index for each household.
OMIpercentile	The OMI percentile is the percentile into which a given agent’s OMI value falls when pooled with all other resident’s OMI values.
looking_to_move	This Boolean determines whether a household agent wants to find a new housing unit. In the current formulation, this is akin to “dislocated,” but is built-in so as to allow other versions of the code to simulate typical housing relocation.
housing_failure	This parameter is a Boolean tracking the failure of the housing unit in which the agents lives, or lived if they have dislocated. Failure here tracks those buildings that will not receive repairs in the two years following an event after experiencing initial damage of DS 2 or higher. This will mean the home will become uninhabitable due to immediate damage or damage exacerbated by deferred repairs.
eschool_closure	This parameter tracks whether the household’s zoned elementary school is experiencing a damage-related closure. This is not pertinent in the current version but is included where data is available on school closures.
mschool_closure	This parameter tracks whether the household’s zoned middle school is experiencing a damage-related closure. This is not pertinent in the current version but is included where data is available on school closures.
hschool_closure	This parameter tracks whether the household’s zoned high school is experiencing a damage-related closure. This is not pertinent in the current version but is included where data is available on school closures.
dislocated	This Boolean tracks if a household’s original housing unit failed resulting in the agent being move to the “none” position. These agents have not yet left the model

	space and may still have jobs in the community while they look for permanent housing in he community.
es_zone	This notes the elementary school zone the household lives in.
ms_zone	This notes the middle school zone the household lives in.
hs_zone	This notes the high school zone the household lives in.
outmigrated	This Boolean variable tracks if a household has permanently left the community and as a result all their interactivity with buildings and other agents in the community is turned off.
timeoutmigrated	This parameter counts how many months a household has been outmigrated.
timedislocated	This parameter counts how many months a household has been dislocated.
repairing	This Boolean is used to indicate if a household is repairing or plans to repair their housing unit within the first two years following the hazard event.
repaired	This Boolean variable tracks the whether repairs have been completed on the housing unit in which the household lives.
ogbldgDS	This offers agents a place to store information on their initial building's damage state. For new agents, this value will be set to "inmigrated" since they have no original structure and to track inmigrated houses more effectively.
workers	This stores the number of workers (0,1,2) that a household has. It is assigned randomly with consideration given to the unemployed female variable.
lookingforwork	This is an integer noting how many workers in the household are without work.

Businesses & Institutions

Businesses and Institutions operate identically for this iteration of the model, so they have been grouped together here for the sake of concision. In other iterations of this tool, it may be necessary to describe each of these agents independently because they do serve different purposes in the community. Businesses in this model encompass everything from small shops to office buildings. Because businesses are not the primary focus of this model, it is assumed that each business building has just a single business in it that will choose to close or stay open based on damage and the availability of workers. Institutions follow these same rules, these institutions are the agents that occupy municipal buildings, churches, and community center spaces. Institutions are thus well suited to provide post-event resources and may not be subject to closure decisions identical to businesses. However, this differentiation in operation and closure mechanisms for these two types of agents is not within the scope of the current model because little is known about the quantitative influence of institution closures on a household's tendency to outmigrate. The number of business and institution agents created in this model is dictated by how many suitable buildings exist in the community in accordance with the archetype mapping in Table 22. The number of workers needed by a business or institution is set randomly at a value between 1-10 for businesses and 1-15 for institutions. Then the total worker count is checked against census values pre-event, and more jobs are assigned or jobs are removed in order to force fit a recreation of the census values for employed workers and unemployment rate. These worker counts could be refined by each archetype in future work. The complete list of parameters assigned to agents and institutions is shown in Table 20.

Table 20: Business-level and Institution-level Parameters

Parameter	Description
id	Agent's unique identifier in the ABM
pos	The building that the agent occupies which comes from "guid" in space dataset
pos_idx	This is the index of the agent's position in the space dataset. It allows for crosswalk of information from agent to space and vice versa.
building_failure	This parameter is a Boolean tracking the failure of the building in which a business or institution is located. Failure here tracks those buildings that will not receive repairs in the two years following an event after experiencing initial damage of DS 2 or higher. This will mean the building will not maintain adequate functionality for business operations due to immediate damage or damage exacerbated by deferred repairs.
n_workers	This tracks the number of workers employed by the agent.
n_workersneeded	This tracks the minimum number of workers needed for the agent to operate in an understaffed capacity without closing.
open	The Boolean stores where the agent is open or closed
timeclosed	If closed, this integer will track how long the agent has been closed and looking for a new location.
repairing	This Boolean is used to indicate if a household is repairing or plans to repair their housing unit within the first two years following the hazard event.
repaired	This Boolean variable tracks whether repairs have been completed on the housing unit in which the household lives.
ogbldgDS	This offers agents a place to store information on their initial building's damage state. For new agents, this value will be set to "inmigrated" since they have no original structure and to track inmigrated houses more effectively.
workerids	This vector stores the ids of the households that have workers employed with this business or institution.

Schools

Schools are only minimally explored in the first implementation of this model. People establish ties to their local school via school zones, but the impact of school functionality loss on outmigration patterns could not be adequately explored in this first implementation of the tool because none of the schools in Mayfield were damaged significantly by the tornado. These agents have nevertheless been included so that the impact of a household-school linkage can be explored further in potential expansions of this model.

5.2.2.5. THE INITIALIZATION

The model will step through the following functions to begin the simulation. A series of callable python functions and built-in Agents.jl functions will parameterize the model and develop the model's space in accordance with a user-provided building dataset file and model input parameter file. Then in the "AddAgents_fromCommunityDataframe" function, the model will allocate households to housing units as well as businesses, institutions, and schools to their respective building types. Workers will be assigned to businesses and institutions directly in the initialization function. In "CheckStructuralConditions," agents will determine if their structure is failing, and based of this failure status, "check_repairing" will allow them to determine whether they plan to repair the building. Then households will determine their propensity to outmigration (OMI) in the "EstablishOutmigration" function. Then, the model will count all entities in the "UpdateModelCounts" function before closing the initialization. More information on each of these functions is provided below.

Function 1: AddAgents_fromCommunityDataframe

Importing the Necessary Data

In order to define the space in which agents will operate, a csv characterizing the community's buildings needs to be imported and augmented to possess the necessary information. This imported dataset requires the fields outlined in Table 21. This import is conducted in a helper python file in which the data is read into a data frame and the percentile of each housing unit's value as well as the building's owner type are determined, except in the case of apartments. The former of these will be essential in calculating household income, and by extension their OMI, later in this initialization and the latter is mapped according to Table 22 to determine which types of agents can live in a given building. For school archetypes, the model will generate two agents. The first will be the school agent which provides community services to school-aged children, and the institution agent will provide jobs for teachers, the administration, and support staff. This building's dual purpose is reflected in Table 22. This python file then returns the community data frame to the primary file to proceed with model initialization.

Table 21: Input Building Dataset Fields and Descriptions

Building Dataset Fields	Description
guid	A unique identifier for each building
archetype	The tornado building archetype as defined in IN-CORE
DS	The initial damage state of the building as defined in IN-CORE
val_struct	The structure's value according to the NSI
name_es	The elementary school zone for each household
name_ms	The middle school zone for each household
name_hs	The high school zone for each household
x	The longitude of the building's centroid
y	The latitude of the building's centroid
Constructed	A string indicating whether the building is an original structure in the building or a new construction. In the case of this example, if it is new, this field indicates at each wave of data collection the building was first completed
apartment	This field flags if the building is an apartment building. In this example, this data was deduced from the EF-scale Damage Indicators, which flag apartment buildings as DI 5s.
Constructtick	An integer indicating at what timestep a newly constructed building will become habitable. Can be calculate with conditional from "Constructed"
school_type	For school buildings, this indicates what type of school it is: elementary, middle, or high school

Table 22: Building archetypes mapped to owner types

Tornado Archetype	Owner Type Mapping
1 residential	Household
2 residential	Household
3 residential	Household
4 residential	Household
5 residential	Household
6 strip mall	Business
7 light industrial	Business
8 heavy industrial	Business
9 elementary and middle schools	School; Institution
10 high schools	School; Institution
11 police and fire stations	Institution
12 hospitals	Business
13 community centers / churches	Institution
14 government buildings	Institution
15 large box stores	Business
16 small box stores	Business
17 manufactured housing units	Household
18 shopping centers	Business
19 office buildings	Business

***Exception: If the Damage Indicator for the building is listed as “apartment,” the building’s archetype is ignored, and the apartment is allocated households.**

The other requisite data input for this model to run is a model parameterization csv file. This rather simple file provides the space for the user to input their desired model parameters as well as the necessary community dependent variables. This file allows for greater interactivity and varied scenario development to provide greater insights regarding the mechanisms that impact outmigration. These variables are grouped and characterized in Table 23.

Table 23: Model Input File Variables

Variable	Value	Description of Variable
	User Defined Variables	
t_outmigration	6	months until dislocated household is offically outmigrated
n_sims	5	number of simulations
simulationperiod	30	months until simulation ends
seed	4062	Seed used in stochastic simulation to allow for result replicability
NeighborInfluence	0.5	Here users pick a value between 0 and 1 for how much neighbors influence each other to move based on their OMI. 0 means individuals will only consider their OMI. This factor only appears in a single version of the model presented towards the end of the methodology section, namely the Neighborhood Influence Model.
	Community Dependent Variables	
household_number_searching	0	Place to store the number of people pre-event who are looking for a new place to live in the community.
n_schools	3	The number of schools in the community
n_unhoused_households	0	Place to store unhoused households – can be assigned to exist prior to the event as well to capture pre-event homelessness etc.
n_housed_households	3881	Value Obtained from Census Data
n_housingunits	4388	Value obtained from Census Data
AnnualOutmigration	12.16	OMM prediction
AnnualInmigration	9.18	IMM prediction
Income1	19.2	Percentage of people making less than $9,999
Income2	27.5	Percentage of people making less than $14,999
Income3	39.2	Percentage of people making less than $24,999
Income4	48.6	Percentage of people making less than $34,999
Income5	65.9	Percentage of people making less than $49,999
Income6	84.6	Percentage of people making less than $74,999
Income7	92.9	Percentage of people making less than $99,999
Income8	97.6	Percentage of people making less than $149,999
Income9	99.2	Percentage of people making less than $199,999

Income10	99.9	Total percentage of people receiving income (set in the model to 100%)
LimEnglish	0.0309	The rate of households with limited English proficiency from the Census
OwnerOcc	0.515	The rate of owner-occupied housing units from the Census
PopDens	66.10	The population density of the community

Determining OMI for Each Agent

It is necessary during initialization of the model to synthetically assign social characteristics to households that are required for generating each household's OMI value, or the score associated with their probability to outmigrate. These factors could have been assigned randomly to fit the known aggregate values for the community based on census data, but this would neglect any correlation between the factors. So a separate correlation analysis was conducted to determine if the relationships between the parameters selected earlier were strong enough to merit tying these values together. The R-squared values for these relationships are shown in Table 24. In order to decide what the R-squared threshold would be for consideration of a relationship, the social science literature was reviewed and academic consensus did not seem to be reached. Although some sources cited R-squared values as low as 0.10 as being acceptable (Ozili, 2022), a reasonable threshold of 0.30 was set here to ensure that the relationships between these variables did not become overly convoluted. Those relationships with an R-squared of greater than 0.30 are shown in bold in Table 24.

Table 24: R-squared values characterizing the strength of the relationships between the factors previously selected for OMM and IMM calculation

	2	3	4	5	6	7	8	9	10
1 - % Limited English	0.001	0.134	0.001	0.008	0.069	0.012	0.000	0.028	0.014
2 - % Females 16+ unemployed		0.166	**0.383**	0.176	0.014	0.261	0.138	0.254	0.000
3 - % Population with no health insur.			0.227	**0.310**	0.000	0.027	0.256	0.155	0.036
4 - % Single Parent HHs				0.079	0.108	**0.379**	0.150	0.232	0.000
5 - Mobile Homes as % of HUs					0.116	0.002	**0.402**	0.264	0.150

6 - % Owner Occupied	0.222	0.059	0.029	0.061
7 - % HHs w/o vehicle		0.054	**0.323**	0.046
8 - % with Bachelors degree or greater			**0.472**	0.199
9 - Median Household Income				0.054
10 - U.S. normalized population density				

The factors that were shown to have a strong enough relationship were then used to generate a parameterization chain starting from a known value. To do this, the NSI-provided structural value of each residential building is processed to generate each residential building's value percentile across the community. Assuming at this point that most households occupy a building that is proportionate to their income relative to the rest of the community, this model links the percentile of the structural value of a home to the percentile of the household's income. From there, census data on income distribution for a given community can be pulled to determine what percentage of the community fall into each income category and thus what income bin corresponds to a given household's income percentile. For example in the case of the income distribution from U.S. Census data for Mayfield that is shown in Figure 42, a household that was assigned to a housing unit whose structural value falls into the 45th percentile of structural values across Mayfield will be assigned a random income within the bin corresponding to the 45th percentile of household incomes. Thus, the household will have an income between \$25,000 and \$34,999. As in the SSM calculation chapter, this income value as well as all other social values discussed here were divided by the U.S. national average for a given year to make the values more independent of time and have values for each social parameter that are on roughly the same scale.

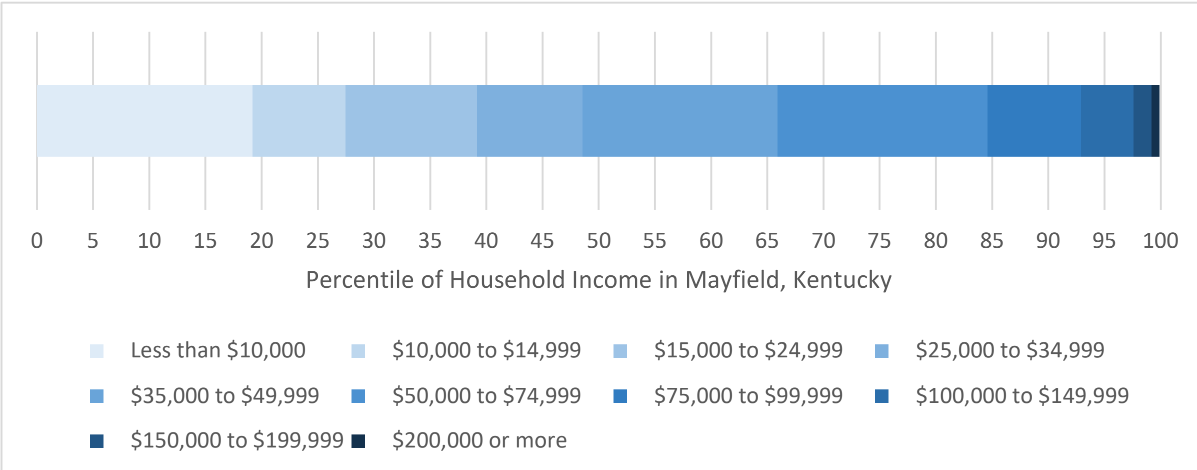

Figure 42: Mayfield income distribution from census data

Once an income is established and standardized against the national average, then many of the household's other nationally standardized social characteristic values can be assigned in accordance with the equations outlined in Table 25. For each parameter, the known parameter is used to obtain the mean conditional value and this mean value becomes the center of a conditional normal distribution with the conditional standard deviation mentioned in the table as well. As an example the mean value for the U.S. normalized probability of not having a vehicle is calculated from the U.S. normalized income in (6). This mean value is then used in tandem with n, a normally distributed random variable of mean zero and standard deviation one, to construct one possible value of the U.S. normalized no vehicle probability for each household to run one realization of the outmigration model as shown in (7). This chaining of parameter calculations is done for those parameters which were shown to have a strong relationship. Then, the parameter's probability is arrived at by taking the U.S. normalized output for each parameter and multiplying by the U.S. nation-wide probability as in (8). Then, for those factors which are Boolean, such as households without a vehicle, a random variable r is sampled from a uniform distribution [0,1) and is compared against the parameter probability to determine if the household will be assigned a value of true (1.0) or false (0.0) for that parameter, which is demonstrated in (9). For the determination of limited English households and households that own their home, which did not have a strong enough relationship with any of the other factors, each household was simply assigned a parameter probability equal to the community's overall percentage of limited English speaking households and households that own their housing unit. Then to determine their Boolean value, (9) is applied to these two variables as well. The probabilistic approach for parameterization introduced at this stage allows for virtually boundless variations of the model to generate a distribution of outmigration results that explore the certainty of these results given the uncertainty of the input parameters.

Table 25: The relationships used to calculate the related social parameters of a household

	Equation of the line (μ)	**Conditional standard deviation (σ)**
Education versus Income	$y = 1.0903x - 0.1812$	0.2422
Mobile Home versus Education	$y = -2.6493x + 3.979$	1.0774
No Health Insurance versus Mobile Home	$y = 0.1558x + 0.8904$	1.5578
No Vehicle versus Income	$y = -0.6793x + 1.3301$	0.2067
Single Parents versus No Vehicle	$y = 0.6115x + 0.5791$	0.1966
Unemployed Females versus Single Parents	$y = 0.868x + 0.1456$	0.2748

$$\mu_{US\ no\ vehic\quad probability} = -0.6793 \times income + 1.3301 \tag{6}$$

$$U.S. normalized\ no\ vehicle\ probability = 0.2067 \times n + \mu_{US\ no\ vehicle\ probability} \tag{7}$$

$$parameter\ probability = U.S. normalized\ parameter\ probability \times National\ probability \tag{8}$$

$$Boolean\ of\ parameter = if(parameter\ probability > r, 1,0) \tag{9}$$

In the case of apartments, the structural value from NSI would not reflect the value of a single housing unit, a simpler rule was imposed by which it was reasonably assumed that apartment residents would fall within the first three bins shown in Figure 42, so the household was randomly assigned an income of less than $24,999. This may need to be further refined in later iterations of the model or adjusted when considering communities with a higher cost of living.

Function 2: CheckStructuralConditions

In this function, agents will experience a simulation of the hazard event and assess their damage, deciding whether to repair or dislocate. The propensity to repair is established in accordance with the repair rates resulting from the Midwest Field Study data for the two years following the event, which has been compiled in Table 26.

As the table indicates, these repair rates have been determined for each initial damage state and are thus allocated to agents and their corresponding buildings to align with the appropriate conditional probability based on the building's initial damage state. If the agent randomly chooses not to perform repairs on their building based on these probabilities, the agent will decide based on the initial damage state if this structure is failing to meet their building needs. If the damage state is "No Damage" or "DS 1 (Slight)," then not performing repairs is unlikely to lead the building to being uninhabitable in two years, while for higher damage states it is likely to inhibit proper functionality of the building. Thus for the higher damage states, the agent will set its "housing_failure" in the case of households, or "building_failure" otherwise, to true and proceed to the following function.

Function 3: check_repairing

For this function, the agents will use the housing_failure/building_failure value assigned in the previous function to determine if repairs will be performed at any point in these first two years. Due to the formulation of building failures discussed above, the repair probability for non-failed buildings two years after the event will be 100% for buildings of initial damage state DS 2, DS 3, or DS 4. Meanwhile, for DS 1 buildings, the repair probability will be set in alignment with the converse of value provided in Table 26 (86.55%), and buildings that are listed as "No Damage" according to the IN-CORE damage states will be assumed to not repair for simplicity as no damage is listed in the data even though the agents may be conducting minor gutter repairs, etc. Once the conditional repair probability based on building failure status has been determined, the agent will either set their "repairing" value to true or false to indicate that they plan to perform repairs at some point during the first two years after the event.

Table 26: Repair Rate and Failure Status by Initial Damage State based on Midwest Field Study Data

Initial Damage State	Midwest Field Study Data: Unrepaired buildings two years after the event	Modeled: Housing failure status	Modeled: Repaired status two years after event assuming housing failure is ***False***
No Damage	0.68%	0% True 100% False	0% Repaired 100% Unrepaired
DS 1 (Slight)	13.45%	0% True 100% False	86.55% Repaired 13.45% Unrepaired
DS 2 (Moderate)	21.64%	21.64% True 78.36% False	100% Repaired 0% Unrepaired
DS 3 (Extensive)	45.97%	45.97% True 54.03% False	100% Repaired 0% Unrepaired
DS 4 (Complete)	92.31%	92.31% True 7.69% False	100% Repaired 0% Unrepaired

This function triggers a nested function before closing. This subfunction, called "clear_building_check," determines if a given building that is not performing repairs will be cleared in accordance with Midwest Field Study clearance rates, shown in Table 27. If the building is selected to be cleared based on its repair status, the probabilities tabulated below, and a degree of randomness, then the building will toggle it's "cleared" parameter to true and no longer be a viable building for inhabitants.

Table 27: Clearance rates assuming no repairs are undertaken

Wave1 DS	**Wave 1**	**Wave 2**	**Wave 3**	**Wave 4**	**Wave 5**	**Wave 6**
model.tick	1	3	6	12	18	24
No Damage	0.00%	0.17%	0.08%	0.28%	0.14%	0.00%
DS 1 (Slight)	0.00%	9.28%	2.06%	1.03%	3.09%	2.06%
DS 2 (Moderate)	0.00%	32.58%	1.12%	10.11%	4.49%	3.37%
DS 3 (Extensive)	0.00%	63.27%	12.24%	7.14%	1.02%	3.06%
DS 4 (Complete)	8.47%	75.14%	5.46%	8.20%	0.00%	0.55%

Function 4: EstablishOutmigration

Household OMI values are used to calculate the new community-wide OMM as well as each households OMI percentile in this function. All non-outmigrated household agents submit their OMI for tabulation and first they are averaged to attain the new timestep's OMM value. Then the tabulated list is sorted and each agent queries the sorted list of all OMIs in the community to determine into which percentile their OMI falls. This is done during initialization as well as at the end of each timestep as the overarching OMM and these percentiles will remain dynamic since agents inmigrate and outmigrate.

Function 5: UpdateModelCounts

This general housekeeping function provides counts for all applicable model-level parameters shown in Table 18. This function provides the primary means by which the outputs of the model are tracked at every time step. Th result of this function constitute the data summaries that are analyzed in the results section of this chapter.

5.2.2.6. THE STEPS

Household Steps

At each timestep, the household agents who have not outmigrated will proceed through the following functions. First, they will determine if they need to dislocate based on structural damage in the "dislocate" function. Then, in the "find_new_house" function, dislocated households will try to find a housing unit that has not experienced structural failure in their price range. Those households that did not dislocate and who have decided that they will perform repairs, will determine in the "perform_repairsHH" whether they will be in the portion of those who completed their repairs during this timestep. The "clear_building_check" function will then determine if the building will be cleared at the given timestep if no repairs are underway, in alignment with Table 27. Next, the households will perform counting tasks in which they tally how long they have been dislocated or outmigrated if they are in one of these states. Then the step ends for this agent. Those functions which have not previously been explained are detailed below.

Function 1: dislocate

At the beginning of every timestep, each household will check if they need to dislocate from their housing unit. This decision will be dictated by the building/ housing failure status of their housing unit as established in the model initialization and their decision to repair the damage. If the housing unit is experiencing structural failure and the household does not begin performing repairs, then the household will leave this housing unit, change their position to the "none" space where households in non-permanent housing arrangements reside in the model. The dislocated households will remain there unless they find another suitable household or if they choose to outmigrate. It is worth noting here that dislocation in this model is not exactly analogous to "dislocation" as it is used elsewhere in the literature. The implementation here does not capture those individuals who dislocate temporarily while performing repairs or with the express intent to return to the same housing unit. Rather, the dislocation used here could be more appropriately described as a sort of pre-outmigration space in which households will need to quickly find another suitable structure at their price point in order to return to the community. They are not residents who will repair their own housing unit and thus are left more at the whim of the actions of the other agents in the model.

Table 28: Household dislocation status based on initial damage state

Initial Damage State	Midwest Field Study Data: Unrepaired buildings two years after the event	Modeled: Household dislocation status
No Damage	0.68%	0% True 100% False
DS 1 (Slight)	13.45%	0% True 100% False
DS 2 (Moderate)	21.64%	21.64% True 78.36% False

DS 3 (Extensive)	45.97%	45.97% True 54.03% False
DS 4 (Complete)	92.31%	92.31% True 7.69% False

Function 2: find_new_house

While executing the "find_new_house" function, household agents that are dislocated or looking to move will search for viable housing units by running three screenings of all housing units. First the housing units will be filtered down to only those that are unoccupied, have not been cleared, are finished being constructed, and have an owner type of household. Once the full set of preliminarily viable structures is found, the household picks one at random for which the original resident had an income within plus or minus 10% of the current resident, as shown in (10). This creates a situation where residents select a home similar to their original home so that their new housing unit will meet their needs without greatly exceeding their budget. Thus, families of four who previously lived in a three-bedroom house are not randomly matching with neither a one-bedroom apartment nor a seven-bedroom house. In situations where the housing unit was empty after initialization, any household agent that is dislocated or looking for a new residence is allowed to move into this structure regardless of income starting at the next timestep. Although this is slightly unrealistic, it also only affects a small subsection of unclaimed housing units. As the final screening before the agent finalizes the move they are making, the housing unit's repair history and damage must be assessed. This is done using a secondary function nested in the primary function called "CheckSingleStructureConditions." In this secondary function, a household simply checks if the housing unit has already failed with a different resident and then does not allow a new agent to move in. These failed structures are prevented from receiving new residents because the data from the Midwest Field Study was used to set the original housing unit failure rate, so allowing later residents to perform repairs on failed housing units would lead to these empirical data values not being followed. If the household agent has chosen a failed house, then they will return to the "none" space, continue to increase their time dislocated, and look for a housing unit in the next time step or outmigrate.

$$New\ Resident\ Income * 0.90 \leq Original\ Household\ Income \leq New\ Resident\ Income * 1.10 \tag{10}$$

Function 3: perform_repairsHH

In the perform_repairsHH function, household agents will collect information about the damage state of their housing unit and if they have decided to repair during the first two years in the initialization of the model, they will begin to conduct repairs on the building at a pace outlined by the repair rates from the Midwest Field Study and tabulated in Table 29. For instance, a building in an initial Damage State of DS1 whose owner has decided to perform repairs will have a 0.4627 probability of being fully repaired by three months after the event. Thus, a random number is generated from zero to one, and if the random number is less than 0.4627, then the housing will repair and update its damage to be "No Damage." Necessarily due to the nature of the input data, these repairs happen on a step function and no repairs are conducted after 24 months.

Table 29: Repair rates from the first six waves of the Midwest Field Study assuming the household has chosen to perform repairs at initialization

Probability of being repaired at each wave of data collection given that the building is repaired by wave 6 (i.e. housing_failure = false)					
Initial Damage State	**Wave 2**	**Wave 3**	**Wave 4**	**Wave 5**	**Wave 6**
DS 1 (Slight)	0.46	0.72	0.94	0.95	1.00
DS 2 (Moderate)	0.53	0.77	0.92	0.97	1.00
DS 3 (Extensive)	0.26	0.65	0.81	0.95	1.00
DS 4 (Complete)	0.14	0.58	0.69	0.81	1.00
No Damage	0.72	1.16	1.28	1.56	1.00

Counting Functions: checktimedislocated and checktimeoutmigrated

Thes functions are rather self-explanatory and simply serve as ways to count how long a household has been dislocated or outmigrated in stand-alone functions to avoid double counting during a given timestep.

Business and Institution Steps

The business and institution agents in the outmigration model perform the following steps at each tick. First, if they do not have a sufficient number of workers, they hire members of the community using the "find_workers" function. Then, those agents that have had to close their previous locations will look for a suitable place to reopen in the "find_new_location" function. They will check if any of their workers have left the community in the "lose_workers_to_outmigration" function and adjust their worker count to reflect these losses. The agents will then determine during the "close" function if they need to close due to a lack of workers or a building failure. If they are not trying to relocate and have an operating location, they will determine if their repairs are complete at a given timestep in the "perform_repairs" function in alignment with the probabilities in Table 29. The agent will then conclude its step. These functions are explained in more detail below.

Function 1: find_workers

At the beginning of the business and institution agents' step, they will look for workers to fill at least the minimum workers needed quota in order to remain operational. This value is theoretical and probabilistic, assigned randomly to all agents of this type to provide sufficient employment for the community. To provide a proxy for economic growth, the businesses and institutions can offer employment to 100% to 120% of the workers needed quota multiplied by the number of years the model has been running, and this multiplier is once again set in accordance with a randomized value. This also allows businesses to have some operational redundancy and simulates the experience of understaffing without closure. If a worker is found, the household to which that worker belongs will have their residents "looking for work" value decreased by one and the business or institution will increase their worker count by one as well as save the id of the household of the new worker to track this relationship.

Function 2: find_new_location

Similar to the find_new_house function for households, the find_new_location function for businesses and institutions will allow those businesses that are closed to look for a new building in which to operate. Unlike households, they do not need to economically match the previous agent that occupied the building, because these agents are not the primary focus of this model and need only sufficiently operate at the periphery of the model. Again, similar to the find_new_house function, if the new building an agent finds is deemed to have failed in the initialization of the model, then the agent will fail to move to this building and have to look again during the next time step.

Function 3: lose_workers_to_outmigration

This simple function tracks if any worker's households have outmigrated. If they have then the business or institution will sever ties with the worker's household by decreasing their worker count by one and removing the household id from the list of worker ids.

Function 4: close

In this function, the business and institution agents will determine if their building has experienced structural failure or if they have gone more than six months since the hazard event without sufficient staff. If either of these situations has occurred, the agent will move to the "none" position to indicate that they have relinquished their property. Then they will set “open” to false, and they will release their workers by increasing the count of workers looking for work in each household by one for those households who have workers that were formally employed at this establishment. Then they will clear their worker ids and begin looking for another location in the next timestep.

Function 5: perform_repairsBS and perform_repairsIN

In its current formulation, this process of performing repairs is identical to how the households perform repairs with identical parameterization because the building repair rates were taken as an aggregate rather than parsing them out by owner type. Doing so would have made some of the higher damage state bins insufficiently small and lead to greater uncertainty. For implementation of this model using IN-CORE recovery models, this differentiation of repair rate by owner type may be more feasible. Nevertheless, for the present model, refer to the repair performance function defined in the household steps for more details.

Model Steps

Above the actions of individual agents exists the actions of the model. The first model step simply calls on the agents to execute their steps. Then, the model calls on the household agents to execute the paramount steps of outmigrating and inmigrating in accordance with the OMM and IMM percentages calculated in (4) and (5). Then model-level parameters are re-calculated, and the model concludes this step. The functions underlying these processes are detailed below.

Function 1: AllAgentsStep

This function triggers the execution of all agent functions outlined above for every household, business, and institution in the community. When all these actions have been performed for the complete dataset, then the model can step through the necessary model-level steps described below before returning to the agents for the next step.

Function 2: Outmigrate

In this function, each household counts the number of outmigrated households and the total number of non-migrants to determine the ever-changing percentage of outmigrated households. If this outmigration percentage is less than the projected outmigration percentage from the OMM linear model equation and thus more households are needed to outmigrate to match this value, the household will check if they are a good candidate because they have been (a) been dislocated for more than 6 months, (b) their OMI score is amongst the highest OMI scores in the community and they are dislocated, (c) all workers in their household have been without work for 6 months, or (d) they simply have a notably high OMI value compared to the rest of the community. These four conditions are checked in order and thus OMI alone will only dictate outmigration if the three previous conditions have resulted in enough dislocated households. The OMI threshold for outmigration is determined by percentile rank of the OMI score in the community, namely if the percentile of their OMI score is greater than the outmigration threshold (100 – OMM). In the case of Mayfield, the OMM is initialized as 12.16% in accordance with (4), making the outmigration threshold 87.84% at the start of the model. In order to prevent far too many candidate outmigrants from leaving the community out of pace with the OMM, the OMM is recalculated, and the total outmigration counts are checked before each conditional for outmigration mentioned above is checked with the exception of the first conditional. This conditional checking for the dislocation time of a given household will be allowed to result in outmigration of that household even if the OMM quota of outmigrated households has already been reached for that time step. This exception is made to simulate the reality that dislocation becomes analogous to outmigration when enough time has passed in this state regardless of how many other households have left the community.

If the household meets any of the conditionals for outmigration, they will add themselves to the cohort of outmigrated agents by moving themselves to the "none" position, increasing the outmigration by one, decreasing the number of households in their building by one, setting the households outmigrated status to true, and turning off their interactivity with the job market. These agents now have no possibility of returning to the community. Although this does not exactly reflect reality and some residents may at some point return, it is beyond the definition of outmigration here to characterize this behavior and model it here. Though it can be said that in the modelling environment, a real-world scenario of a household returning would appear as a household inmigrating for the first time with similar characteristics to a household that previously outmigrated, which is already occurring organically in the model due to the formulation of the input parameters. This would be functionally equivalent in the model and not impact results. These four conditions are checked in order and thus OMI alone will only dictate outmigration if the three previous conditions have not resulted in enough outmigrated households.

Function 3: Inmigrate

The inmigration step adds agents to the model to meet the inmigration quotas set by the IMM multi-regression equation. First the total number of inmigrated households is tabulated and the inmigration percentage is calculated. Then the adjusted IMM value is determined to be the average of all non-outmigrated households' IMI values. The deficit between the total percentage of inmigrated households in the model and the updated IMM value multiplied the number of years the model has been running will determine the number of new households to be added at this time step.

The required new households will be added and given social characteristics in the same way the original agents were; however, these agents will not start in a house that dictates their income.

Rather, they will be assigned an income randomly in proportions relative to the national income distribution. Once added to the model, the agents look for a place to live using the find_new_house function. Then finally the model increases its count of inmigrated households by one.

Function 4: EstablishOutmigration

As mentioned in initialization, this function is re-run at the end of every timestep to update the community-wide OMM value as well as each household's OMI percentile rank.

Function 3: UpdateModelCounts

This function is a simple tallying function described in the model initialization section.

5.2.3. Model Variations

5.2.3.1. The Jobless Model

The role of business and institutional ties in the form of job allocations are helpful in getting a clearer sense of the way workers are moving through, in, and out of a community. This in turn provides one dimension of each household's experience in the community and whether they choose to stay, move in, or move out. Although these ties add a level of realism to the model, it introduces another set of assumptions regarding number of employees at each business, types of businesses, how and when businesses will decide to close or get rid of some workers, and the bounds of how many workers they can hire versus the minimum number of workers with which they can still function. The result is larger bonds of uncertainty but also more potential to identify emergent characteristics. Thus, for the sake of completeness, a version of the model that does not account for job allocation is also provided, and its results are analyzed alongside the other model variations. This simplified model is termed the Base Model in the later results section and the available code files. This model represents the least intricate but also most validated version of the model.

5.2.3.2. The Neighborhood Influence Model

One other model variation of interest was explored once other models had been completed. This variation allows neighboring households to influence each other regarding propensity for outmigration. This both seems to address an intuitive process by which neighbors influence each other's actions while also helping to address the question that is born out of the shift from aggregate metrics to individualized metrics. Namely, this concern would be that it is not the agents with predictive traits for outmigration from the aggregate community-level model that are outmigrating but rather these agents are by some roundabout means influencing others to outmigrate. For instance, it is shown that households are moving out at a higher rate when female unemployment is high across the community, but that does not make it logically conditional that households with unemployed females are the ones moving out. Rather they could be influencing the incredibly complex system of a community to generate more outmigrants by some non-apparent mechanism. It is apparent that these are factors of predictive power, but it is not known how or why they have predictive power. Consequently, the incorporation of neighborhood influence helps to address this lingering question. If it is assumed that it is individuals with the predictive traits who are moving out, then neighborhood influence could be set to zero. Whereas if it is assumed that in fact these factors are changing the likelihood of other nearby agents to outmigrate, then the neighborhood influence could be set all the way to one, indicating that each agent needs to rely only on the aggregate score of its neighbors to decide whether to outmigrate. This more scenario-based and user-modified model variation is only built onto the Base Model, not the job allocation variation, to avoid layering more uncertainty onto the business-level uncertainties already described in the previous paragraph.

5.3. Results and Discussion

The following results are broken into three separate analyses which are divided in alignment with the code modifications mentioned above. First, Model A is the model which incorporates job loss into the decision to outmigrate. Model B allows agents to consider their outmigration probability as well as that of their neighbors in the decision to move away. And finally, Model C provides a simplified version of the code without jobs or neighborhood ties that involves the fewest assumptions but also does not reflect the impact of these other notable community ties to place.

For validation of these results, the ACS 5-year data has been compiled for Mayfield city, Kentucky. This dataset is chosen because there really is not another alternative dataset for this set of parameters at this resolution; however, it is unclear how much of a hazard's impact is reflected in the values summarized by ACS 5-year data. This dataset is not ideal for comparison because the 5-year dataset will by definition consider years from before the event, but the 1-year dataset is not available for a community of this size. The provided values can also not be taken entirely at face value as their margin of error is quite significant and sometimes larger than the data itself, as is the case with limited English household values. This is all to say that alignment with this dataset is not seen as a complete validation or invalidation of this model. In fact, it is more so the trends in the ACS data rather than the exact values that are considered here for alignment. Table 30 shows the summary of this data for Mayfield. Because the event month was December 2021 and thus almost perfectly between the two years, the estimated annual values were averaged to get a December value for consideration. In the following graphs, the values from Table 30 have been included as a dashed black line with the margin of error from the census data represented as a gray bounded range. In these analyses, December 2021 is mapped to the 0 timestep as this is the month of the event that impacted the community of Mayfield. Each graph below is generated for the parameters noted in the equations for OMM and IMM stated earlier with the addition of mean household income to better characterize the economic distribution of the impact and the exception of the percentage of mobile homes as this physical feature cannot be directly mapped here with consideration of the building input dataset. For each factor, the percentage change is shown is it was projected by each of the three models (A, B, and C). For Model B, the neighbor influence value is set to a trial value of 0.5. A set of 25 runs were executed for each model, the results of which are displayed as lines of various muted colors in the Figures below. Multiple iterations were run to convey some of the uncertainty present in these factors as the model unfolds for each simulation. For a full and robust stochastic analysis of these parameters, a larger set of runs would be necessary, but with the two goals in this analysis being the comparison of the three models and the validation of these models against already imprecise census data, a stochastic analysis was unnecessarily computationally expensive. Though this may certainly be of interest to incorporate in future iterations of this model.

Table 30: ACS 5-year data for Mayfield city, Kentucky

Months since event	0	6	12	18	Trend
date	Dec. 2021	June 2022	Dec. 2022	June 2023	
median household income	39901	42589	42024	41458	Up
% with bachelor's degree or greater	21.5	22	22	22	Up or Stationary
% of households with limited English proficiency	0.56	0	0.41	0.82	Up
% of Females 16+ in workforce who are unemployed	5.04	4.15	3.92	3.69	Down
% of households w/o health insurance	7.2	7.1	6.15	5.2	Down
% of single parent households	13.8	13.4	13.8	14.1	Up or Stationary
% owner occupied housing units	55.9	58.2	61.2	64.1	Up
% of households w/o a vehicle	11.4	12.3	11.7	11.2	Down or Stationary
count of total households	4195	4175	4214	4252	Up or Stationary

5.3.1. Income

5.3.1.1. MEDIAN HOUSEHOLD INCOME

The model data for median household income does not possess the inflection points found in the census data, though this is not an immediate disqualifier as the census data is considering the past five years in this data, and thus, an inflection point may reflect a value shift earlier in the record. As for matching the general trend, Model C does well to generate a spray of results to either side of the final census data point at 18 months. Model A sufficiently matches the general trend as does Model B.

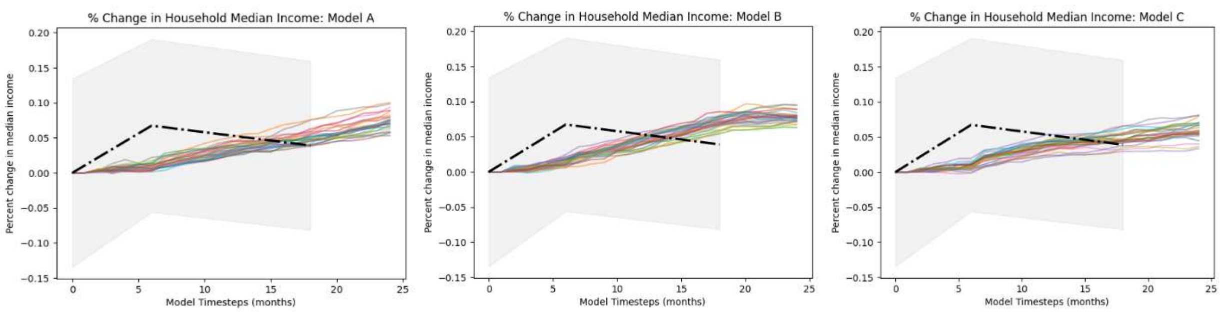

Figure 43: Change median household income projections for each model

In contrast with the inflection point inconsistencies seen in median household income, mean household income sees similar line behavior from the census data and the models' results. However, the models drag slightly below the census expected value. For Model C, this low skew intensifies slightly at about 14 months. Model B similarly has strong agreement with the census value until a flattening out, but for Model B, this flattening comes after the census data record ends. Meanwhile Model A maintains a trajectory seemingly identical to that of the census data. The importance of this difference in behavior is likely minor as the margin of error is still such that all three models reside comfortably within its bounds. Nevertheless, the trends shown in these figures demonstrate a heartening level of agreement between the census data and the models generated in this work.

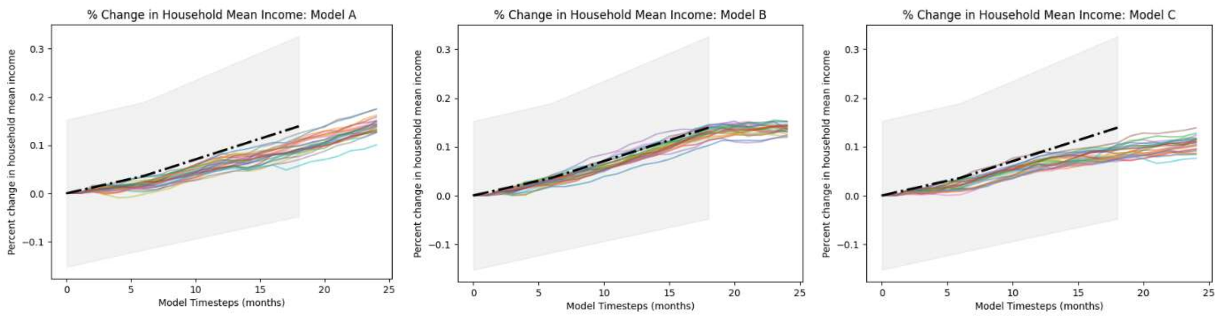

Figure 44: Change mean household income projections for each model

5.3.2. Educational Attainment

For educational attainment, the models show a plausible track of growth in educational attainment that adheres closely to the growth trajectory shown in the census data for the first six timesteps and then drifts slightly beyond the expected value from the census in the later timesteps, but is still well within the bounds of possibility and does not represent a meaningful departure from the census data. If one model were to be selected as having more predictive strength than the others, it may be Model C. Though the differences are marginal enough that if this factor alone were used to rank the performance of these models, then a stochastic analysis would be needed. However, other factors are included here and may provide further insights.

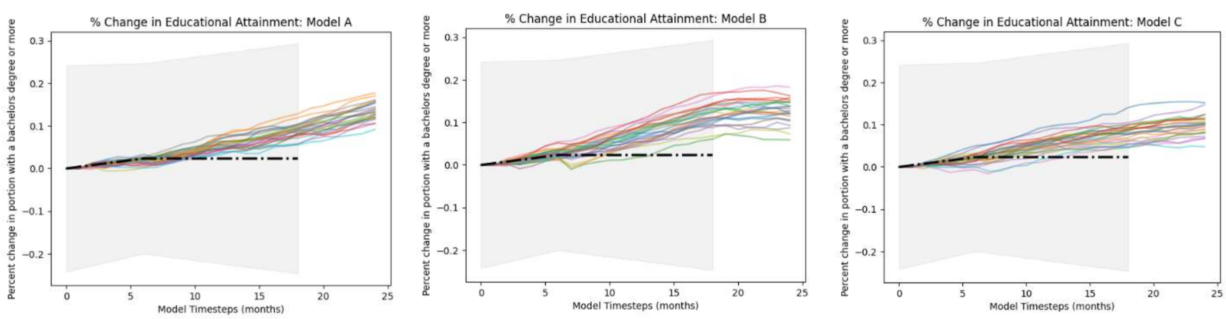

Figure 45: Change in educational attainment projections for each model

5.3.3. Unemployed Females

The margin of error from the census data is so large for this parameter as to make a deep investigation of it's trends somewhat fruitless. To quickly summarize, all three models predicted relative stationarity of this community characteristic with a quite marginal downward skew. The census data's expected value projected a slightly more pronounced negative value. This parameter's results do not eliminate any model from consideration.

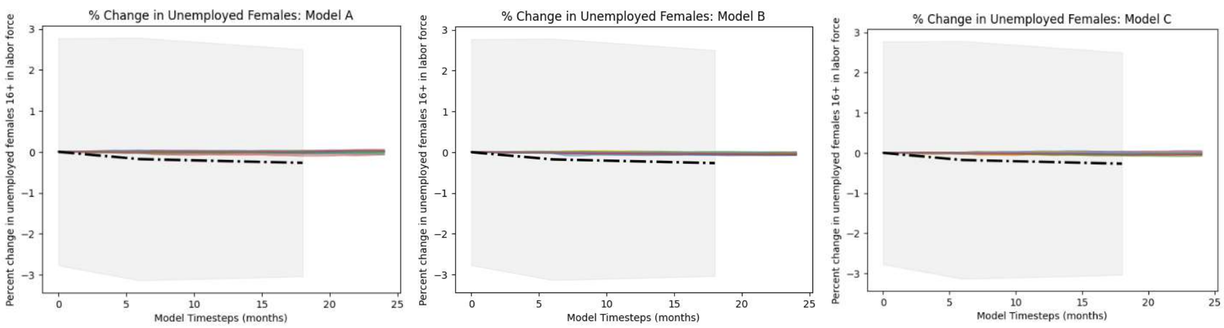

Figure 46: Change in unemployed females projections for each model

5.3.4. Limited English-Speaking Households

Due to the size of this community, the margin of error, and the incredibly small size of their limited English-speaking households, this factor's volatility in the census data is not surprising. By contrast, the model outputs appear almost stationary, but with the census value crisscrossing the model values, this parameter's results are seen as satisfactory for all models albeit slightly less impactful of a result than some of the others.

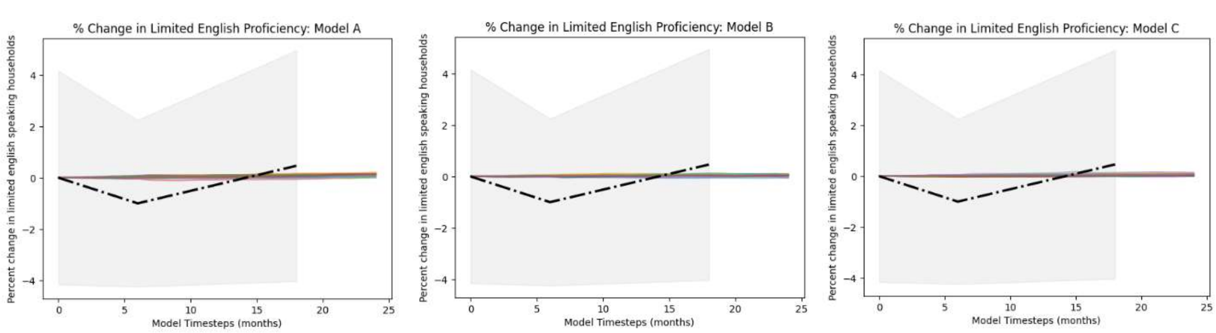

Figure 47: Change in limited English proficiency projections for each model

5.3.5. Households without Health Insurance

This factor was shown to have a slightly negative trend in the census data, and this trend is mimicked quite satisfactorily by the three model outputs below. Indeed, in all three models, the outputs have created an evenly distributed cloud around the census expected value for the entire monitored period.

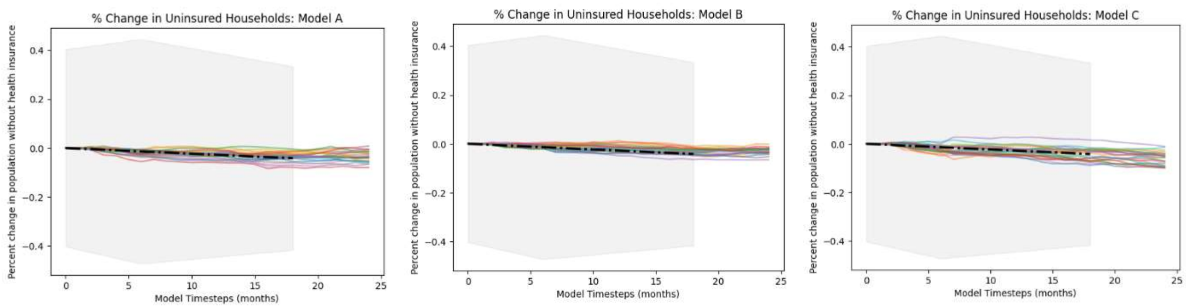

Figure 48: Change in uninsured household projections for each model

5.3.6. Single Parent Households

The change in single parent households showed fluctuation rather than the consistent trend seen with some of the other parameters. Overall, this value rose slightly but only after seeing a dip at the 6-month mark. The models seem to characterize the full range of values experienced quite well, though they all end on a slightly negative trend rather than a slightly positive one. These slight variations in trend are likely negligible, but this will only become clear when the 2024 ACS data is released.

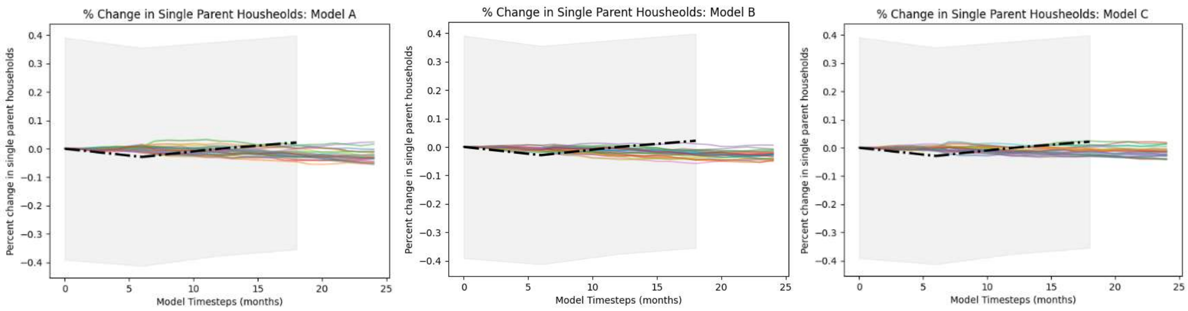

Figure 49: Change in single parent household projections for each model

5.3.7. Owner Occupied Housing Units

This parameter introduces some of the most drastic changes seen in the dataset following the event and similarly yields some of the most evidence in determining viable or non-viable models as these models have largely behaved well for other factors presented here. Model B showed almost no positive trend and is quickly eliminated from consideration by lying outside the bounds of the margin of error for this parameter. Model A and Model C adhered quite well to the census value for the first six months, but at that time, they began to diverge from this expected value. Model A lost almost all its upward trend and began to take on a more constant value. Although Model C stagnated for approximately six months to a year, by the end of the simulation the upward trend had returned to Model C's results roughly on pace with the expected census value. Furthermore, Model C stayed within the margin of error bounds of the census data. Because the results for most other parameters exhibit imperceptibly small variations, this parameter seems the most well suited to decide the ideal model for this case study. Namely, that is Model C.

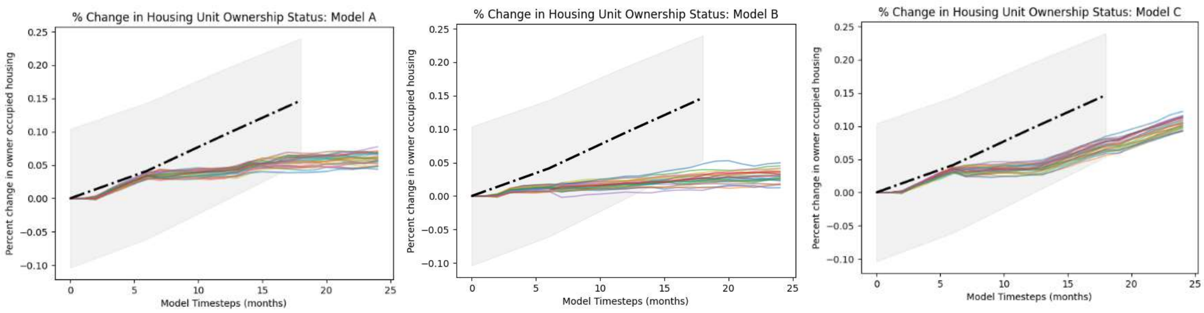

Figure 50: Change in ownership status projections for each model

5.3.8. Households without a Vehicle

All models for this factor adhered to the general trend of the data and accurately predicted the 18-month mark expected value. They, however, did not fully echo the inflection pattern shown in the census data, but as mentioned previously, this inflection point may or may not have actually occurred at that moment but rather may have been carried over from values in previous years given the formulation of this dataset. Thus, this feature is unnecessary to capture, and this graph shows another alignment between prediction and estimated reality.

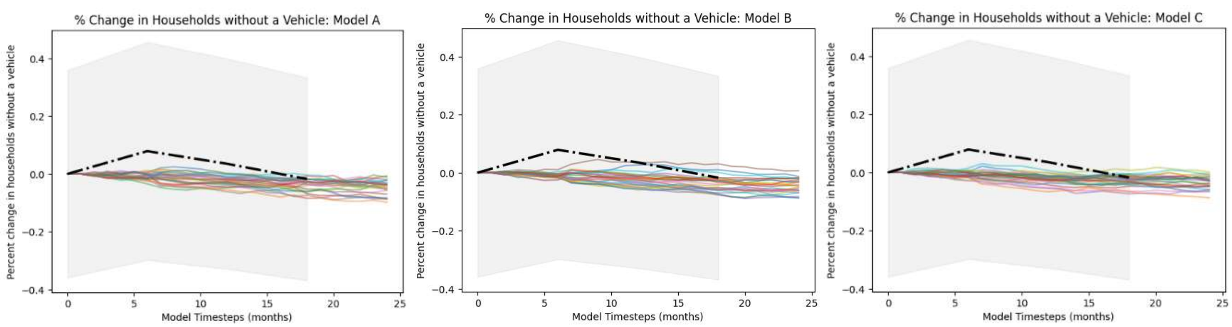

Figure 51: Change in households without a vehicle projections for each model

5.3.9. Total Count of Households

The change in total household count is dictated in the model by the IMM and OMM calculations as well as the rules for outmigration. As such the model will show greater volatility as this represents a volatile time for the impacted community. This sort of volatility will not be present in the census data by virtue of how that data is collected, processed, and represented to the end user. Nevertheless, the alignment between the models and years-long trends should be somewhat aligned. Due to the unavailability of the 2024 data at the time of this writing, the more long-term agreement between the census expected values and the model outputs cannot be entirely confirmed. However, it is clear in the figures below that for Model A and Model C the change in household count remained within the census data bounds and the population was returning to generally returning to pre-event levels in alignment with the census data expected value, which showed slight decline and very low-level growth during this time period. Meanwhile, Model B saw an early outmigration followed by a continuous increase in population, exceeding the census expected value but not breaking through the margin of error bounds.

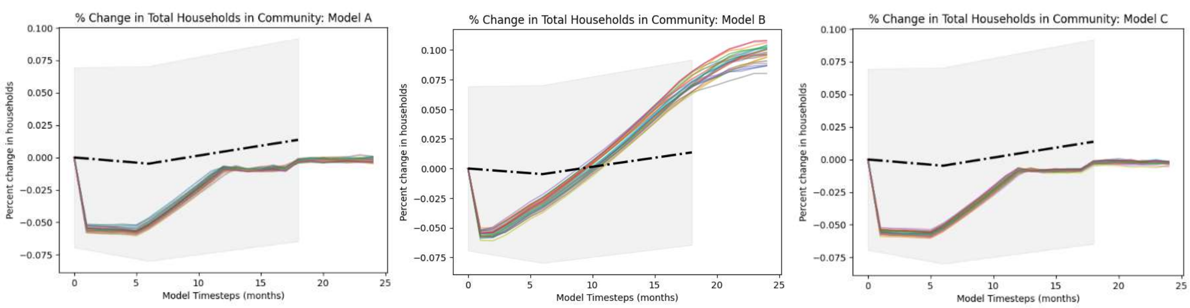

Figure 52: Change in total household projections for each model

Taken as a sum total of all these results, Model A and Model C align notably well with census data for this community and this analysis period with Model A faring slightly better overall. By contrast, Model B has a few noteworthy discrepancies that make it a difficult model to select. Though it is worth noting that in other iterations where the neighbor influence value for Model B was set to a lower value of say 0.2, the model behaved quite closely to Model C, indicating that there is not much to be gained from the model aside from what Model C already has to offer.

As stated previously, the census data has several limitations especially at this scale and was used because it is the only consistent data source available to analyze this type of social data for this area. However, this multi-factor alignment and predictive power across a variety of parameters for a hindcasted event indicates that this outmigration model merits further exploration and implementation.

5.4. Areas for Future Work

The model presented here shows good results for the percent change of the factors outlined above as this was the primary goal in understanding not only population change but also how a community's characteristics change after an event and the ways this might impact its resilience. However, it is worth noting that a lingering question is a misalignment of some of the factors initialization values from census bounds of confidence. For instance, the current model does not generate enough households, consistently generating about a hundred households less than the lower bound of the census data. This is likely due to the handling of household generation which is done in alignment with the number of structures designed as housing units from the NSI dataset. Two sources of error exist here that likely contributed to the model initialization values being artificially low. First are the errors inherent to the NSI dataset and second is the fact that structures with multiple housing units that did not classify as apartments were difficult to capture and thus not considered. However, it is known from data collection that at least one neighborhood of this particular community was made up of duplexes, and they are a frequent enough occurrence in American cities to be necessary to consider in future iterations of the model. This first point drives home once again the utmost importance of access to high-quality data in order to provide meaningful results to communities about the risks they face.

Secondly, the time window for this analysis was necessarily limited by the amount of time that has passed since the event. Doing more long-term projections of migration patterns using this tool will be helpful as the OMM and IMM linear calculations were trained on four years of data post-event, and it would be impactful to know if these results hold at longer timescales. Similarly, it is necessary to test the viability of this tool in other communities to ensure that the performance of the model is generalizable, or if it is not, determining where and when this generalizability breaks down so that the applicability of its implementation is not overstated.

Lastly, there are several means by which to expand this model to incorporate additional factors or mechanisms of interest that may impact the outmigration a community experiences. As such, exploring these options further would certainly be beneficial. For a starting point, although the implementation of business ties used here did not yield the most meaningful results, that is not to say that business ties couldn't be formulated in a different manner in the model that would prove quite successful. The business model shown in Model A here simply shows one possible implementation of this model feature and does not represent an exhaustive exploration of business' impact on outmigration.

5.5. Summary and Conclusions

As discussed in the framing of this investigation, multi-regression models and other top-down approaches are reasonably well-suited to characterize and predict aggregate values across a community. However, when trying to better understand the characteristics of the individuals that make up these aggregate values, top-down approaches are less helpful. Certainly, having a combination of top-down and bottom-up approaches utilized in this implementation allowed for the benefits to each modeling approach to be realized. From the top-down approach came the anchoring in reality that is sometimes missing in ABMs. Meanwhile, bottom-up ABM allowed for the relationship between different factors, their trends over time, and their impact on community-wide outcomes to be analyzed. This dovetailing of methods to inform more holistic results is to be a path forward for complex problems such as outmigration and other topics associated with natural hazards research. However, the balance of complexity and simplicity should always be considered in modeling not only due to the computational expense of complex models, but also because in some situations as seen here, the simplest model (Model C) can at times yield the best results especially when more complex models incorporate more assumptions.

6. SUMMARY AND CONCLUSIONS

6.1. Summary and Conclusions

This dissertation puts forth a series of methods and best practices by which interdisciplinary research can better integrate data from diverse fields, allowing the learnings in one discipline to inform the findings in another. It is the goal of this dissertation to leverage data integration techniques to generate more complex, yet still replicable, characterizations of communities, the risks they face, and the assets they possess to combat these risks. To encapsulate this intended dynamism, the proposed concept of social susceptibility is brought to the fore once again before discussing the more concrete contributions outlined for each chapter. This concept, which was introduced explicitly in Chapter 2 but truly underpins much of this dissertation, predicates a susceptibility *to* something, not simply a state variable like vulnerability. It is possible to arrive at more productive conversations that inspire action and policy-driven decision-making by pivoting the conversation away from the value judgement inherent in calling communities "vulnerable" or the implication that a community somehow needs to change to become not vulnerable. If for instance, as is the case in this work, a community is found based on empirical evidence to be socially susceptible to negative long-term outcomes after a hazard event, the community can place greater emphasis on developing a robust resilience plan and adhering to this plan. By the same logic, if a later metric is developed to characterize a community's social susceptibility to increases in economic inequality following an event, then new local policy may need to be drafted to help curb this post-event shift and promote equity. Thus, the specificity of "social susceptibility to…" could be mistaken to represent a minor shift in verbiage but in fact constitutes a key shift in the author's thinking that informed the remainder of this dissertation and is hoped to be a shift at least entertained by the reader as well. It is within this context that the latter contributions are noted.

Alongside this definition of social susceptibility, Chapter 2 details a social susceptibility metric with predictive strength for long-term outcomes. This social susceptibility metric emphasizes external validation techniques to ensure that the resultant value has clearly delineated and actionable real-world implications. This social susceptibility metric shows greater predictive strength than commonly used metrics in the field when considering adverse long-term community outcomes. The results of this analysis highlight aging populations, households with limited English proficiency, single parent households, households without a vehicle, and income inequality as impactful factors for long-term recovery outcomes. In conjunction with the social susceptibility metric, a decision support tool was designed and implemented using the Midwest field study as an example to ensure that the work being done was replicable and controlled. This tool is highly practical in nature and provides a template for incorporating social factors into damage assessment work in a standardized way. This data standardization was a request often encountered in the literature.

Another common refrain in the literature was for more high-quality, high-resolution data in a field in which data often comes at great cost to people and communities. In order to contribute to the alleviation of this need, a longitudinal dataset is under development that is being made available on DesignSafe-CI for researchers to validate damage and recovery models as well as explore new research question (J. van de Lindt et al., 2025). The collection, curation, and cleaning of this dataset is reviewed in Chapter 3. Due to the geographically large study area and the isolated nature of the different communities, this data also provides insight into the different trajectories that communities can have after disasters so that the drivers of these trajectories can begin to be understood. Thus, it provides another means by which the SSM provided in Chapter 2 can be assessed.

In Chapter 4, the task of data integration pivots from collection and curation to model implementation. As such, Chapter 4 presents a generalizable sheltering model to determine ideal community tornado shelter locations. Developed out of knowledge gained and observations made during the Midwest Field Study, the use case of this sheltering model presented in this dissertation puts forth a concept for public-private partnership to allow more access to community tornado shelters in places that need this alternative. Outside of this use case, this model can also be used for any set of shelters and sheltering agents, thus providing the necessary visualization and easily understandable scenario analysis to allow for more effective siting of future community shelters.

Chapter 5 presents an agent-based model aimed at predicting outmigration and the demographic patterns of population migration following a hazard event. This is done by collating the techniques and outputs of previous endeavors. This model extends the agent-based modeling concepts introduced in Chapter 4 to this more exploratory topic. The multi-regression approaches utilized in Chapter 2 are revived to underpin this model with empirical data. Then the outmigration model is run for the community of Mayfield, Kentucky, and the data collected in Chapter 3 is used to develop the community dataset and building stock recovery dataset on which the model has been run. This population outmigration model demonstrates the strength of agent-based modeling to explore parts of hazards research that are difficult to define and even harder to quantify by allowing for greater investigation into the mechanisms that drive community behavior post-event.

6.2. Contributions

The contributions provided throughout this work have been condensed and provided here for efficiency. The primary contribution offered in Chapter 2 is a metric and protocol by which to select communities and triage data collection efforts after a hazard event in order to maximize what can be learned from a resilience and recovery perspective in longitudinal field study investigations. In Chapter 3, the key contributions offered are a structured dataset that tracks the building stock recovery of 6 distinct communities to validate repair and reconstruction rates for tornado recovery models as well as an extensive log of publicly available video data from the Midwest field study for diverse disciplines to use and learn from without the need for extensive funding and personnel. The work in Chapter 4 offers two over-arching contributions. The first is a generalizable sheltering model that would allow communities to determine a socially relevant and well-situated community tornado shelter location. Then the second is an application of this model validating a theory for shelter placement that would allow thousands more people access to a community tornado shelter in Alabama alone with similar results expected in many Southern states. And finally, the primary contribution of the research presented in Chapter 5 is a validated and predictive model for long-term outmigration and demographic changes after a hazard event that leverages many of the key learnings and data outputs from the previous chapters in order to serve as a culminating product of this dissertation.

6.3. Recommendation for Future Work

Each of the previous chapters have detailed areas for future work that are pertinent to the respective topics; however, a prioritized tabulation of the most vital areas for future work is also provided here for completeness. First and foremost in the discussion of furthering interdisciplinary data integration, the robustness of the data available ought to continue to grow and be supported by diverse teams of researchers. The datasets like the one put forth in Chapter 2 of this dissertation must grow not only in prevalence but also accessibility. In the field of hazards research, the data collected is inextricably intertwined with the suffering of real people and communities. Thus, in this field even more so than others, it is prudent to allow for as much knowledge as possible to be derived from this data in support of communities.

Second, unbiased and empirically validated metric development should continue to be of utmost importance. Metrics have long been a topic of great debate in the discussion of social vulnerability and for good reason. A person is not as simple as a single value let alone whole communities of people. It is reductionist and yet necessary at this point in our conceptualization of communities to have a means by which to make policy decisions, triage efforts, understand risks, and prioritize equity. Until a more holistic yet generalizable approach can be developed, metrics seem to still be the most promising path forward even if the story they tell is only part of the story. The more that can be done to further develop this partial story, the more beneficial research will be to real people and communities. It is worth noting that although theory is also important in development of these metrics, theory is inherently not devoid of personal perspective and thus should be addressed and consulted with an understanding of the implicit bias that has lived in the positions of power that have developed social theory for generations.

Lastly, it is important to continue the exploration of models that strive to sufficiently characterize the mechanisms by which recovery happens in a community. As datasets grow and modeling capability grows alongside it, so to does the potential for models to adequately capture the most "wicked" problems in hazards research. The next class of social models should further develop human behavior and in so doing allow for the formation of emergent qualities in the end results that cannot be captured by a single multi-regression. As the developed models are able to more accurately and efficiently represent the reality in which a hazard event occurs, the less necessary it will be for communities to learn about the risks they face by having to live through them.

REFERENCES

Cited Directly in Text

Bakkensen, L. A., Fox-Lent, C., Read, L. K., & Linkov, I. (2017). Validating Resilience and Vulnerability Indices in the Context of Natural Disasters. *Risk Analysis*, *37*(5), 982–1004. https://doi.org/10.1111/risa.12677

Beccari, B. (2016). A Comparative Analysis of Disaster Risk, Vulnerability and Resilience Composite Indicators. *PLoS Currents*. https://doi.org/10.1371/currents.dis.453df025e34b682e9737f95070f9b970

Bracken, L. J., & Oughton, E. A. (2006). 'What do you mean?' The importance of language in developing interdisciplinary research. *Transactions of the Institute of British Geographers*, *31*(3), 371–382. https://doi.org/10.1111/j.1475-5661.2006.00218.x

Brink, M., Hengeveld, G. M., & Tobi, H. (2020). Interdisciplinary measurement: A systematic review of the case of sustainability. *Ecological Indicators*, *112*, 106145. https://doi.org/10.1016/j.ecolind.2020.106145

Bromham, L., Dinnage, R., & Hua, X. (2016). Interdisciplinary research has consistently lower funding success. *Nature*, *534*(7609), 684–687. https://doi.org/10.1038/nature18315

Brotzge, J., & Donner, W. (2013). The Tornado Warning Process: A Review of Current Research, Challenges, and Opportunities. *Bulletin of the American Meteorological Society*, *94*(11), 1715–1733. https://doi.org/10.1175/BAMS-D-12-00147.1

Bruzzese, J.-M., Usseglio, J., Goldberg, J., Begg, M. D., & Larson, E. L. (2020). Professional development outcomes associated with interdisciplinary research: An integrative review. *Nursing Outlook*, *68*(4), 449–458. https://doi.org/10.1016/j.outlook.2020.03.006

Chen, X., Meaker, J. W., & Zhan, F. B. (2006). Agent-Based Modeling and Analysis of Hurricane Evacuation Procedures for the Florida Keys. *Natural Hazards*, *38*(3), 321–338. https://doi.org/10.1007/s11069-005-0263-0

Collins, H. M., & Evans, R. (2007). *Rethinking expertise*. University of Chicago Press.

Community. (2024). Dollar General Newsroom. https://newscenter.dollargeneral.com/company-facts/community

Corfidi, S. F., Weiss, S. J., Kain, J. S., Corfidi, S. J., & Rabin, R. M. (2010). Revisiting the 3-4 April 1974 Super Outbreak of Tornadoes. *Weather and Forecasting*, *25*(2), 465–510.

Crawford, P. S., Graettinger, A., Powell, L., Awondo, S., Back, E., & Spector, S. (2017). Five Years after the April 27, 2011, Tuscaloosa Tornado: A Study in Community Resilience. *AEI 2017*, 758–769. https://doi.org/10.1061/9780784480502.063

Cutter, S. L., Boruff, B. J., & Shirley, W. L. (2003). Social Vulnerability to Environmental Hazards: Social Vulnerability to Environmental Hazards. *Social Science Quarterly*, *84*(2), 242–261. https://doi.org/10.1111/1540-6237.8402002

Datseris, G., Vahdati, A. R., & DuBois, T. C. (2024). Agents.jl: A performant and feature-full agent-based modeling software of minimal code complexity. *SIMULATION*, *100*(10), 1019–1031. https://doi.org/10.1177/00375497211068820

Davidson, R. A. (2015). Integrating disciplinary contributions to achieve community resilience to natural disasters. *Civil Engineering and Environmental Systems*, *32*(1–2), 55–67. https://doi.org/10.1080/10286608.2015.1011627

Deloitte Digital. (2020). *LightOSM* [Computer software]. https://github.com/DeloitteOptimalReality/LightOSM.jl

Dollar General. (2021). *202 Annual Report & 2021 Proxy Statement* [Annual Report]. https://ir-api.eqs.com/media/document/f2450c48-4e6d-481a-a2d6-f22ec55789f8/assets/Dollar_General_2020_Annual_Report.pdf?disposition=inline

Eigenbrode, S. D., O'rourke, M., Wulfhorst, J. D., Althoff, D. M., Goldberg, C. S., Merrill, K., Morse, W., Nielsen-Pincus, M., Stephens, J., Winowiecki, L., & Bosque-Pérez, N. A. (2007). Employing Philosophical Dialogue in Collaborative Science. *BioScience*, *57*(1), 55–64. https://doi.org/10.1641/B570109

Evans, C. M., Adams, R. M., & Peek, L. (2021a). *CONVERGE training modules: Collecting and sharing perishable data*. https://converge.colorado.edu/resources/training-modules

Faber, M. H., Giuliani, L., Revez, A., Jayasena, S., Sparf, J., & Mendez, J. M. (2014). Interdisciplinary Approach to Disaster Resilience Education and Research. *Procedia Economics and Finance*, *18*, 601–609. https://doi.org/10.1016/S2212-5671(14)00981-2

Fang, F. C., & Casadevall, A. (2011). Reductionistic and Holistic Science. *Infection and Immunity*, *79*(4), 1401–1404. https://doi.org/10.1128/IAI.01343-10

Fatemi, F., Ardalan, A., Aguirre, B., Mansouri, N., & Mohammadfam, I. (2017). Social vulnerability indicators in disasters: Findings from a systematic review. *International Journal of Disaster Risk Reduction*, *22*, 219–227. https://doi.org/10.1016/j.ijdrr.2016.09.006

FEMA. (2022). *Community Resilience Indicator Analysis: 2022 Update*. https://www.fema.gov/sites/default/files/documents/fema_2022-community-resilience-indicator-analysis.pdf

Fielding, N. G. (2012). Triangulation and Mixed Methods Designs: Data Integration With New Research Technologies. *Journal of Mixed Methods Research*, *6*(2), 124–136. https://doi.org/10.1177/1558689812437101

Flynn, W., & Islam, T. (2019). Spatial and temporal analysis of the 27 April 2011 tornado outbreak in Central Alabama. *GeoJournal*, *84*(2), 533–544. https://doi.org/10.1007/s10708-018-9874-3

Ganapti, N. E., & Mostafavi, A. (2018). Cultivating Metacognition in Each of Us: Thinking About 'Thinking' in Interdisciplinary Disaster Research. *Risk Analysis*. https://doi.org/10.1111/risa.13226

Ge, Y. "Gurt," Zobel, C. W., Murray-Tuite, P., Nateghi, R., & Wang, H. (2021). Building an Interdisciplinary Team for Disaster Response Research: A Data-Driven Approach. *Risk Analysis*, *41*(7), 1145–1151. https://doi.org/10.1111/risa.13280

Gilligan, J. M. (2021). Expertise Across Disciplines: Establishing Common Ground in Interdisciplinary Disaster Research Teams. *Risk Analysis*, *41*(7), 1171–1177. https://doi.org/10.1111/risa.13407

Gomes, L. C., Bianchi, F. J. J. A., Cardoso, I. M., Schulte, R. P. O., Arts, B. J. M., & Fernandes Filho, E. I. (2020). Land use and land cover scenarios: An interdisciplinary approach integrating local conditions and the global shared socioeconomic pathways. *Land Use Policy*, *97*, 104723. https://doi.org/10.1016/j.landusepol.2020.104723

Green, S., & Andersen, H. (2019). Systems science and the art of interdisciplinary integration. *Systems Research and Behavioral Science*, *36*(5), 727–743. https://doi.org/10.1002/sres.2633

Hinkel, J., & Klein, R. J. T. (2009). Integrating knowledge to assess coastal vulnerability to sea-level rise: The development of the DIVA tool. *Global Environmental Change*, *19*(3), 384–395. https://doi.org/10.1016/j.gloenvcha.2009.03.002

IRS Statistics of Income Division. (2020). *Statistics of Income Tax Stats—Migration Data* [Dataset]. https://www.irs.gov/statistics/soi-tax-stats-migration-data

Jakobsen, C. H., Hels, T., & McLaughlin, W. J. (2004). Barriers and facilitators to integration among scientists in transdisciplinary landscape analyses: A cross-country comparison. *Forest Policy and Economics*, *6*(1), 15–31. https://doi.org/10.1016/S1389-9341(02)00080-1

Johnston, B., & Van De Lindt, J. (2024). Weighing structural damage and social susceptibility: A decision-making tool to perform longitudinal studies of geographically large hazard events. *Risk Analysis*, risa.14284. https://doi.org/10.1111/risa.14284

Johnston, B., & van de Lindt, J. W. (2022). Barriers and Possibilities for Interdisciplinary Disaster Science Research: Critical Appraisal of the Literature. *Natural Hazards Review*, *23*(1), 04021053. https://doi.org/10.1061/(ASCE)NH.1527-6996.0000520

Kallemeyn, L. M., Hall, J. N., & Gates, E. (2020). Exploring the Relevance of Complexity Theory for Mixed Methods Research. *Journal of Mixed Methods Research*, *14*(3), 288–304. https://doi.org/10.1177/1558689819872423

Kendra, J., & Nigg, J. (2014). Engineering and the social sciences: Historical evolution of interdisciplinary approaches to hazard and disaster. *Engineering Studies*, *6*(3), 134–158. https://doi.org/10.1080/19378629.2014.978335

Krocak, M. J., Allan, J. N., Ripberger, J. T., Silva, C. L., & Jenkins-Smith, H. C. (2021). An analysis of tornado warning reception and response across time: Leveraging respondent's confidence and a nocturnal tornado climatology. *Weather and Forecasting*. https://doi.org/10.1175/WAF-D-20-0207.1

Kröger, M., & Schäfer, M. (2016). Scenario development as a tool for interdisciplinary integration processes in sustainable land use research. *Futures*, *84*, 64–81. https://doi.org/10.1016/j.futures.2016.07.005

Kuligowski, E. D., Lombardo, F. T., Phan, L. T., Levitan, M. L., & Jorgensen, D. P. (2013). *Draft final report, National Institute of Standards and Technology (NIST) Technical investigation of the May 22, 2011 tornado in Joplin, Missouri* (NIST NCSTAR 3; p. NIST NCSTAR 3). National Institute of Standards and Technology. https://doi.org/10.6028/NIST.NCSTAR.3

Kulp, S. A., & Strauss, B. H. (2019). New elevation data triple estimates of global vulnerability to sea-level rise and coastal flooding. *Nature Communications*, *10*(1), 4844. https://doi.org/10.1038/s41467-019-12808-z

Lafuerza, L. F., Dyson, L., Edmonds, B., & McKane, A. J. (2016). Staged Models for Interdisciplinary Research. *PLOS ONE*, *11*(6), e0157261. https://doi.org/10.1371/journal.pone.0157261

Lattuca, L. R. (2001). *Creating interdisciplinarity: Interdisciplinary research and teaching among college and university faculty* (1st ed). Vanderbilt University Press.

Le Sueur, H., Bruce, I. N., Geifman, N., & on behalf of the MASTERPLANS Consortium. (2020). The challenges in data integration – heterogeneity and complexity in clinical trials and patient registries of Systemic Lupus Erythematosus. *BMC Medical Research Methodology*, *20*(1), 164. https://doi.org/10.1186/s12874-020-01057-0

Lynam, T., Damayanti, R., Rialine Titaley, C., Suharno, N., Bradley, M., & Krentel, A. (2020). Reframing Integration for Mixed Methods Research. *Journal of Mixed Methods Research*, *14*(3), 336–357. https://doi.org/10.1177/1558689819879352

Miller, T. R., Baird, T. D., Littlefield, C. M., Kofinas, G., Chapin Iii, F. S., & Redman, C. L. (2008). Epistemological Pluralism: Reorganizing Interdisciplinary Research. *Ecology and Society*, *13*(2), art46. https://doi.org/10.5751/ES-02671-130246

Moezzi, M., & Peek, L. (2021). Stories for Interdisciplinary Disaster Research Collaboration. *Risk Analysis*, *41*(7), 1178–1186. https://doi.org/10.1111/risa.13424

Nastar, M., Boda, C. S., & Olsson, L. (2018). A critical realist inquiry in conducting interdisciplinary research: An analysis of LUCID examples. *Ecology and Society*, *23*(3), art41. https://doi.org/10.5751/ES-10218-230341

Nateghi, R., Sutton, J., & Murray-Tuite, P. (2021). The Frontiers of Uncertainty Estimation in Interdisciplinary Disaster Research and Practice. *Risk Analysis*, *41*(7), 1129–1135. https://doi.org/10.1111/risa.13337

National Academy of Sciences, National Academy of Engineering, & Institute of Medicine. (2004). The drivers of interdisciplinary research. In *Facilitating Interdisciplinary Research* (pp. 26–40). National Academies Press. https://doi.org/10.17226/11153

National Institute of Standards and Technology. (2023, June 1). *Community Resilience center of Excellence*. Centers of Excellence. https://www.nist.gov/coe/community-resilience-center-excellence

National Science Foundation. (n.d.). *Disaster Resilience Research Grants: Program Solicitation NSF 22-593*. Retrieved June 1, 2023, from https://new.nsf.gov/funding/opportunities/disaster-resilience-research-grants-drrg

National Science Foundation. (2016). *10 big ideas for future NSF investments*. https://www.nsf.gov/about/congress/reports/nsf_big_ideas.pdf

National Weather Service. (2021). *NWS Storm Damage Summaries: Dec 10-11, 2023 Tornado Outbreak*. https://www.weather.gov/crh/dec112021

National Weather Service. (2022). *December 10-11, 2021 Tornado Outbreak*. https://www.weather.gov/ohx/20211211

NCDC Storm Events Database. (1996). [Dataset]. National Centers for Environmental Information. https://doi.org/gov.noaa.ncdc:C00510

Newgreen, D. F., Landman, K. A., & Osborne, J. M. (2019). Addressing Interdisciplinary Difficulties in Developmental Biology/Mathematical Collaborations: A Neural Crest Example. In Q. Schwarz & S. Wiszniak (Eds.), *Neural Crest Cells* (Vol. 1976, pp. 21–36). Springer New York. https://doi.org/10.1007/978-1-4939-9412-0_2

NHERI. (n.d.). *CONVERGE Overview*. CONVERGE Natural Hazards Center. Retrieved July 1, 2023, from https://converge.colorado.edu/

Niche. (2021). https://www.niche.com/places-to-live/search/best-places-to-live/?geoip=true

Norris, F. H., Stevens, S. P., Pfefferbaum, B., Wyche, K. F., & Pfefferbaum, R. L. (2008). Community Resilience as a Metaphor, Theory, Set of Capacities, and Strategy for Disaster Readiness. *American Journal of Community Psychology*, *41*(1–2), 127–150. https://doi.org/10.1007/s10464-007-9156-6

Ntontis, E., Drury, J., Amlôt, R., Rubin, G. J., & Williams, R. (2020). Endurance or decline of emergent groups following a flood disaster: Implications for community resilience. *International Journal of Disaster Risk Reduction*, *45*, 101493. https://doi.org/10.1016/j.ijdrr.2020.101493

OpenStreetMap Contributors. (2024). *Planet dump retrieved from https://planet.osm.org* [Dataset]. https://www.openstreetmap.org

Ozili, P. K. (2022). The Acceptable R-Square in Empirical Modelling for Social Science Research. *SSRN Electronic Journal*. https://doi.org/10.2139/ssrn.4128165

Paducah National Weather Service. (2021). *The Violent Tornado Outbreak of December 10-11, 2021*. National Weather Service. https://www.weather.gov/pah/December-10th-11th-2021-Tornado

Paducah National Weather Service. (2022). *Summary of snow event on January 6, 2022*. National Weather Service. https://www.weather.gov/pah/Jan062022_Snow

Peacock, W. G., Dash, N., & Zhang, Y. (2007). Sheltering and Housing Recovery Following Disaster*. In H. Rodríguez, E. L. Quarantelli, & R. R. Dynes, *Handbook of Disaster Research* (pp. 258–274). Springer New York. https://doi.org/10.1007/978-0-387-32353-4_15

Peek, L., & Guikema, S. (2021). Interdisciplinary Theory, Methods, and Approaches for Hazards and Disaster Research: An Introduction to the Special Issue. *Risk Analysis*, *41*(7), 1047–1058. https://doi.org/10.1111/risa.13777

Peek, L., Tobin, J., Adams, R. M., Wu, H., & Mathews, M. C. (2020). A Framework for Convergence Research in the Hazards and Disaster Field: The Natural Hazards Engineering Research Infrastructure CONVERGE Facility. *Frontiers in Built Environment*, *6*, 110. https://doi.org/10.3389/fbuil.2020.00110

Peek, L., Tobin, J., Lindt, J. W., & Andrews, A. (2021). Getting Interdisciplinary Teams into the Field: Institutional Review Board Preapproval and Multi-Institution Authorization Agreements for Rapid Response Disaster Research. *Risk Analysis*, *41*(7), 1204–1212. https://doi.org/10.1111/risa.13740

Peterson, G. D., Cumming, G. S., & Carpenter, S. R. (2003). Scenario Planning: A Tool for Conservation in an Uncertain World. *Conservation Biology*, *17*(2), 358–366. https://doi.org/10.1046/j.1523-1739.2003.01491.x

Prevatt, D. O., van de Lindt, J. W., Back, E. W., Graettinger, A. J., Pei, S., Coulbourne, W., Gupta, R., James, D., & Agdas, D. (2012). Making the Case for Improved Structural Design: Tornado Outbreaks of 2011. *Leadership and Management in Engineering*, *12*(4), 254–270. https://doi.org/10.1061/(ASCE)LM.1943-5630.0000192

Railsback, S. F., & Grimm, V. (2019). *Agent-based and individual-based modeling: A practical introduction* (Second edition). Princeton University Press.

Reilly, A. C., Dillon, R. L., & Guikema, S. D. (2021). Agent-Based Models as an Integrating Boundary Object for Interdisciplinary Research. *Risk Analysis*, *41*(7), 1087–1092. https://doi.org/10.1111/risa.13134

Repko, A. F. (2012). *Interdisciplinary research: Process and theory* (2. ed). SAGE Publ.

Richardson, L., Long, E., Goodfellow, C., Milicev, J., & Gardani, M. (2024). Starting an undergraduate degree amid the COVID-19 pandemic: A mixed-method egocentric network study on student loneliness. *PLOS ONE*, *19*(2), e0297953. https://doi.org/10.1371/journal.pone.0297953

Riley, E. V. (2024, May 30). The people racing to build shelters outside tornado alley. *BBC*. https://www.bbc.com/future/article/20240530-the-people-racing-to-build-shelters-outside-tornado-alley

Ritchie, L. A., & Gill, D. A. (2007). SOCIAL CAPITAL THEORY AS AN INTEGRATING THEORETICAL FRAMEWORK IN TECHNOLOGICAL DISASTER RESEARCH. *Sociological Spectrum*, *27*(1), 103–129. https://doi.org/10.1080/02732170601001037

Rodela, R., & Alašević, D. (2017). Crossing disciplinary boundaries in environmental research: Interdisciplinary engagement across the Slovene research community. *Science of The Total Environment*, *574*, 1492–1501. https://doi.org/10.1016/j.scitotenv.2016.08.144

Sapat, A. (2021). Lost in Translation? Integrating Interdisciplinary Disaster Research with Policy Praxis. *Risk Analysis*, *41*(7), 1232–1239. https://doi.org/10.1111/risa.13198

Sawada, Y., & Hanazaki, R. (2020). Socio-hydrological data assimilation: Analyzing human–flood interactions by model–data integration. *Hydrology and Earth System Sciences*, *24*(10), 4777–4791. https://doi.org/10.5194/hess-24-4777-2020

Schmidtlein, M. C., Shafer, J. M., Berry, M., & Cutter, S. L. (2011). Modeled earthquake losses and social vulnerability in Charleston, South Carolina. *Applied Geography*, *31*(1), 269–281. https://doi.org/10.1016/j.apgeog.2010.06.001

Sherman-Morris, K., Houston, J. B., & Subedi, J. (2021). Theoretical Matters: On the Need for Hazard and Disaster Theory Developed Through Interdisciplinary Research and Collaboration. *Risk Analysis*, *41*(7), 1059–1065. https://doi.org/10.1111/risa.13223

Shvaiko, P., & Euzenat, J. (2013). Ontology Matching: State of the Art and Future Challenges. *IEEE Transactions on Knowledge and Data Engineering*, *25*(1), 158–176. https://doi.org/10.1109/TKDE.2011.253

Silberstein, M. (2002). Reduction, emergence, and explanation 2002. "Reduction, emergence, and explanation." In The Blackwell guide to the philosophy of science, edited by P. K. Machamer and M. Silberstein, 80–107. Oxford, UK: Blackwell. In P. K. Machamer & M. Silberstein (Eds.), *The Blackwell guide to the philosophy of science* (pp. 80–107). Blackwell.

Smith, A. B. (2020). *U.S. Billion-dollar Weather and Climate Disasters, 1980—Present (NCEI Accession 0209268)* [Dataset]. NOAA National Centers for Environmental Information. https://doi.org/10.25921/STKW-7W73

Smith, G., Martin, A., & Wenger, D. E. (2018). Disaster Recovery in an Era of Climate Change: The Unrealized Promise of Institutional Resilience. In H. Rodríguez, W. Donner, & J. E. Trainor (Eds.), *Handbook of Disaster Research* (pp. 595–619). Springer International Publishing. https://doi.org/10.1007/978-3-319-63254-4_28

Spielman, S. E., Tuccillo, J., Folch, D. C., Schweikert, A., Davies, R., Wood, N., & Tate, E. (2020). Evaluating social vulnerability indicators: Criteria and their application to the Social Vulnerability Index. *Natural Hazards*, *100*(1), 417–436. https://doi.org/10.1007/s11069-019-03820-z

Subedi, J., Houston, J. B., & Sherman-Morris, K. (2021). Interdisciplinary Research as an Iterative Process to Build Disaster Systems Knowledge. *Risk Analysis*, *41*(7), 1072–1077. https://doi.org/10.1111/risa.13244

Tierney, K. J. (2018). *Disasters: A sociological approach*. Polity Press.

Tobi, H. (2014). Measurement in interdisciplinary research: The contributions of widely-defined measurement and portfolio representations. *Measurement*, *48*, 228–231. https://doi.org/10.1016/j.measurement.2013.11.013

Tobi, H., & Kampen, J. K. (2018). Research design: The methodology for interdisciplinary research framework. *Quality & Quantity*, *52*(3), 1209–1225. https://doi.org/10.1007/s11135-017-0513-8

United States Census Bureau. (2020, December 10). *American Community Survey 5-Year Data (2009-2019)*. https://www.census.gov/data/developers/data-sets/acs-5year.html

Uprichard, E., & Dawney, L. (2019). Data Diffraction: Challenging Data Integration in Mixed Methods Research. *Journal of Mixed Methods Research*, *13*(1), 19–32. https://doi.org/10.1177/1558689816674650

US Army Corps of Engineers. (2022). *National Structure Inventory* [Dataset]. https://nsi.sec.usace.army.mil/downloads/

van de Lindt, J., Johnston, B., Wang, W. (Lisa), Crawford, S., Do, T., Yan, G., Dao, T., & Harati, M. (2025). *Midwest Tornado December 2021 Field Study* [Dataset]. DesignSafe-CI. https://doi.org/in preparation

Van De Lindt, J. W., Gillis Peacock, W., & Mitrani-Reiser, J. (2018). *Community resilience-focused technical investigation of the 2016 Lumberton, North Carolina flood: Multi-disciplinary approach* (NIST SP 1230; p. NIST SP 1230). National Institute of Standards and Technology. https://doi.org/10.6028/NIST.SP.1230

van de Lindt, J. W., Gillis Peacock, W., & Mitrani-Reiser, J. (2018). *Community resilience-focused technical investigation of the 2016 Lumberton, North Carolina flood: Multi-disciplinary approach* (NIST SP 1230; p. NIST SP 1230). National Institute of Standards and Technology. https://doi.org/10.6028/NIST.SP.1230

van de Lindt, J. W., Wang, W. (Lisa), Johnston, B., Crawford, P. S., Yan, G., Dao, T., Do, T., Skakel, K., Harati, M., Nguyen, T., Robinson, C., & Croope, S. (in review). Social Susceptibility Driven Longitudinal Tornado Reconnaissance Methodology: 2021 Midwest Quad-State Tornado Outbreak. *ASCE OPEN: Multidisciplinary Journal of Civil Engineering*.

Wang, H., Mostafizi, A., Cramer, L. A., Cox, D., & Park, H. (2016). An agent-based model of a multimodal near-field tsunami evacuation: Decision-making and life safety. *Transportation Research Part C: Emerging Technologies*, *64*, 86–100. https://doi.org/10.1016/j.trc.2015.11.010

Wang, W. "Lisa," Van De Lindt, J. W., Johnston, B., Crawford, P. S., Yan, G., Dao, T., Do, T., Skakel, K., Harati, M., Nguyen, T., Umeike, R., Croope, S., & R. Barbosa, A. (2024). Application of Multidisciplinary Community Resilience Modeling to Reduce Disaster Risk: Building Back Better. *Journal of Performance of Constructed Facilities*, *38*(3), 04024012. https://doi.org/10.1061/JPCFEV.CFENG-4650

Weckwerth, W. (2011). Green systems biology—From single genomes, proteomes and metabolomes to ecosystems research and biotechnology. *Journal of Proteomics*, *75*(1), 284–305. https://doi.org/10.1016/j.jprot.2011.07.010

Weighill, D., Tschaplinski, T. J., Tuskan, G. A., & Jacobson, D. (2019). Data Integration in Poplar: 'Omics Layers and Integration Strategies. *Frontiers in Genetics*, *10*, 874. https://doi.org/10.3389/fgene.2019.00874

Wisner, B., Blaikie, P., Cannon, T., & Davis, I. (1994). *At Risk: Natural Hazards, People's Vulnerability and Disasters*. Taylor & Francis.

Wong-Parodi, G., & Small, M. J. (2021). A Decision-Centered Method to Evaluate Natural Hazards Decision Aids by Interdisciplinary Research Teams. *Risk Analysis*, *41*(7), 1118–1128. https://doi.org/10.1111/risa.13261

Yosso, T. J. (2005). Whose culture has capital? A critical race theory discussion of community cultural wealth. *Race Ethnicity and Education*, *8*(1), 69–91. https://doi.org/10.1080/1361332052000341006

Zapko, K. A., Ferranto, M. L., Brady, C., Corbisello, A., Hill, D., Mullen, R., DeFiore-Golden, P. J., & Martin, L. (2015). Interdisciplinary Disaster Drill Simulation: Laying the Groundwork for Further Research: *Nursing Education Perspectives*, *36*(6), 379–382. https://doi.org/10.5480/14-1544

Papers Used in Data Integration Literature Review

Adad, M., Semlali, E., El-Ayachi, M., & Ibannain, F. (2020). "Supporting land data integration and standardization through the LADM standard: Case of Morocco's country profile MA-LADM." *Land Use Policy*, 97, 104762.

Bolduc, B., Hodgkins, S., Varner, R. K., Crill, P. M., McCalley, C. K., Chanton, J. P., Tyson, G. W., Riley, W. J., Palace, M., Duhaime, M. B., Hough, M. A., IsoGenie Project Coordinators, IsoGenie Project Team, A2A Project Team, Saleska, S., Sullivan, M. B., & Rich, V. I. (2020). "The IsoGenie database: an interdisciplinary data management solution for ecosystems biology and environmental research." *PeerJ*, 8, e9467.

Brink, M., Hengeveld, G. M., & Tobi, H. (2020). "Interdisciplinary measurement: A systematic review of the case of sustainability." *Ecological Indicators,* 112, 106145.

Feng, S., Krueger, A. B., Oppenheimer, M., & Schneider, S. H. (2010). "Linkages among climate change, crop yields and Mexico–US cross-border migration." *Proceedings of the National Academy of Sciences - PNAS*, 107(32), 14257–14262.

Fielding, N. G. (2012). "Triangulation and Mixed Methods Designs: Data Integration with New Research Technologies." *Journal of Mixed Methods Research*, 6(2), 124–136.

Gomes, L. C., Bianchi, F. J .J. A., Cardoso, I. M., Schulte, R. P. O., Arts, B. J. M., & Fernandes Filho, E.I. (2020). "Land use and land cover scenarios: An interdisciplinary approach integrating local conditions and the global shared socioeconomic pathways." *Land Use Policy*, 97, 104723.

Green, S., & Anderson, H. (2019). "Systems science and the art of interdisciplinary integration." *Systems Research and Behavioral Science,* 36 (5), 727-743.

Hinkel, J., & Klein, R. (2009). "Integrating knowledge to assess coastal vulnerability to sea-level rise: The development of the DIVA tool." *Global Environmental Change*, 19(3), 384–395.

Hinkel, J., Lincke, D., Vafeidis, A. T., Perrette, M., Nicholls, R. J., Tol, R. S. J., Marzeion, B., Fettweis, X., Lonescu, C., & Levermann, A. (2014). "Coastal flood damage and adaptation costs under 21st century sea-level rise." *Proceedings of the National Academy of Sciences - PNAS*, 111(9), 3292–3297.

James, P., Jankowska, M., Marx, C., Hart, J., Berrigan, D., Kerr, J., Hurvitz, P., Hipp, J., & Laden, F. (2016). "'Spatial Energetics': Integrating Data From GPS, Accelerometry, and GIS to Address Obesity and Inactivity." *American Journal of Preventive Medicine*, 51(5), 792–800.

Jang, E. E., McDougall, D. E., Pollon, D., Herbert, M., & Russell, P. (2008). "Integrative Mixed Methods Data Analytic Strategies in Research on School Success in Challenging Circumstances." *Journal of Mixed Methods Research*, 2(3), 221–247.

Kathiravelu, P., Sharma, A., Galhardas, H., Roy, P. V., & Veiga, L. (2019) “On-demand big data integration.” *Distributed and Parallel Databases*, 37, 273–295.

Kröger, M., & Schäfer, M. (2016). “Scenario development as a tool for interdisciplinary integration processes in sustainable land use research.” *Futures,* 84A, 64-81.

Kulp, S. A., & Strauss, B. H. (2019). “New elevation data triple estimates of global vulnerability to sea-level rise and coastal flooding.” *Nature Communications*, 10, 4844.

Le Sueur, H., Bruce, I.N., Geifman, N. et al. (2020). “The challenges in data integration – heterogeneity and complexity in clinical trials and patient registries of Systemic Lupus Erythematosus.” *BMC Medical Research Methodology*, 20, 164.

Lindfors, E., van Dam, J. C. J., Lam, C. M. C., Zondervan, N. A., Martins dos Santos, V. A. P., & Suarez-Diez, M. (2018) “SyNDI: synchronous network data integration framework.” *BMC Bioinformatics*, 19, 403.

Muto-Fujita, A., Takemoto, K., Kanaya, S., Nakazato, T., Tokimatsu, T., Matsumoto, N., Kono, M., Chubachi, Y., Ozaki, K., & Kotera, M. (2017). “Data integration aids understanding of butterfly–host plant networks.” *Scientific Reports*, 7, 43368.

Nastar, M., Boda, C.S., & Olsson, L. (2018). “A critical realist inquiry in conducting interdisciplinary research: an analysis of LUCID examples.” *Ecology and Society*, 23(3), 41.

Nateghi, R., Sutton, J., & Murray-Tuite, P. (2019). “The Frontiers of Uncertainty Estimation in Interdisciplinary Disaster Research and Practice.” *Risk Analysis.*

Newgreen, D.F., Landman, K.A., & Osborne, J.M. (2019) Addressing Interdisciplinary Difficulties in Developmental Biology/Mathematical Collaborations: A Neural Crest Example. In Schwarz Q. and Wiszniak S. (Eds.) Neural Crest Cells. Methods in Molecular Biology. (Vol. 1976). Humana Press, New York, NY.

Rigaud, K. K., de Sherbinin, A., Jones, B., Bergmann, J., Clement, V., Ober, K., Schewe, J., Adamo, S., McCusker, B., Heuser, S., Midgley, A. (2018.) *Groundswell: Preparing for Internal Climate Migration.* World Bank, Washington, DC.

Sawada, Y., & Hanazaki, R. (2020). “Socio-hydrological data assimilation: Analyzing human–flood interactions by model–data integration.” *Hydrology and Earth System Sciences*, 24(10), 4777-4791.

Silberstein, M. (2002). Reduction, emergence, and explanation. In P. K. Machamer & M. Silberstein (Eds.), *The Blackwell guide to the philosophy of science*. (pp. 80-107). Cambridge: Blackwell.

Tobi, H. (2014). “Crossing disciplinary boundaries in environmental research: Interdisciplinary engagement across the Slovene research community.” *Measurement,* 48, 228-231.

Uprichard, E., & Dawney, L. (2019). “Data Diffraction: Challenging Data Integration in Mixed Methods Research.” *Journal of Mixed Methods Research*, 13(1), 19–32.

Wang, H., Mostafazi, A., Cramer, L. A., Cox, D., & Park, H. (2016). “An agent-based model of a multimodal near-field tsunami evacuation: Decision-making and life safety.” *Transportation Research Part C: Emerging Technologies*, 64, 86-100.

Weckwerth, W. (2011). “Green systems biology — From single genomes, proteomes and metabolomes to ecosystems research and biotechnology.” *Journal of Proteomics*, 75(1), 284–305.

Weighill, D., Tschaplinski, T., Tuskan, G., & Jacobson, D. (2019). “Data Integration in Poplar: ‘Omics Layers and Integration Strategies.” *Frontiers in Genetics*, 10, 874.

Zhu, P., Zhunag, Q., Archontoulis, S. V., Bernacchi, C., & Müller, C. (2019). “Dissecting the nonlinear response of maize yield to high temperature stress with model-data integration.” *Global Change Biology*, 25 (7), 2470-2484.