

Article

Modeling Community Resilience Under Prolonged Disruption: An Agent-Based Framework Integrating Social Connectivity, Migration, and Policy-Driven Allocation

Joshua Hatfield ¹ , Sudipta Chowdhury ^{2,*} and Ammar Alzarrad ³ 

¹ Department of Computer Science, Marshall University, Huntington, WV 25755, USA; hatfield308@marshall.edu

² Department of Mechanical and Industrial Engineering, Marshall University, Huntington, WV 25755, USA

³ Department of Civil Engineering, Marshall University, Huntington, WV 25755, USA; alzarrad@marshall.edu

* Correspondence: chowdhurys@marshall.edu

Abstract

Communities under prolonged disruptions operate as interconnected socio-technical systems in which the effectiveness of any response depends not only on local conditions but also on the structural relationships that link communities to one another. This study introduces an agent-based response framework for evaluating policy-driven intervention strategies across such systems. Each community is described by its population, economic conditions, and access to critical services, and is linked to other communities through a social connectivity network that defines the pathways for population movement and channels the spread of disruption stress between regions. The agent-based model then tracks how vulnerable each community is by combining its local conditions with the conditions of the communities it is most connected to, and it measures the toll of any disruption through a single social cost metric that weighs lost access to healthcare, retail, and food services. The framework is instantiated using county-level COVID-19 data for Illinois, treated as an exogenous hazard input, and evaluated through Monte Carlo simulation across risk-averse, risk-neutral, risk-seeking, adaptive, and no-aid policy regimes. Compared with the no-aid baseline, the highest-intensity (risk-averse) regime produced the lowest social cost and the highest level of assistance, while all intervention regimes resulted in lower migration. Adaptive managerial decision-making was shown to offer no consistent advantage over simple proactive rules, suggesting that consistency and speed of allocation, rather than sophistication, drive system-wide outcomes.

Keywords: socio-technical system; disruption; managerial decision-making; agent-based model; Monte Carlo simulation



Academic Editor: Domenico Ursino

Received: 14 May 2026

Revised: 22 June 2026

Accepted: 24 June 2026

Published: 29 June 2026

Copyright: © 2026 by the authors.

Licensee MDPI, Basel, Switzerland.

This article is an open access article distributed under the terms and conditions of the [Creative Commons Attribution \(CC BY\)](https://creativecommons.org/licenses/by/4.0/) license.

1. Introduction

Natural disasters, disease outbreaks, and infrastructure failures impose severe and often prolonged burdens on communities, disrupting access to healthcare, food systems, retail services, and the broader supply networks that sustain daily life [1–3]. The consequences of such events extend well beyond the communities at their epicenter. Because modern communities are structurally interdependent, linked through shared infrastructure, economic relationships, and dense social network; disruption originating in one locality propagates outward, placing compounded pressure on neighboring regions and eroding the capacity of entire systems to function [4,5]. When these cascading effects are not contained quickly,

they generate secondary disruptions in communities that were not initially at the center of the event, intensify population displacement, and accelerate the deterioration of critical facilities across a region [6,7]. The scale and pace of this propagation mean that the window for effective intervention is narrow, and the cost of a delayed or misdirected response is high.

This interconnectedness introduces a fundamental complexity into crisis response. The vulnerability of any given community is shaped not only by its own internal conditions, i.e., its population density, economic capacity, and infrastructure, but also by the conditions of the communities it is most closely linked to. A community with moderate internal exposure may nonetheless face severe disruption if it is tightly connected to regions experiencing high disease burden or infrastructure failure. Conversely, a geographically isolated community may be partially insulated from cascading effects even when nearby regions are severely affected. This network-mediated exposure means that resilience and vulnerability are fundamentally relational properties: they cannot be accurately characterized by examining communities in isolation [8,9]. At the same time, the behavioral responses of populations; most notably migration away from high-burden areas alter the distribution of need across the network over time, creating feedback dynamics that further complicate the resource allocation problem [10–14]. Decision-makers operating in these environments must therefore contend with a system that is simultaneously interconnected, dynamic, and highly sensitive to the timing and targeting of intervention.

A substantial body of research has investigated crisis response and resource allocation in disaster and epidemic settings from multiple complementary perspectives. Social vulnerability indices have been developed to classify communities according to their relative susceptibility to hazardous events, thereby providing a structural foundation for prioritization and intervention planning [15,16]. In parallel, willingness-to-pay (WTP) frameworks have been used to quantify the societal burden associated with service disruptions by assigning relative importance weights to different categories of critical infrastructure and essential services [17–19]. Research has also examined how risk perception shapes stakeholder decision-making, demonstrating that the assessment and communication of risk can significantly influence the effectiveness of intervention strategies and crisis-response behavior [20–26]. Building on these insights, resilience-oriented operations research has incorporated risk-averse and chance-constrained formulations to support infrastructure restoration and emergency logistics planning under uncertainty [27–31]. At the modeling level, agent-based approaches have been widely applied to simulate community-level responses to infectious disease outbreaks, enabling the representation of behavioral heterogeneity among individuals and institutions operating under uncertain conditions [32–34]. Complementing these approaches, stochastic simulation and Monte Carlo methods have been employed to evaluate intervention performance and system behavior under parameter uncertainty in epidemic and disaster environments [35,36]. This body of work has advanced the understanding of individual components of crisis response, resilience, and uncertainty-aware decision-making.

Despite this progress, a critical gap persists. Network-oriented and agent-based studies have modeled how disruption spreads across connected populations, and resource-allocation studies have modeled how aid is targeted, but these two lines of work have largely developed separately: allocation frameworks typically assess local need and distribute resources without representing how inter-community connectivity shapes the propagation of disruption and the effectiveness of intervention over time, while connectivity-focused models rarely evaluate competing allocation strategies under that propagation [37–39]. Recent resilience-oriented studies have further emphasized the importance of interconnected socio-technical systems, adaptive infrastructure planning, cascading

regional effects, and network-aware disruption modeling in understanding large-scale crisis behavior [40–59]. This independence assumption understates the degree to which conditions in one community are determined by conditions in adjacent ones, and it limits the ability of decision-support tools to capture the cascading regional dynamics that characterize real-world crises. Furthermore, while risk classification and vulnerability indices provide useful snapshots of community exposure, they rarely inform a systematic comparison of how alternative allocation strategies perform across a range of resource constraints, intervention timings, and connectivity configurations. The result is that decision-makers facing large-scale, regionally cascading disruptions lack a quantitative basis for choosing between competing response strategies and understanding the conditions under which each is most effective.

Epidemic crises provide a representative context in which interdependent infrastructure stress, population movement, and policy-driven intervention occur simultaneously. An infectious disease propagates along the same social and economic pathways that connect communities under normal conditions, meaning that social connectivity is often simultaneously the medium of transmission and the structure through which response resources must flow [60–62]. The Social Connectivity Index (SCI), derived from Meta’s Data for Good platform, quantifies the strength of inter-county social ties. It has been used as a predictor of disease trajectory similarity and behavioral compliance with public health interventions [63,64], and separately as a predictor of inter-area migration destination choice [65,66]. The empirical analysis additionally incorporates publicly available demographic, epidemiological, and connectivity datasets to capture regional variation in population structure, COVID-19 burden, and inter-county social relationships [67–69]. Integrating such connectivity information into a resource allocation framework; rather than treating it as a background feature offers a principled basis for identifying which communities face the greatest networked exposure and directing aid accordingly. The COVID-19 pandemic provides a rich and well-documented empirical setting for this purpose, with county-level records of disease burden, population mobility, and service disruption that reveal systematic disparities rooted in social and infrastructural conditions [70,71]. Critically, this study does not model disease transmission dynamics. COVID-19 case and death counts are instead treated as exogenous inputs derived from historical data, and the framework is designed to analyze how communities and stakeholders respond to observed disruption through migration, resource allocation, and economic adjustment; not how the disease itself spreads.

More specifically, this study addresses the identified gap by developing an agent-based response framework that integrates social connectivity, dynamic migration, risk-attitude-based policy regimes, and a WTP-based social cost metric into a unified simulation environment. The framework is instantiated using county-level COVID-19 data for the state of Illinois and evaluated through Monte Carlo simulation across five policy regimes: risk-averse, risk-neutral, risk-seeking, adaptive, and no-aid baseline. By holding external disruption inputs constant across scenarios and varying only the allocation policy, the study isolates the effect of decision logic on regional outcomes including social cost, population movement, and assistance reach.

The principal methodological contribution is the integration of a dynamic, connectivity-weighted Vulnerability Index within a controlled environment for comparing allocation policies under a common hazard trajectory. The Vulnerability Index combines observed local case burden with connectivity-weighted exposure from socially connected communities and evolves as conditions change across the simulation, and it is embedded in a framework that holds an exogenous hazard trajectory constant while varying only the allocation policy. Prior agent-based, mobility, resilience, and network-oriented studies have each addressed

components of this problem, modeling disruption propagation or aid targeting separately; the contribution here is to combine dynamic network-mediated vulnerability assessment with the systematic comparison of risk-attitude allocation regimes under identical disruption conditions, with connectivity entering as one component of that environment rather than as the dominant driver of the results. This integration enables three findings that the separated literatures do not directly establish. First, the highest-intensity (risk-averse) regime consistently shows the lowest social cost among the intervention regimes across a broad range of modeling assumptions. Second, intervention timing and prioritization efficiency exert a stronger influence on regional outcomes than total aid volume, with diminishing returns emerging once critical community needs are met. Third, adaptive strategies do not consistently improve upon simpler proactive rules, indicating that consistency and speed of allocation, rather than policy complexity, are the structural determinants of resilience under prolonged disruption.

Note that while the empirical instantiation draws on COVID-19 data, the framework is intended to be transferable in principle to other hazard types, including natural disasters, infrastructure outages, and supply-chain shocks with similar propagation patterns, subject to validation in those settings. The remainder of the paper is organized as follows. Section 2 presents the conceptual model and its underlying assumptions. Section 3 develops the formal model formulation, including the Vulnerability Index, policy and decision mechanisms, and the social cost evaluation approach. Section 4 describes the simulation design and case study. Section 5 reports robustness, sensitivity, and temporal results. Sections 6 and 7 discuss the findings and outline limitations and directions for future research.

2. Conceptual Framework

2.1. Overview and Design Rationale

The framework developed in this study is designed to support the evaluation of policy-driven resource allocation strategies in regional systems subject to prolonged external disruption. Its central purpose is not to predict the evolution of a hazard, but to analyze how communities and centralized decision-making agents respond to observed disruption through the reallocation of limited resources, the adjustment of economic activity, and the redistribution of population across an interconnected regional network. This distinction is deliberate. In most real crisis-response environments, stakeholders do not control the hazard itself, instead they act on its observed consequences. The framework is therefore formulated as a response model, in which the disruption signal is treated as an exogenous input and all modeled dynamics concern the behavioral, economic, and mobility responses that follow from it.

Three core design choices shape the structure of the framework. First, communities are represented as nodes in a network rather than as independent units, reflecting the empirical reality that disruption propagates through inter-community connections and that the vulnerability of any given community is partly determined by the conditions of its neighbors. Second, the performance of the system is evaluated through a unified social cost metric that aggregates the burden of service disruption across multiple facility types into a single comparable measure, enabling systematic comparison across communities and policy scenarios. Third, resource allocation is governed by behavioral policy regimes rather than optimization, reflecting the heuristic and institutionally constrained nature of real-world emergency decision-making. Figure 1 summarizes the conceptual structure of the proposed framework, showing how observed hazard inputs, inter-community connectivity, vulnerability assessment, policy allocation, and social cost evaluation interact across sequential decision intervals. Each of these choices is elaborated in the subsections that follow.




Figure 1. Conceptual workflow of the proposed agent-based response framework. The model integrates exogenous hazard inputs, SCI-based social connectivity, vulnerability indexing, policy-driven aid allocation, migration feedback, and WTP-weighted social cost evaluation to assess community resilience outcomes across policy scenarios.

2.2. Community Representation and Networked Interdependence

Each community in the framework is characterized by a set of state variables describing its demographic composition, economic capacity, and the availability of critical service facilities. These variables evolve over discrete decision intervals in response to both external disruption and the policy interventions applied by the centralized agent. Communities are not modeled as isolated entities; they are embedded within a connectivity structure that encodes the strength of social and economic ties between pairs of communities. This structure governs two key dynamics: the propagation of disruption signals across the network, and the movement of population between communities in response to changing local conditions.

The connectivity structure is operationalized using Meta’s Social Connectivity Index (SCI), which quantifies inter-county social ties as a weighted county-to-county network. Each edge represents the relative strength of social connectedness between a pair of counties and provides the structural basis for modeling network-mediated exposure and migration pathways. Each edge in this graph represents the relative strength of interaction between a pair of counties, and the bidirectional nature of the index reflects the symmetric character of social connection. This network provides the structural basis for computing a Vulnerability Index (VI) for each community: a composite measure that combines local disruption intensity with connectivity-weighted exposure from neighboring communities. Communities that are both internally burdened and tightly connected to other high-burden areas receive higher vulnerability scores, reflecting their elevated priority for resource allocation.

Migration is modeled as a sequential, population-conserving redistribution process along SCI-defined connections. For each origin–destination pair, the migration rate is determined by normalized differences in hazard intensity and social cost, the origin community’s baseline migration tendency, the migration scaling parameter, and the aid-adjustment factor. Positive differences direct movement toward communities with lower disruption and social burden, while similar conditions across communities leave baseline migration tendencies as the primary driver. The rate is bounded to $([0, 1])$, and migrants are removed from the origin population and added to the destination population. Within each decision interval, migration is evaluated after population has been updated for growth and observed mortality. Although migration conserves population during redistribution, total regional population may still change because growth and mortality are applied beforehand. These movements subsequently alter service demand, vulnerability, and future allocation conditions across the connected regional system.

2.3. Social Cost as the Performance Metric

The primary outcome measure of the framework is social cost, which refers to the aggregate burden imposed on a community’s population due to disrupted access to critical services. Social cost is evaluated across three categories of facilities: healthcare, retail, and food services. Each category is assigned a WTP weight that reflects its relative importance to community functioning [19]. These weights do not represent direct monetary valuations; instead, they function as relative scaling coefficients that allow heterogeneous service disruptions to be aggregated into a single interpretable metric. Consequently, social cost values reported throughout the manuscript represent aggregated weighted burden units rather than direct monetary currency values and should therefore be interpreted comparatively across policy scenarios rather than as literal dollar-denominated economic losses. A community with a high concentration of healthcare facilities, for example, will incur higher social cost when those facilities are disrupted than a community of comparable size with fewer such facilities because healthcare carries the highest relative weighting in the formulation.

This formulation serves two purposes. At the community level, it translates the operational status of critical facilities into a measure of population-level burden that can be tracked over time. At the system level, it provides a consistent basis for comparing outcomes across communities and across policy scenarios, allowing the analysis to identify which allocation strategies reduce aggregate regional burden most effectively and under what conditions. Social cost is computed at every decision interval, making it a dynamic rather than static measure, capturing how burden accumulates or diminishes as the disruption evolves and as policy interventions take effect.

2.4. Policy Regimes and the Decision-Making Agent

Resource allocation in the framework is performed by a centralized agent that distributes assistance across communities at each decision interval. The agent does not solve an optimization problem; instead, it follows one of several predefined behavioral policy regimes that translate observed community conditions into allocation decisions. This design reflects the institutional reality of emergency management, in which allocation rules are often shaped by organizational priorities, political constraints, and simplified assessment criteria rather than formal optimization.

Four policy regimes are evaluated. The *risk-averse* regime applies the highest assistance-effectiveness posture uniformly across all aided communities, converting available capacity into support at the strongest rate and reflecting a precautionary stance toward worst-case outcomes. The *risk-neutral* regime applies an intermediate effectiveness posture, and the

risk-seeking regime applies the lowest, reflecting a delayed or minimal intervention stance. The three static regimes therefore differ in how effectively a common baseline capacity is translated into assistance, not in how that capacity is targeted across communities. The *adaptive* regime is the only regime that reclassifies individual communities during the simulation, raising the effectiveness posture for communities whose hazard or vulnerability exceeds the system mean. A fifth *no-aid* scenario serves as a counterfactual baseline, capturing system evolution in the complete absence of external intervention. By holding all external inputs constant and varying only the policy regime, the framework enables direct comparison of how different allocation logics shape system-level outcomes under identical disruption conditions.

2.5. Scope and Generalizability

The framework is designed to analyze response dynamics rather than hazard generation. The external disruption is represented as an observed signal; operationalized in this study using historical COVID-19 case and death data for Illinois counties that informs the hazard intensity input to the model at each decision interval. No assumptions are made about how the disruption itself evolves; its trajectory is determined entirely by the historical record. This design choice isolates the analysis from the considerable uncertainties associated with epidemic or disaster forecasting, allowing the framework to focus on the comparative performance of allocation strategies under a common disruption trajectory.

Although the empirical instantiation is grounded in COVID-19 data, the conceptual structure of the framework does not depend on the characteristics of any specific hazard type. The community representation, connectivity structure, social cost formulation, and policy regime architecture are each defined in terms that generalize across disruption contexts. The framework is therefore intended to be transferable in principle, subject to validation with appropriate data inputs, to the analysis of resource allocation under natural disasters, prolonged infrastructure outages, and supply-chain disruptions, that is, settings in which an external event generates cascading service disruptions across an interconnected regional system and in which a centralized agent must allocate limited resources under uncertainty.

3. Model Formulation

3.1. System Structure and State Representation

The proposed framework is formulated as a discrete-time, agent-based system designed to capture the evolution of interconnected communities under conditions of external stress. The model consists of three primary components: (i) community units that represent geographically defined populations, (ii) a centralized decision-making entity responsible for allocating limited resources, and (iii) an exogenous hazard signal that reflects observed disruption intensity. The system evolves over a sequence of discrete decision intervals indexed by t , allowing the model to represent dynamic adjustments in response to changing conditions.

Each community $i \in \mathcal{N}$ is represented as an aggregate entity characterized by a set of state variables that describe its demographic, economic, and infrastructural conditions. These include population ($\text{population}_{i,t}$), economic attributes such as employment and income, and the availability of critical service facilities. The aggregation at the community level is a deliberate modeling choice that enables tractable representation of large-scale systems while preserving the key structural features necessary to analyze system-level responses. Communities are embedded within a connectivity structure that governs interactions across the system. This structure captures interdependencies among communities and

enables the representation of non-local effects, such as migration flows and the propagation of stress signals across connected regions.

The decision-making component is modeled as a centralized agent that observes system conditions and allocates resources across communities at each decision interval. This agent does not optimize a global objective function; rather, it operates according to predefined policy regimes that reflect different behavioral orientations toward risk. This formulation is intended to represent stylized decision-making processes observed in practice, where policy actions are often guided by heuristics or institutional rules rather than fully optimized solutions. By holding the external conditions constant and varying the policy regime, the framework enables systematic comparison of alternative allocation strategies and their implications for system performance.

The external hazard is represented as an exogenous signal that varies across communities and over time, operationalized using observed indicators derived from historical data. Two county-level quantities are drawn from the data at each interval t : the number of newly reported COVID-19 cases $C_{i,t}$ and the number of reported COVID-19 deaths $D_{i,t}$. Hazard intensity is defined from the case count,

$$\text{hazard}_{i,t} = \frac{C_{i,t}}{\text{population}_{i,t}}, \quad (1)$$

the ratio of newly reported cases to the current county population. Deaths enter the model only through the population update,

$$\text{population}_{i,t} = \text{population}_{i,t-1} \cdot g - D_{i,t},$$

where g is the baseline growth factor. Case counts drive the hazard term and, through it, vulnerability and policy response. Deaths reduce the population at risk, which influences migration, vulnerability, and allocation at subsequent intervals through the changed population base. Both series are obtained from daily USAFacts county reports and aggregated to the bimonthly decision intervals of the simulation. The label “observed impact” used in the simulation procedure refers to the aggregated case count $C_{i,t}$ within an interval and is not a separate variable. Case incidence serves as the primary hazard indicator because the framework targets resilience and intervention effectiveness and does not model disease progression.

The framework does not model disease transmission and does not include epidemiological state variables such as susceptible, exposed, infectious, or recovered populations. Observed case incidence is treated as an exogenous disruption signal because rising case prevalence accompanies greater healthcare demand, workforce absenteeism, supply-chain strain, and reductions in service availability. These effects fall on the healthcare, retail, and food-service categories that define the model’s social-cost formulation, which makes case incidence a practical proxy for broader community disruption.

Time progression in the model is governed by discrete decision intervals, during which the system state is updated sequentially. At each interval, population levels are adjusted based on growth and reported deaths $D_{i,t}$, the hazard signal is updated, and derived quantities such as vulnerability and social cost are computed. This sequential update structure provides a transparent and modular representation of system dynamics, enabling the integration of multiple interacting components while maintaining clarity in how state transitions occur over time.

3.2. Networked Risk and Community Dynamics

The framework represents communities as elements of an interconnected system in which local conditions depend on internal characteristics and on interactions with other communities. A network-based representation of social connectivity governs the propagation of external stress signals and the movement of populations across the system, allowing the model to characterize *networked risk*: the exposure of a community is shaped by its own conditions together with those of its connected neighbors. This relational view is consistent with recent evidence that disruption and recovery in community systems propagate through inter-community dependencies and do not remain confined to individual localities [4,6,7].

To quantify community-level exposure, the model defines a *Vulnerability Index (VI)* that combines local disruption burden with connectivity-weighted external signals. For community i at time t ,

$$VI_{i,t} = C_{i,t} \left(\frac{\sum_{j \in N(i)} SCI_{ij,t} (1 + \lambda H_{j,t})}{\sum_{j \in N(i)} SCI_{ij,t}} \right) \quad (2)$$

where $C_{i,t}$ denotes the observed county-level case burden during interval t , $SCI_{ij,t}$ represents the social connectivity between communities i and j , and $H_{j,t}$ is a binary hazard indicator equal to 1 when the hazard intensity in community j exceeds the system median and 0 otherwise. The parameter λ controls the influence of high-risk neighboring communities on the connectivity-weighted exposure term. The neighborhood set $N(i)$ contains the counties connected to community i through the SCI network. Because the SCI data include both inter-county and self-referential ties, the exposure term reflects a combination of local and network-mediated disruption conditions, with $C_{i,t}$ carrying the local burden explicitly and the SCI-weighted sum carrying the influence of connected communities.

The index expresses two sources of risk within a single quantity. Local case incidence captures disruption inside the focal community, and the weighted exposure term captures pressure originating from socially connected communities that are experiencing elevated disruption. A community with high local case burden and strong ties to other high-risk regions receives a larger vulnerability score and is prioritized more aggressively under risk-based allocation. Dividing the weighted exposure term by the total SCI weight prevents communities with large aggregate connectivity from receiving high scores merely because of network size. The hazard indicator $H_{j,t}$ is binary and threshold-based. It separates counties experiencing elevated disruption from those at routine levels and limits sensitivity to short-term fluctuations in reported case incidence, which keeps the exposure term focused on whether a neighbor is currently a high-risk source rather than on the precise magnitude of its case count. The threshold is the system-wide median of hazard intensity at each interval, so classification is relative to the current state of the system and not fixed at an absolute case level. This relative definition suits the comparative design of the study, since a county counts as a high-risk neighbor when its hazard exceeds the prevailing system level at that interval, and the same rule applies across all policy scenarios. A median split partitions counties by relative standing by construction, so during low-disruption intervals roughly half of all counties are labeled high hazard even though their absolute case burden is low. This behavior is intended. The indicator marks which neighbors are comparatively elevated within the network, while the continuous local term $C_{i,t}$ in Equation (2) carries absolute burden, so a county with low local cases receives a low vulnerability score even when several of its neighbors sit above the interval median. The two terms play distinct roles: $H_{j,t}$ encodes relative neighbor standing, and $C_{i,t}$ anchors the index to absolute local conditions.

Relationship to Existing Vulnerability Frameworks

The proposed Vulnerability Index relates to several established approaches in disaster resilience, public health preparedness, and community risk assessment. Widely used examples include the Centers for Disease Control and Prevention (CDC) Social Vulnerability Index (SVI) [72], Cutter’s Social Vulnerability framework [73], and the Baseline Resilience Indicators for Communities (BRIC) [74]. These indices identify populations likely to be disproportionately affected by disasters, disease outbreaks, and infrastructure disruptions, and they have informed preparedness planning and resource targeting across a range of hazard contexts.

The CDC SVI characterizes community susceptibility using demographic and socioeconomic indicators such as income, age distribution, disability status, housing conditions, and transportation access [72]. Cutter’s framework similarly emphasizes how social and economic characteristics shape a community’s capacity to prepare for, respond to, and recover from disruptive events [73]. BRIC broadens this perspective across social, economic, institutional, infrastructural, environmental, and community-capacity dimensions [74]. These frameworks draw on indicators that change slowly, which suits them to characterizing baseline vulnerability and limits their ability to track the evolving conditions of an active crisis.

Recent work has begun to address this temporal limitation. Longitudinal modifications of the CDC SVI extend the index across historical census periods so that vulnerability can be examined as societal composition changes over time [75]. Mobility-based formulations move further in this direction by deriving vulnerability from population movement and updating it across the pre- and post-disaster phases of an event [76]. These developments reflect a broader shift toward vulnerability measures that respond to changing conditions during a disruption.

The Vulnerability Index used here belongs to this emerging class of dynamic measures and is distinguished by its explicit treatment of inter-community connectivity. Vulnerability evolves across the simulation as observed disruption intensity changes, so the index tracks the trajectory of the crisis directly. Inter-community social connectivity enters through the SCI, allowing a community’s score to reflect network-mediated exposure from connected regions. Population redistribution feeds back into both connectivity and local conditions, so migration alters vulnerability trajectories over the course of the simulation. The index complements established vulnerability frameworks by supplying a connectivity-aware, time-varying measure suited to policy evaluation within interconnected socio-technical systems under prolonged disruption.

The SCI provides the structural basis for inter-community interaction. *SCI* values act as weights on a dense county-to-county network that encodes the relative strength of interaction between community pairs [61,63]. To reflect changing system conditions, connectivity weights evolve with population. Connectivity between communities i and j at time $t + 1$ updates as

$$SCI_{ij,t+1} = SCI_{ij,t} \cdot (1 + \Delta_i + \Delta_j), \quad (3)$$

where Δ_i and Δ_j denote proportional population changes in the respective communities. This update is a modeling approximation that encodes the intuition that shifts in population distribution can alter the intensity of inter-community interaction. It introduces a mechanism for adjusting network influence over time and does not assert an empirically estimated relationship.

SCI values are constrained to remain nonnegative after each update and are not explicitly re-normalized between intervals. Normalization instead occurs inside the Vulnerability Index through the weighted-average structure of Equation (2), where the SCI-weighted exposure term is divided by the total SCI weight of each focal community. Meta’s SCI is

a static measure of baseline social connectedness, and the update rule allows population redistribution to modify interaction intensity while preserving the original SCI network as the baseline structure. Counties gaining population exert greater interaction influence, and counties losing population exert less. Because the update is proportional to interval-level population change and bounded below at zero, it keeps connectivity nonnegative while permitting limited temporal adjustment in connectivity strength. The updated SCI values are interpreted as a simulation construct and not as directly observed changes in social relationships. The update changes the relative weights within a focal county's neighborhood, since the destination term Δ_j varies across neighbors, so it is not a pure common rescaling. The internal normalization in Equation (2) removes only the component of the drift that is common to all of a county's edges, the overall scale of its total SCI weight, while preserving the relative differences among neighbors that the weighted average is intended to capture. Larger relative SCI weights therefore continue to exert greater influence on exposure by design, and the nonnegativity constraint keeps the updated weights well defined.

Influence across county pairs is heterogeneous even though the network is dense. Counties with stronger social ties contribute more heavily to vulnerability through the SCI-weighted exposure term. For migration, SCI defines the feasible movement pathways, while migration volume follows from hazard differences, social-cost differences, baseline migration tendencies, and aid conditions. The use of social connectivity to structure relocation pathways is supported by evidence that the Facebook Social Connectedness Index predicts inter-area migration destination choice, including during the COVID-19 period [65], and that social-network structure causally shapes where migrants relocate [66]. The uniform-connectivity ablation reported in Section 5 removes this heterogeneity and isolates whether the results depend on the specific SCI weighting structure.

Migration between connected communities is modeled as a sequential, population-conserving redistribution process. For an origin community i and destination community j , the migration rate is

$$R_{ij,t} = \left(\widetilde{Pos}_{ij,t} + \widetilde{SC}_{ij,t} + B_i \right) \cdot \gamma \cdot A_t \quad (4)$$

where $\widetilde{Pos}_{ij,t}$ is the normalized difference in hazard intensity between counties i and j , $\widetilde{SC}_{ij,t}$ is the normalized difference in social cost, B_i is the baseline migration tendency of community i , realized at each interval as the risk-adjusted value $\widetilde{B}_{i,t}$, γ is the migration scaling parameter, and A_t is the aid-adjustment factor. Migration rates are bounded to $[0, 1]$ to prevent unrealistic population transfers. When two communities present similar hazard and social-cost conditions, the normalized difference terms approach zero and the baseline tendency B_i governs movement, which keeps destination selection well defined when origin and destination share comparable disruption profiles.

The number of migrants transferred along a connection is

$$M_{ij,t} = \left(P_{i,t} - P_{i,t-1}^{assist} \right) R_{ij,t} \quad (5)$$

where $P_{i,t}$ is the current population of community i and $P_{i,t-1}^{assist}$ is the portion of that population assisted in the previous interval. Migration is evaluated before the current-interval assistance step (Algorithm 1), so the most recent available assisted count is the one carried from interval $t - 1$. Subtracting this previously assisted population lowers migration pressure in communities that received effective intervention in the prior interval while keeping the redistribution population-conserving. Positive values of $\widetilde{Pos}_{ij,t}$ and $\widetilde{SC}_{ij,t}$ arise when the origin carries greater hazard burden or social cost than the destination, so movement is directed toward communities with comparatively lower disruption and lower social burden, which serve as the more attractive destinations in the network. Migration is

evaluated only along existing SCI connections, so social connectedness sets the feasible movement pathways and shapes migration volume indirectly through the structure of available destinations. SCI enters the destination choice as an eligibility filter and not as a flow weight: it determines which origin–destination pairs are admissible, while the migrated volume along each admissible pair is set entirely by $R_{ij,t}$ in Equation (4), which contains no SCI term. Two neighbors with identical hazard and social-cost gaps therefore receive equal migration pull regardless of their SCI magnitude.

Algorithm 1: Monte Carlo agent-based simulation procedure

Data: County-level demographic and economic data, SCI network, and observed county-level case counts $C_{i,t}$ and death counts $D_{i,t}$.

Result: Time-series measures of migration, social cost, vulnerability, assistance, and sufficiency.

Initialize communities with population, economic attributes, facility counts, and SCI connections;

foreach *policy scenario* **do**

for $r \leftarrow 1$ **to** N *Monte Carlo replications* **do**

 Reset all community states;

foreach *decision interval* t **do**

Population update (reduce by reported deaths);

$$population_{i,t} = population_{i,t-1} \cdot g - D_{i,t};$$

Connectivity update: adjust $SCI_{ij,t}$ from interval population changes;

Hazard intensity;

$$hazard_{i,t} = \frac{C_{i,t}}{population_{i,t}};$$

Vulnerability Index: compute $VI_{i,t}$ from local case burden $C_{i,t}$ and connectivity-weighted exposure from socially connected communities;

Migration: compute flows along SCI connections as a population-conserving redistribution, using the previous-interval assisted population $P_{i,t-1}^{assist}$ in Equation (5);

Policy stance: classify each community using hazard intensity and $VI_{i,t}$;

Assistance: allocate population and business assistance under the selected policy scenario;

Economic update: compute capability, demand, and sufficiency;

Social cost: compute aggregated social cost from WTP weights and facility counts;

end

 Record simulation outputs for replication r ;

end

 Aggregate results across all N replications;

end

Several safeguards keep the redistribution physically consistent with migration behavior observed in disaster settings [11,14]. Bounding the rate to $[0, 1]$ prevents any single transfer from exceeding the available source population at the time it is evaluated. Migration is applied sequentially, so once migrants move from an origin to a destination, both populations update before the next transfer is computed, which prevents cumulative outflows from driving a county below feasible levels during redistribution. All transfers are within-network, so individuals leaving one county enter another SCI-connected

county and remain in the regional system. Assistance influences migration through the aid-adjustment factor, through the previous-interval assisted population subtracted in Equation (5), and through subsequent changes in social cost and policy conditions; current-interval assistance is computed after migration and therefore affects relocation only at the following interval.

Together, the Vulnerability Index, the dynamic connectivity update, and the migration process define the mechanisms through which networked interactions drive system evolution. They allow the model to represent local and non-local effects and to trace how stress propagates across connected communities and how behavioral responses redistribute population and risk over time.

3.3. Policy and Decision Mechanism

A centralized decision-making agent distributes assistance across communities at each decision interval. The agent applies a set of predefined behavioral rules that represent stylized decision environments in which allocation follows heuristics and institutional priorities rather than a fully specified objective function. At each interval the agent evaluates two indicators for every community: hazard intensity and the Vulnerability Index. Together these characterize the observed stress level of a community. The static regimes do not use these indicators to target individual communities; they hold one effectiveness posture for all aided communities throughout the run. The adaptive regime alone uses the two indicators to reclassify communities over time, as stated in Table 1.

Table 1. Adaptive Policy Classification Rules.

Condition	Assigned Policy
$Hazard_i \geq \overline{Hazard}$ or $VI_i \geq \overline{VI}$	Risk-Averse
$Hazard_i < \overline{Hazard}$ and $VI_i < \overline{VI}$	Risk-Seeking
Three consecutive averse intervals	Risk-Neutral

Under each static regime, every aided community receives the baseline capacity a and the same regime-specific effectiveness band, so the regimes set how strongly that capacity is converted into assistance rather than which communities are favored. The risk-averse band is highest, the risk-neutral band intermediate, and the risk-seeking band lowest, which produces progressively weaker mitigation of social burden across the three. These regimes are behavioral strategies that enable systematic comparison of how the strength of intervention shapes system outcomes, and they are not intended as optimal policies.

The adaptive regime adjusts allocation as system conditions evolve. At each interval it reclassifies every community: a community whose hazard intensity or Vulnerability Index exceeds the current system-wide mean receives a risk-averse posture and more aggressive assistance, and a community below both means receives a risk-seeking posture and lower assistance. A community held in the risk-averse posture for three consecutive intervals is reassigned to a neutral posture, which prevents persistent over-prioritization of the same locations.

Assistance enters the system through population support, business support, and the resulting mitigation of service disruptions, which alter economic capacity and social cost at later intervals and reshape migration incentives across the network. Holding external conditions identical across scenarios isolates the effect of the decision logic itself on system outcomes. Assistance is computed in three steps: each aided county starts from the same intervention capacity, the policy regime sets how effectively that capacity converts into support, and a random draw within the regime's effectiveness range determines the realized assistance for that county and interval. When aid is active, every aided county receives the

same baseline intervention-capacity parameter a , set to $a = 1000$ in the baseline scenario. The parameter a is an abstract intervention-capacity value, not a count of individuals, so the reported “People Assisted” figures are the estimated number of individuals reached after capacity passes through this assistance function, summed across all aided counties; this is why a per-county capacity of $a = 1000$ produces system-wide totals in the tens of thousands. The no-aid scenario sets assistance to zero for every county.

The capacity pool is not divided among counties; each county draws on the full parameter independently. The policy regime then fixes an assistance-effectiveness interval $[L_p, U_p]$, which sets how strongly that capacity is converted into support:

$$[L_p, U_p] = \begin{cases} [0.5, 0.7], & \text{risk-averse,} \\ [0.3, 0.5], & \text{risk-neutral,} \\ [0.1, 0.3], & \text{risk-seeking,} \\ 0, & \text{no-aid.} \end{cases}$$

A higher interval means capacity converts into assistance more effectively. The risk-averse regime therefore has the strongest effectiveness, applied uniformly to all aided counties, and the risk-seeking regime the weakest. For each county i and interval t , a single draw $U(L_p, U_p)$ is taken uniformly from the regime’s effectiveness interval. The assistance computation is discretized into S sub-trials per interval, and the draw is converted into an integer success count

$$K_{i,t} = \lceil S \cdot U(L_p, U_p) \rceil,$$

so $K_{i,t}$ of the S sub-trials deliver effective assistance and the remaining $S - K_{i,t}$ deliver only a low baseline level. A higher effectiveness interval yields a larger $K_{i,t}$ and therefore more effective support. Population assistance is the average over these S sub-trials of the support delivered in each:

$$PA_{i,t} = \frac{K_{i,t}(C_{i,t}a + 0.5(P_{i,t} - C_{i,t})a) + (S - K_{i,t})(0.2 P_{i,t}a)}{S},$$

where $C_{i,t}$ is the number of new cases and $P_{i,t}$ is the current county population. In a successful sub-trial, the case population receives full support and the remaining population receives partial support at coefficient 0.5; in an unsuccessful sub-trial, the whole population receives only the low baseline coefficient 0.2. The corresponding number of individuals effectively reached is

$$P_{i,t}^{assist} = \left\lceil \frac{K_{i,t} (P_{i,t} - C_{i,t}) 0.5}{S} + \frac{(S - K_{i,t}) P_{i,t} 0.2}{S} + C_{i,t} \right\rceil.$$

Business assistance is computed the same way, using an independent draw from the same interval and county business counts in place of population.

3.4. Economic and Social Cost Evaluation with Simulation Design

The evaluation of system performance combines economic capacity modeling with a unified social cost metric. The economic component represents baseline productive capacity and demand, while the social cost formulation quantifies the burden associated with disruptions to critical services. Economic conditions within each community are represented using aggregate measures of capability, demand, and sufficiency. Capability reflects the productive potential of a community based on its population and economic

attributes, while demand represents the level of economic activity required to sustain normal operations. For community i at time t , capability and demand are defined as

$$\text{capability}_{i,t} = \text{population}_{i,t} \cdot (1 - \text{unemploymentRate}_i) \cdot \text{medianIncome}_i, \quad (6)$$

$$\text{demand}_{i,t} = \text{population}_{i,t} \cdot \text{medianIncome}_i \cdot \text{growthFactor}. \quad (7)$$

The ratio of available resources to required demand is captured through a sufficiency measure,

$$\text{sufficiency}_{i,t} = \frac{\text{capability}_{i,t} + \text{businessAssist}_{i,t} + \text{populationAssist}_{i,t}}{\text{demand}_{i,t}} \times 100, \quad (8)$$

which indicates whether a community can maintain economic functionality given internal capacity and external assistance. The assistance terms $\text{businessAssist}_{i,t}$ and $\text{populationAssist}_{i,t}$ are expressed on the same income-equivalent scale as capability so that the ratio is dimensionally consistent.

The primary outcome measure of the framework is *social cost*, which quantifies the burden imposed on a community due to disruptions in access to critical services. Social cost is a weighted aggregation of service disruptions across facility categories, including healthcare, retail, and food systems. For community i at time t ,

$$\text{socialCost}_{i,t} = \left[\sum_k \left(\text{facilityCount}_{k,i} \cdot \text{WTP}_k \cdot T \right) \right] \cdot \left(P_{i,t} - P_{i,t}^{\text{assist}} \right) \cdot A_t, \quad (9)$$

where k indexes facility types, WTP_k is the weight associated with each service category, and T is the duration of the decision interval. The WTP_k values act as relative weighting coefficients that differentiate the social importance of service categories and do not represent direct monetary valuations, which allows heterogeneous service disruptions to be aggregated into a single comparable metric. Because the metric aggregates weighted burden units across all facilities and the full residual exposed population without rescaling, the social-cost values are large in absolute magnitude and carry no standalone units; they are meaningful only in comparison across policy scenarios and intervals.

Aid enters social cost through the residual exposed population $P_{i,t} - P_{i,t}^{\text{assist}}$, so counties with larger effective assistance carry lower residual burden. This residual population is then scaled by the aid-adjustment factor A_t . Under active aid, $A_t = 1000/a$, so the baseline $a = 1000$ gives $A_t = 1.0$. The no-aid scenario is treated as a separate branch: the multiplicative form $1000/a$ does not apply, $P_{i,t}^{\text{assist}} = 0$ because no county receives assistance, and the residual-burden factor is set directly to $A_t = 2.0$ to represent unmitigated disruption pressure. The factor A_t is a model-based adjustment linking the remaining exposed population and intervention capacity to social burden, and does not represent an outage probability or a monetary discount.

4. Simulation Design and Case Study

4.1. Case Study Description

To demonstrate the applicability of the proposed framework, the model is instantiated using a county-level case study of the state of Illinois. This setting provides a diverse set of communities with varying demographic, economic, and connectivity characteristics, enabling evaluation of policy behavior across heterogeneous conditions. The analysis spans the period from March 2020 to May 2023, with the simulation implemented over bimonthly decision intervals, resulting in approximately 20 time steps. This temporal structure allows the model to capture both short-term fluctuations and longer-term system adjustments under sustained external stress.

The empirical instantiation integrates multiple data sources to parameterize the system. Social connectivity between counties is derived from Meta's SCI, which provides a network-based measure of interaction intensity between geographically separated populations. Hazard intensity is operationalized using observed COVID-19 case and death data obtained from USAFacts [77], which serve as externally observed indicators of disruption. Socio-economic attributes, including population, employment, and income characteristics, are obtained from publicly available U.S. Census data [78]. These datasets collectively enable the construction of a data-driven representation of community conditions while preserving the general structure of the proposed framework.

County-level COVID-19 case and mortality data were obtained from USAFacts and aggregated from daily observations to the bimonthly decision intervals used by the simulation. This aggregation was performed to align the temporal resolution of the empirical inputs with the decision-making structure of the model. To reduce the influence of reporting irregularities and uncertainty observed during the earliest stages of the pandemic, the analysis period was restricted to portions of the dataset exhibiting more stable reporting behavior across Illinois counties. Because USAFacts occasionally contains delayed or irregular county-level reporting during the earliest months of the pandemic, no attempt was made to interpolate missing observations. Instead, reported values were aggregated directly within each bimonthly interval, which reduced the influence of short-term reporting artifacts while preserving the observed temporal pattern of disruption. No epidemiological forecasting or statistical smoothing procedures were applied beyond temporal aggregation because the objective of the framework is comparative policy evaluation rather than reconstruction of disease dynamics. Consequently, reported case observations were used directly as exogenous disruption indicators within the simulation.

The Illinois case study provides a realistic and data-rich environment for evaluating how different policy regimes influence system behavior across interconnected communities. By maintaining consistent model structure and parameters across scenarios, the analysis enables direct comparison of policy outcomes under identical external conditions, thereby isolating the impact of decision logic on social cost, migration, and assistance distribution.

4.2. Parameterization

Model parameters are specified to keep conditions consistent across simulation scenarios and to support controlled comparison of alternative policy regimes. All parameters are held constant across scenarios unless explicitly varied, so differences in system outcomes can be attributed to the underlying decision logic and not to exogenous variation in model inputs. The parameters fall into three groups: weighting coefficients used in the social cost formulation, network influence parameters, and behavioral assumptions governing growth and migration. Table 2 summarizes the primary parameters used in the analysis. The WTP values act as relative weighting coefficients that aggregate heterogeneous service disruptions into a single comparable metric and carry no direct monetary interpretation.

The connectivity amplification parameter $\lambda = 0.5$ supplies a moderate level of network influence within the Vulnerability Index. Setting λ to zero removes connectivity amplification, and large values let neighboring high-risk communities dominate local vulnerability calculations. The baseline value places network effects at a level that meaningfully shapes vulnerability while leaving local disruption conditions controlling. The sensitivity analysis in Section 5 shows that the principal policy conclusions hold across alternative values of λ . Migration behavior is governed by the county-specific baseline tendency B_i and the global responsiveness multiplier γ that appear in Equation (4). The baseline tendency is

read from the county-level migration attribute in the initialization file, base_migration_i , and transformed as

$$B_i = \frac{\text{base_migration}_i}{10}, \quad (10)$$

which gives each county its own behavioral baseline for relocation rather than imposing a uniform rate across communities. The baseline tendency entering Equation (4) is the risk-adjusted value

$$\tilde{B}_{i,t} = B_i \eta_{i,t}, \quad \eta_{i,t} \sim \begin{cases} \mathcal{U}(0.8, 1.1), & \text{risk-averse,} \\ \mathcal{U}(0.85, 1.15), & \text{risk-neutral,} \\ \mathcal{U}(0.9, 1.2), & \text{risk-seeking,} \end{cases} \quad (11)$$

so B_i is scaled by a small stochastic factor set by the county's current risk posture. With this substitution, the migration rate $R_{ij,t}$ in Equation (4) combines the normalized hazard-intensity difference $\tilde{Pos}_{ij,t}$, the normalized social-cost difference $\tilde{SC}_{ij,t}$, and $\tilde{B}_{i,t}$, scaled by γ and the aid-adjustment factor A_t . The baseline multiplier is $\gamma = 1.0$, and the aid-adjustment factor follows the implementation

$$A_t = \begin{cases} 2.0, & \text{no-aid policy or aid unavailable,} \\ 1000/a, & \text{otherwise,} \end{cases} \quad (12)$$

with a the baseline intervention capacity of Section 3.3. The rate is bounded to $[0, 1]$ as in Equation (4), and the migrant count along $i \rightarrow j$ follows Equation (5), with the result floored to whole individuals in implementation.

Because migration behavior is context-dependent and difficult to calibrate consistently across crisis settings, B_i and γ are treated as stylized behavioral assumptions rather than directly estimated parameters. The global sensitivity analysis varies the migration multiplier across

$$\gamma \in \{0.75, 1.0, 1.25\}, \quad (13)$$

to evaluate whether the relative policy conclusions remain stable under lower, baseline, and higher responsiveness to disruption gradients.

The service-weighting coefficients $WTP_{\text{health}} = 85$, $WTP_{\text{retail}} = 55$, and $WTP_{\text{food}} = 40$ reflect the relative societal burden of disruptions to each service category. Healthcare carries the highest weight because the resilience literature identifies healthcare systems as critical societal infrastructure whose disruption propagates into public health and emergency response capacity, with downstream economic consequences [79]. Interruptions in medical access also bear directly on morbidity and mortality. Retail and food services carry substantial weights for their contribution to economic continuity and daily community operations, with disruption effects that are generally less immediate than healthcare outages. These values are relative weighting coefficients rather than direct monetary willingness-to-pay estimates; only their ratios affect the social-cost computation, so the absolute magnitudes set an arbitrary common scale. Sensitivity analysis on the weights shows that the absolute social cost magnitude shifts with the chosen values while the relative ranking of policy scenarios holds, which supports the comparative conclusions.

These choices define a baseline scenario against which alternative assumptions are evaluated. The robustness analyses show that absolute outcome magnitudes change under alternative parameter settings while the relative ranking of policy regimes holds, indicating that the primary conclusions do not hinge on any single parameter choice.

Table 2. Sample Major Model Parameters and Selection Rationale.

Parameter	Description	Rationale
$WTP_{health} = 85$	Relative weight for healthcare services	Highest weight, reflecting the role of healthcare access in community stability, emergency response capability, and public health outcomes during prolonged disruptions.
$WTP_{retail} = 55$	Relative weight for retail services	Reflects the role of retail infrastructure in economic continuity and access to essential goods, placed below healthcare in relative importance.
$WTP_{food} = 40$	Relative weight for food services	Reflects the burden of disruptions to food access and supply, treated as substantial and placed below healthcare-related disruptions.
$\lambda = 0.5$	Connectivity-based risk weighting parameter	Moderate amplification value that lets neighboring high-risk communities influence vulnerability while leaving local hazard conditions controlling.
$g = 1.02$	Per-interval growth factor	Applied identically across all scenarios as a fixed reference dynamic rather than an empirical population-growth estimate for Illinois. Because it enters every scenario equally, it does not affect the relative ranking of policy regimes, which is the object of comparison.
γ	Migration scaling parameter	Global multiplier applied to the combined migration signal in Equation (4), setting the overall responsiveness of relocation to disruption gradients. The baseline value is 1.0; sensitivity testing varies $\gamma \in \{0.75, 1.0, 1.25\}$.
B_i	Baseline county migration tendency	County-specific term in Equation (4), set as $B_i = \text{base_migration}_i / 10$ from the initialization file and adjusted by risk attitude at each interval.
$a = 1000$	Baseline resource allocation level	Reference intervention capacity used across policy scenarios and varied during robustness testing to evaluate sensitivity to resource availability.

4.3. Simulation Procedure

The simulation is conducted over a sequence of discrete decision intervals spanning the study period, with each interval representing a bimonthly update of system conditions. For each scenario, the model is initialized using county-level demographic, economic, and connectivity data, and the system state is reset at the beginning of each simulation run to ensure consistency across replications.

At each decision interval, the system evolves through a structured sequence of updates. First, community populations are adjusted to reported deaths. The hazard signal is then updated using the exogenous input data, followed by recalculation of derived quantities, including the Vulnerability Index and connectivity adjustments. Migration flows are subsequently determined based on relative system conditions, allowing population redistribution across communities. The centralized decision-making agent then allocates resources according to the policy regime defined for the scenario, after which economic indicators and social cost are computed for each community. This sequence captures the dynamic interaction between system state, behavioral responses, and policy intervention.

To account for stochastic variability in migration behavior and resource allocation, each scenario is evaluated using a Monte Carlo simulation approach. Specifically, a minimum of $n = 1000$ independent replications are conducted per scenario. For each replication, key outcome variables—including social cost, migration, assistance levels, and economic sufficiency—are recorded at each decision interval. The stochastic elements of the model

are the policy-effectiveness draw $U(L_p, U_p)$ that sets the success count $K_{i,t}$, the analogous independent draw governing business assistance, and the risk-attitude migration multiplier $\eta_{i,t}$ that scales the baseline migration tendency in Equation (11); all are drawn from uniform distributions over their policy-specific ranges. Each of the $n = 1000$ replications per scenario uses an independent random stream, and deterministic inputs (case and death series, facility counts, economic attributes) are held fixed across replications. Interval-level confidence intervals are computed across replications at each decision interval, while scenario-level intervals first collapse each replication across the simulation horizon and then compute the interval across replications, so time points are not treated as independent samples. Aggregated results are summarized using statistical measures such as mean values, confidence intervals, and interquartile ranges (see Section 5) The simulation procedure follows the sequential update logic described in Algorithm 1.

5. Results

5.1. Robustness and Sensitivity Analysis

To evaluate whether the observed policy behavior remains stable under varying modeling assumptions, a series of robustness and sensitivity experiments were conducted. These experiments systematically altered core model parameters related to social cost weighting, resource availability, intervention timing, connectivity structure, migration dynamics, and growth assumptions. Across all experiments, outcomes were evaluated using mean social cost aggregated across Monte Carlo replications ($n = 1000$ per scenario). The relative ordering of policy performance held across configurations, with the risk-averse regime producing the lowest social cost in nearly every tested case. Adaptive allocation did not improve on the simpler risk-averse policy despite its greater decision complexity. The observed policy effects therefore followed from the broader dynamics of the framework and did not depend on any single parameter selection.

Vulnerability Index Threshold Sensitivity. The Vulnerability Index uses the binary indicator $H_{j,t}$ to identify whether a neighboring community is experiencing elevated hazard intensity. In the baseline specification, this indicator is defined relative to the system-wide median hazard intensity at each decision interval. To evaluate whether the reported policy conclusions were sensitive to this median-based threshold, the model was re-run using alternative high-hazard thresholds at the 60th, 75th, and 90th percentiles. Table 3 summarizes the resulting policy stability across these threshold definitions. As expected, increasing the threshold used to define $H_{j,t}$ reduced the mean Vulnerability Index because fewer neighboring counties were classified as high-hazard contributors to network exposure. However, the full policy ranking remained unchanged across all tested thresholds. In each case, the risk-averse regime produced the lowest mean social cost, followed by the adaptive, risk-neutral, risk-seeking, and no-aid scenarios. This indicates that the main policy conclusion is robust to the selected high-hazard threshold and is not an artifact of the original median-based Vulnerability Index specification.

Table 3. Policy ranking stability under alternative Vulnerability Index hazard-threshold definitions.

VI Threshold	Mean VI	Best Policy	Best Policy Unchanged?	Full Ranking Unchanged?
60th percentile	2224.34	Risk-Averse	Yes	Yes
75th percentile	2090.17	Risk-Averse	Yes	Yes
90th percentile	1945.02	Risk-Averse	Yes	Yes

WTP Robustness Analysis. The WTP weighting structure influenced the computation of social cost by modifying the relative importance assigned to healthcare, retail, and food sys-

tems. To evaluate whether policy outcomes were sensitive to these assumptions, alternative weighting schemes were examined, including equal-weight settings, healthcare-dominant weighting, and food-dominant weighting. In Figure 2, the x-axis represents the alternative service-weighting configurations used within the social cost formulation, including the baseline weighting structure, equal weighting across service categories, healthcare-dominant weighting, and food-dominant weighting. The y-axis represents the mean social cost aggregated across all Monte Carlo replications for each policy regime, where lower values indicate reduced cumulative disruption burden across the modeled communities. Figure 2 showed that changes in WTP priorities altered the absolute magnitude of social cost but did not change the relative ordering of the policy regimes. Across all weighting configurations, the risk-averse strategy consistently produced the lowest social cost, whereas the no-aid baseline generated substantially higher values. Under healthcare-dominant weighting, overall social cost increased across all scenarios because disruptions to medical infrastructure received higher penalties. Food-dominant weighting produced lower overall cost magnitudes than the healthcare-focused settings. Because the policy ordering held across every weighting configuration, the comparative performance of the regimes was insensitive to the selected weighting structure, which matters here because the WTP values are illustrative coefficients and not empirically calibrated economic valuations.




Figure 2. Mean social cost across alternative WTP weighting configurations.

Aid Stress Testing. To evaluate sensitivity to intervention capacity, the per-county capacity parameter a was varied across three levels: $a = 500$, $a = 1000$, and $a = 2000$. This parameter is applied to each aided county rather than divided from a shared regional pool, so it sets per-county intervention capacity, not a total budget. In Figure 3, the x-axis represents the per-county capacity parameter a in effect during each simulation interval, with 500, 1000, and 2000 corresponding to low-, moderate-, and high-capacity scenarios. The y-axis represents the mean social cost aggregated across all Monte Carlo replications, where lower values indicate greater overall system resilience. Figure 3 presents the resulting social cost outcomes across the evaluated policy regimes. The results demonstrated a strong inverse relationship between available aid and system-wide social cost. Increasing aid availability substantially reduced social cost across all intervention-based strategies, whereas the no-aid baseline remained unchanged because no assistance was allocated under that scenario. Under low-budget conditions, the differences among intervention

policies became more pronounced, with the risk-averse allocation strategy producing the largest reduction in social burden relative to the alternative regimes. The reduction in social cost also exhibited diminishing marginal effects as aid availability increased. The transition from 500 to 1000 units of aid produced a substantially larger reduction in social cost than the transition from 1000 to 2000. This diminishing pattern follows in part from the reciprocal aid-adjustment factor $A_t = 1000/a$, which scales social cost inversely with a , so the change from 500 to 1000 moves A_t from 2.0 to 1.0 while the change from 1000 to 2000 moves it only from 1.0 to 0.5. The model does not separately represent the progressive satisfaction of critical needs, so the result is read as a property of the capacity-scaling specification rather than as an independently estimated saturation effect.




Figure 3. Mean social cost under varying aid levels.

Intervention Timing Sensitivity. Intervention timing represented an important factor in the crisis response process because delayed action allowed stress conditions to propagate across interconnected communities before mitigation measures were introduced. To evaluate this effect, aid allocation was delayed by multiple decision intervals prior to intervention. In Figure 4, the x-axis represents the number of delayed decision intervals before aid distribution begins. A value of “Delay = 0 Intervals” indicates that intervention started immediately at the onset of the disruption, whereas “Delay = 2 Intervals” and “Delay = 5 Intervals” indicate that aid allocation was postponed for two and five simulation decision intervals, respectively. The y-axis represents the mean social cost aggregated across all Monte Carlo replications. As shown in Figure 4, delayed intervention consistently increased social cost across all policy regimes. Immediate intervention (0-interval delay) produced the lowest observed social cost, whereas delays of two and five intervals progressively worsened overall system outcomes. The no-aid baseline remained unchanged because intervention was never initiated under that scenario. The widening separation among the policy trajectories showed that delayed response let vulnerability and migration pressure intensify before corrective action took effect, an effect most pronounced in the risk-seeking and neutral regimes. The risk-averse strategy held the lowest social cost across all timing conditions, though its advantage narrowed as delays grew. The magnitude of these timing effects was comparable to that of total aid availability, particularly in highly connected environments where stress propagated quickly across communities.

Network Ablation Experiments Social connectivity is a central mechanism in the framework for representing inter-community influence and migration. Two ablation experiments

isolated its contribution to system behavior: connectivity amplification was removed by setting $\lambda = 0$, and the SCI-based weights were replaced with uniform connectivity across communities. In Figure 5, the x-axis gives the connectivity assumption applied in the Vulnerability Index calculation. “Full” denotes the complete SCI-based structure, “Lambda Zero” sets $\lambda = 0$, and “Uniform” replaces the SCI weights with equal values across communities. The y-axis gives mean social cost aggregated across Monte Carlo replications. The relative ordering of the policy regimes held across all three configurations, and the social cost of the intervention-based strategies varied only modestly whether the full network, uniform connectivity, or no amplification was applied.




Figure 4. Mean social cost under varying intervention delay intervals.




Figure 5. Mean social cost under alternative network connectivity assumptions.

These ablations also bear on the choice of threshold for $H_{j,t}$. Setting $\lambda = 0$ collapses the connectivity factor in Equation (2) to one, so the Vulnerability Index reduces to the local case burden $C_{i,t}$ and the neighbor-hazard contribution vanishes; replacing the SCI weights with uniform values keeps the indicator but removes the heterogeneous weighting that scales each neighbor’s contribution. Both are stronger perturbations of the exposure mechanism than a shift in the median threshold, yet the policy ranking held under each,

so the comparative conclusions do not depend on the specific exposure classification, including the median that defines $H_{j,t}$, or on the specific connectivity weighting. The SCI component still captures the non-local interactions and migration-driven redistribution that would otherwise be absent from the model, and the interpretation of this result for the role of social connectivity is developed in Section 6.

Global Parameter Sensitivity. A broader sensitivity analysis was conducted by simultaneously varying multiple model parameters, including the connectivity amplification coefficient (λ), migration multipliers, growth factors, and baseline aid levels. In Figure 6, the x-axis represents the parameter configurations evaluated during the global sensitivity analysis, including variations in λ , migration multipliers, growth factors, and baseline aid levels. The y-axis represents the mean social cost aggregated across all Monte Carlo replications. Each plotted point corresponds to the average system-wide social cost under the specified parameter configuration. Figure 6 showed that variations in λ , migration multipliers, and growth factors produced comparatively smaller changes in social cost than variations in aid allocation levels. Baseline aid availability generated the largest shifts in system outcomes across all evaluated policy regimes, while migration and connectivity parameters affected behavior to a limited extent relative to the intervention-based effects. The no-aid scenario stayed comparatively stable across all configurations, since in the absence of intervention system behavior followed the exogenous hazard conditions. The risk-averse regime maintained the lowest social cost across every parameter variation.




Figure 6. Global parameter sensitivity analysis across multiple model assumptions.

Key Insights. These experiments showed that the relative performance of the policy regimes held across a broad range of modeling assumptions and parameter configurations. The risk-averse strategy consistently produced the lowest or near-lowest social cost regardless of variations in weighting structures, aid availability, intervention timing, network connectivity assumptions, or global parameter settings. Although the absolute magnitude of system outcomes varied across experimental conditions, the comparative ordering of the policy regimes remained highly consistent throughout the analysis. The sensitivity experiments further indicated that the observed effectiveness of proactive intervention was not strongly dependent on any single parameter specification or modeling assumption. Several recurring behavioral patterns also emerged across the experiments, including diminishing

marginal reductions in social cost at higher aid levels, substantial sensitivity to intervention timing, and comparatively lower sensitivity to migration and connectivity scaling parameters. These patterns suggested that the dominant system behavior was influenced more strongly by intervention structure than by secondary parameter variation within the explored parameter space. Across the robustness analysis, the framework behaved consistently, and early, proactive allocation produced more resilient system behavior under sustained external stress.

5.2. Temporal and Aggregate Policy Effects

To examine how alternative policy regimes influenced system evolution, a controlled simulation experiment was conducted across multiple scenarios under identical initial conditions, parameter settings, and exogenous inputs. Each scenario represented a distinct allocation strategy, enabling direct comparison of how different decision logics shaped system behavior over time. Time-series trends for social cost, migration, and population assistance were analyzed across all scenarios using Monte Carlo simulation outputs ($n = 1000$), with uncertainty represented through 95% confidence intervals.

In Figure 7, the x-axis represents the sequential bimonthly decision intervals spanning March 2020 through May 2023, while the y-axis represents the mean cumulative social cost aggregated across Monte Carlo replications. Lower values indicate reduced disruption burden and greater overall system stability. Figure 7 illustrated the temporal evolution of social cost across policy regimes. A clear separation between scenarios emerged early in the simulation and persisted throughout the study horizon. The risk-averse regime consistently produced the lowest social cost across all decision intervals, while the no-aid baseline yielded the highest values. These results suggested that applying the highest-intensity intervention posture from the outset reduced the accumulation of system stress over time, whereas limited or absent intervention allowed disruption effects to propagate across the connected system.




Figure 7. Mean Social Cost Across Policy Scenarios (95% CI).

In Figure 8, the x-axis represents the sequential bimonthly simulation decision intervals, while the y-axis represents the average number of individuals relocating between communities during each interval. Higher values indicate increased population displacement in response to elevated system stress. Figure 8 presented migration dynamics across policy scenarios. The no-aid condition resulted in the highest levels of population movement, indicating stronger relocation responses under elevated stress conditions. All intervention regimes reduced migration substantially relative to no-aid, which indicates that sustained

intervention stabilized local conditions and lowered mobility-driven responses, with differences among the intervention regimes comparatively small. Across all scenarios, realized migration stayed a small fraction of total regional population, on the order of a fraction of one percent per interval, so the mechanism did not generate implausibly large displacement.




Figure 8. Migration Across Policy Scenarios (95% CI).

In Figure 9, the x-axis represents the sequential bimonthly simulation decision intervals, while the y-axis represents the average number of individuals receiving assistance across all Monte Carlo replications. Higher values indicate greater levels of policy-driven intervention and aid distribution. Figure 9 showed the evolution of population receiving assistance across policy regimes. The risk-averse scenario consistently delivered the highest level of assistance throughout the simulation horizon, while risk-seeking and no-aid strategies provided substantially lower levels of support. The relatively stable separation between assistance trajectories indicated that the observed policy effects emerged from sustained differences in allocation behavior rather than isolated short-term fluctuations. Across the temporal results, the higher-intensity intervention postures produced more stable system behavior over time and delivered more assistance system-wide; under the static regimes this reflects a stronger effectiveness posture applied uniformly to all aided counties rather than spatial targeting of the most vulnerable.

To quantify the overall magnitude of policy effects, outcomes were also aggregated across all decision intervals and Monte Carlo replications. Table 4 summarizes the average system behavior associated with each policy regime under identical experimental conditions.

The aggregate results quantified substantial differences in system performance across policy regimes. Among the intervention regimes, the risk-averse posture produced the lowest social cost and the highest assistance. Relative to the no-aid baseline, the risk-averse social cost is approximately 69% lower and intervention migration is approximately 49% lower, but both no-aid gaps are partly mechanical: the no-aid scenario applies $A_t = 2.0$ against $A_t = 1.0$ for active aid, which doubles the no-aid residual-burden factor in Equation (8) and the no-aid migration rate in Equation (4) before any assisted-population effect. The among-intervention comparisons, which all share $A_t = 1.0$, isolate the behavioral effect and are the basis for the policy conclusions. The adaptive, risk-neutral, and risk-seeking regimes showed comparatively similar outcome ranges despite their differences in allocation behavior. Economic sufficiency remained relatively stable across all scenarios, indicating that short-term system instability was driven more strongly by

disruption propagation and migration dynamics than by large variations in aggregate economic capacity. The no-aid sufficiency interval rounds to zero because no assistance or unemployment-mitigation effects are introduced in that scenario, leaving sufficiency nearly deterministic under the fixed exogenous input trajectory. The sufficiency intervals for the intervention scenarios are likewise very small because sufficiency is dominated by the deterministic capability and demand terms. Migration differs only marginally across the four intervention regimes because the dominant drivers, the exogenous hazard and social-cost gradients, are common to all of them, while the regime enters the migration rate only through the smaller risk-attitude multiplier $\eta_{i,t}$ and the shared aid-adjustment factor.




Figure 9. Population Assistance Across Policy Scenarios (95% CI).

Table 4. Scenario-level summary of model outcomes across Monte Carlo replications.

Scenario	Social Cost (Mean $\pm$ 95% CI)	Migration (Mean $\pm$ 95% CI)	People Assisted (Mean $\pm$ 95% CI)	Sufficiency (Mean $\pm$ 95% CI)
No Aid	$2.304 \times 10^{16} \pm 3.223 \times 10^{11}$	$30,802.02 \pm 2.39$	0.00 ± 0.00	95.28 ± 0.000000
Adaptive	$7.482 \times 10^{15} \pm 2.497 \times 10^{12}$	$15,674.95 \pm 1.43$	$43,737.97 \pm 11.80$	95.83 ± 0.000070
Risk Averse	$7.195 \times 10^{15} \pm 2.501 \times 10^{12}$	$15,730.28 \pm 1.38$	$49,728.94 \pm 11.79$	96.01 ± 0.000044
Risk Neutral	$7.870 \times 10^{15} \pm 2.487 \times 10^{12}$	$15,695.18 \pm 1.35$	$42,363.85 \pm 11.76$	95.89 ± 0.000044
Risk Seeking	$8.539 \times 10^{15} \pm 2.487 \times 10^{12}$	$15,658.60 \pm 1.24$	$35,035.33 \pm 11.74$	95.77 ± 0.000044

To evaluate whether the observed differences between policy regimes were statistically meaningful, pairwise comparisons were conducted by bootstrap resampling of the scenario-level replication outcomes, drawing 10,000 resamples per comparison; the reported 95% confidence intervals and p -values in Table 5 are derived from the resulting bootstrap distribution of the mean difference. Although the differences among the intervention-based policies were substantially smaller than the difference between any intervention policy and the no-aid baseline, all pairwise comparisons reached statistical significance. With $n = 1000$ replications and deliberately separated effectiveness bands, this significance primarily reflects the internal stability of the simulation under its assumptions rather than evidence of real-world policy difference, and it should be read alongside the effect sizes in Table 4.

Three structural patterns emerged from the simulation results. First, the Vulnerability Index varied only modestly across scenarios, which indicates that vulnerability tracked exogenous disruption intensity and the underlying network structure more than the policy in force, with policy effects surfacing downstream through migration dynamics and assistance distribution. Second, economic sufficiency held steady even where social cost and migra-

tion differed sharply, which locates the short-term instability in disruption propagation and population mobility rather than in aggregate economic capacity. Third, the adaptive regime did not outperform the risk-averse strategy; the direct comparison in Table 5 shows the adaptive regime carried significantly higher social cost than the risk-averse regime, so within the modeled system the added policy complexity yielded no gain over a consistently applied strong intervention posture, which was enough to stabilize interconnected community dynamics under sustained external stress.

Table 5. Pairwise comparisons for mean social cost across different policy scenarios.

Comparison	Mean Difference	95% CI	p-Value	Significant
Risk Averse vs. Risk Neutral	-6.746×10^{14}	$[-6.781 \times 10^{14}, -6.711 \times 10^{14}]$	<0.001	Yes
Risk Averse vs. Risk Seeking	-1.344×10^{15}	$[-1.347 \times 10^{15}, -1.340 \times 10^{15}]$	<0.001	Yes
Risk Averse vs. No Aid	-1.585×10^{16}	$[-1.5855 \times 10^{16}, -1.5849 \times 10^{16}]$	<0.001	Yes
Risk Averse vs. Adaptive	-2.861×10^{14}	$[-2.870 \times 10^{14}, -2.840 \times 10^{14}]$	<0.001	Yes
Adaptive vs. Risk Neutral	-3.880×10^{14}	$[-3.916 \times 10^{14}, -3.845 \times 10^{14}]$	<0.001	Yes
Adaptive vs. Risk Seeking	-1.057×10^{15}	$[-1.061 \times 10^{15}, -1.054 \times 10^{15}]$	<0.001	Yes

6. Discussion

The findings of this study indicate that intervention structure plays a central role in shaping regional resilience outcomes during prolonged disruptions. Across all experimental settings, proactive and risk-averse allocation strategies consistently produced lower social cost than reactive or no-intervention approaches. More importantly, this relative ordering remained stable under substantial variation in aid availability, intervention timing, weighting assumptions, migration parameters, and network connectivity structures. Such consistency suggests that the observed policy behavior was not driven by isolated parameter choices, but instead reflected persistent characteristics of the modeled response dynamics.

One noteworthy finding is that the adaptive policy did not consistently outperform the simpler risk-averse allocation strategy. Although the adaptive regime continuously adjusts its prioritization decisions in response to evolving hazard and vulnerability conditions, the results suggest that rapid and consistent intervention may be more important than dynamic reassessment in this particular context. Because the adaptive policy periodically reallocates attention away from previously prioritized communities as conditions change, assistance may become distributed more broadly across the system before the most vulnerable locations fully stabilize. The static risk-averse strategy instead applies the highest effectiveness posture uniformly to all aided communities for the full horizon, which sustains a higher assistance level system-wide and yields lower overall social cost across most experimental conditions. This finding suggests that under prolonged disruption, consistency of intervention may provide greater resilience benefits than increased policy complexity.

One of the most important observations emerging from the analysis concerns the role of intervention timing. Delays in aid deployment consistently amplified system-wide social burden, particularly in highly connected environments where migration and inter-community influence accelerated the propagation of stress conditions. In practical terms, the results imply that emergency response effectiveness depends not only on the quantity of available resources, but also on how rapidly those resources can be mobilized and distributed. Even relatively moderate delays allowed vulnerability conditions to accumulate across connected communities before mitigation measures became active. This behavior highlights the operational importance of pre-positioned resources, adaptive logistics planning, and rapid coordination mechanisms during the early stages of disruption.

The aid experiments revealed a similarly important pattern. Although increasing aid availability reduced social cost across all intervention-based policies, the rate of improvement declined at higher resource levels. The largest reductions in social burden occurred during the transition from severely constrained aid conditions to moderate resource availability, whereas additional increases beyond that point produced comparatively smaller improvements. This diminishing-return behavior suggests that resilience gains are strongly associated with the ability to stabilize critical needs early in the disruption cycle. For policy-makers operating under fiscal or logistical constraints, the findings indicate that improving allocation efficiency and prioritization strategy may generate greater system-level benefit than simply increasing total resource expenditure.

The study also demonstrated the importance of representing interconnected regional dynamics within crisis-response models. Traditional aid-allocation frameworks frequently evaluate communities independently, implicitly assuming that disruptions remain spatially localized. The proposed framework instead modeled migration-driven redistribution effects and inter-community influence through the SCI-based connectivity structure. The network ablation experiments showed that the relative policy ordering remained stable across alternative connectivity assumptions, so the connectivity mechanisms shape where and how disruption propagates across communities without altering the comparative ranking of allocation strategies. Such effects are particularly relevant in real-world disruptions where population displacement, infrastructure outages, and resource shortages propagate across neighboring regions rather than remaining confined to a single locality.

The network ablation experiments yield a further insight. The SCI-based connectivity structure shapes migration pathways, vulnerability propagation, and the spatial distribution of disruption, yet the ranking of policy performance held when alternative connectivity assumptions were introduced. Social connectivity therefore operates as a moderating mechanism within the framework and not as the dominant driver of resilience outcomes: it governs how disruption spreads and which communities absorb secondary effects, while the effectiveness of a given intervention strategy depends more strongly on allocation posture, intervention timing, and resource availability. The network component contributes explanatory detail about where and when burden accumulates, and the stability of the rankings across specifications indicates that the advantage of proactive intervention reflects a system-level characteristic of the model rather than a feature of the SCI formulation.

From a managerial perspective, the framework provides several practical insights for emergency managers and resilience planners. First, the consistent performance of the risk-averse strategy suggests that applying a strong intervention posture early and uniformly may substantially reduce aggregate regional burden during prolonged disruptions; under the adaptive regime, which alone targets by hazard and vulnerability, the same level of targeting did not improve on this uniform posture. Second, the strong sensitivity to intervention timing indicates that investments in rapid deployment capability, emergency coordination infrastructure, and decentralized resource staging may produce significant resilience benefits even when total aid availability is limited. Third, the relatively lower sensitivity to migration and connectivity scaling parameters suggests that intervention structure and operational prioritization exert stronger influence on overall system performance than secondary behavioral assumptions within the explored parameter space.

Beyond emergency response operations, the framework may also support broader resilience-oriented planning applications. Because the modeling environment captures cascading stress propagation and migration-driven redistribution dynamics, it may, with further validation, be adaptable to disaster-relief coordination, emergency shelter allocation, healthcare resource distribution, evacuation planning, infrastructure restoration prioritization, and regional recovery strategies. The framework is intended for disruptions

characterized by prolonged cascading effects, including hurricanes, wildfires, prolonged utility outages, climate-related displacement events, and large-scale supply-chain disruptions. In such settings, localized intervention decisions frequently generate secondary consequences across interconnected regions, making system-wide evaluation increasingly important for policy development.

These results should be read in light of the framework's intended scope. The model is a controlled simulation environment for comparing the relative performance of allocation strategies under common disruption conditions, and its outputs are meaningful as comparisons across policy regimes rather than as forecasts of realized social cost, migration, or assistance for any specific event. Within that scope, the study demonstrated that integrating interconnected social dynamics with intervention-oriented policy modeling can provide meaningful insight into regional resilience behavior under sustained external stress.

7. Conclusions, Limitations, and Future Research Directions

This study presented a simulation-based framework for evaluating crisis-response and aid-allocation strategies under conditions characterized by interconnected social stress, migration-driven redistribution, and cascading regional disruption. The framework is intended as a controlled environment for comparing allocation strategies under common disruption conditions rather than as a predictive or operational planning tool. Within this scope, the model isolates the effect of decision logic by holding external disruption inputs constant across scenarios, so its outputs are meaningful as comparisons across policy regimes rather than as forecasts of realized social cost, migration, or assistance for any specific region or event. The reported social-cost, migration, and assistance values therefore support conclusions about the relative ordering of allocation strategies and the conditions under which each performs best, and they are not intended as absolute estimates for a particular hazard. By integrating aid allocation dynamics, inter-community connectivity, migration behavior, and infrastructure-sensitive social cost modeling within a unified simulation environment, the study provided a mechanism for analyzing how intervention structure and timing influence regional resilience outcomes during prolonged disruptions.

From a theoretical perspective, the study contributed to the growing body of resilience and crisis-response literature by explicitly incorporating interconnected community dynamics into the evaluation of intervention policies. Where conventional allocation models frequently assess communities independently, the proposed framework captured how localized disruptions propagated through migration and influence mechanisms across connected regions. The results demonstrated that policy effectiveness cannot be evaluated solely through localized outcomes because intervention decisions generated secondary redistribution effects throughout the broader system. The study therefore extended existing resilience-oriented modeling approaches by emphasizing the importance of cascading regional interactions in shaping system-wide social burden under sustained stress conditions.

The study also contributed to the practical understanding of crisis-response policy design and emergency management planning. Across all robustness and sensitivity experiments, proactive and risk-averse intervention strategies consistently produced lower social cost than reactive or no-aid approaches. The findings further demonstrated that intervention timing exerted substantial influence on system performance, particularly within highly connected environments where vulnerability conditions propagated rapidly across communities. The aid experiments also revealed diminishing marginal improvements at higher resource levels, suggesting that effective prioritization and early deployment may provide greater resilience benefit than simply increasing total aid availability. These findings provide actionable insights for emergency managers, resilience planners, and policymakers seeking to design resource-allocation strategies under constrained operational conditions.

The framework is designed for potential transferability across a range of disruption contexts in which cascading regional interactions influence recovery behavior. Potential applications include disaster-relief coordination, emergency shelter allocation, health-care resource distribution, evacuation planning, infrastructure restoration prioritization, and regional resilience assessment under hurricanes, wildfires, prolonged utility outages, supply-chain disruptions, and climate-related displacement events. Because the framework captures interconnected migration and stress-propagation dynamics, it is intended to support evaluation of policy behavior in large-scale disruptions where localized decisions generate broader regional consequences, pending validation in those settings.

Several limitations nevertheless define the scope of the present study and motivate future research. The migration and social influence mechanisms adopted in the framework were intentionally stylized and were not calibrated using region-specific empirical behavioral datasets. Similarly, the WTP weighting structure represented relative infrastructure importance rather than empirically validated economic valuation. The Social Connectedness Index is derived from Facebook friendship ties and under-represents older, lower-income, and less digitally connected populations; because these groups are often among the most vulnerable, the connectivity-weighted component of the Vulnerability Index may understate their networked exposure, which qualifies the interpretation of the network-mediated term and its use in prioritization. The connectivity update in Equation (3) adjusts SCI weights with population change but holds the network topology fixed across the simulation horizon, so the set of inter-county ties does not itself evolve, and the weight update is a modeling approximation rather than an empirically estimated change in social connectedness. Real-world disruptions also alter transportation accessibility, communication infrastructure, and mobility behavior in ways the present formulation does not capture. The framework also does not impose a shared regional budget constraint of the form $\sum_i a_{i,t} \leq B_t$: each aided county receives the baseline capacity a independently, so the regimes differ in how effectively capacity is converted into assistance rather than in how a scarce common pool is divided among competing counties. The comparison is therefore one of allocation posture and effectiveness rather than constrained allocation under a fixed budget, and incorporating an explicit budget constraint with county-level decision variables is an important extension. In addition, the analysis focused on comparative policy behavior under synthetic stress scenarios rather than operational forecasting for a specific geographic region or disaster event. The present study demonstrates the framework using a single empirical setting, county-level COVID-19 data for Illinois, and transferability to other hazard types such as hurricanes, wildfires, infrastructure outages, and supply-chain disruptions remains a design goal to be validated with case studies specific to those contexts.

Future research may extend the framework in several directions. Incorporating empirically calibrated mobility and demographic datasets would improve the realism of migration and redistribution dynamics. Network topology adaptation, in which inter-county ties form or dissolve as infrastructure accessibility changes during a disruption, could extend the fixed-edge structure used here. Additional extensions may include stochastic hazard evolution, multi-agent coordination behavior, adaptive intervention strategies, and integration with real-time infrastructure and mobility data streams. Further work may also examine competing policy objectives, equity-oriented allocation mechanisms, and cross-regional coordination strategies under uncertain resource availability.

Author Contributions: Conceptualization, S.C.; methodology, S.C. and J.H.; formal analysis, J.H. and S.C.; investigation, J.H. and A.A.; Data Curation, J.H.; writing—original draft preparation, J.H.; writing—review and editing, S.C. and A.A. All authors have read and agreed to the published version of the manuscript.

Funding: This research received no external funding.

Data Availability Statement: The raw data underlying this study are drawn from publicly available sources. The processed input files, parameter settings, and simulation code are not publicly deposited at this time because they support ongoing follow-on work, and will be made available by the authors on reasonable request.

Conflicts of Interest: The authors declare no conflicts of interest.

References

1. Yang, P.; Ng, T.L.; Cai, X. Reward-Based Participant Management for Crowdsourcing Rainfall Monitoring: An Agent-Based Model Simulation. *Water Resour. Res.* **2019**, *55*, 8122–8141. [\[CrossRef\]](#)
2. Guidotti, R.; Chmielewski, H.; Unnikrishnan, V.; Gardoni, P.; McAllister, T.; Van De Lindt, J. Modeling the resilience of critical infrastructure: The role of network dependencies. *Sustain. Resilient Infrastruct.* **2016**, *1*, 153–168. [\[CrossRef\]](#) [\[PubMed\]](#)
3. Chowdhury, S.; Zhu, J. An Agent-Based Model for Resilient Power Distribution System Restoration in Dynamic Environments. In *Computing in Civil Engineering 2021*; American Society of Civil Engineers: Reston, VA, USA, 2022. [\[CrossRef\]](#)
4. Anderson, M.; Brunner, L.; Logan, T. Functional Isolation: The compounding burden on community capacity from cascading infrastructure failures and access disruption. *Int. J. Disaster Risk Reduct.* **2025**, *118*, 105190. [\[CrossRef\]](#)
5. Dulin, S.; Mitoulis, S.A.; Bredikhin, A.; Treyz, E.; Leung, B.; Dykes, J.; Karpeles, O.; Gurav, S.; Karhunen, A.; Linkov, I. Rethinking infrastructure design from component failure to systemic resilience. *Nat. Commun.* **2025**, *16*, 9681. [\[CrossRef\]](#) [\[PubMed\]](#)
6. Yin, K.; Li, B.; Mostafavi, A. Deep Learning-driven Community Resilience Rating based on Intertwined Socio-Technical Systems Features. *arXiv* **2025**, arXiv:2311.01661. [\[CrossRef\]](#)
7. Liu, C.F.; Mostafavi, A. Network dynamics of community resilience and recovery: New frontier in disaster research. *Int. J. Disaster Risk Reduct.* **2025**, *123*, 105489. [\[CrossRef\]](#)
8. Kim, K.; Kang, J.Y.; Hwang, C. Identifying Indicators Contributing to the Social Vulnerability Index via a Scoping Review. *Land* **2025**, *14*, 263. [\[CrossRef\]](#)
9. Wainger, L.A.; Weber, M.A.; Price, E.W. Modifying Social Vulnerability Indices to Complement Physical Exposure Risk Analyses. *Nat. Hazards Rev.* **2025**, *26*, 04024048. [\[CrossRef\]](#)
10. Perez, L.; Dragicevic, S.; Gaudreau, J. A geospatial agent-based model of the spatial urban dynamics of immigrant population: A study of the island of Montreal, Canada. *PLoS ONE* **2019**, *14*, e0219188. [\[CrossRef\]](#) [\[PubMed\]](#)
11. Nourali, Z.; Shortridge, J.E.; Bukvic, A.; Shao, Y.; Irish, J.L. Simulation of Flood-Induced Human Migration at the Municipal Scale: A Stochastic Agent-Based Model of Relocation Response to Coastal Flooding. *Water* **2024**, *16*, 263. [\[CrossRef\]](#)
12. Marini, M.; Chokani, N.; Abhari, R.S. Agent-Based Model Analysis of Impact of Immigration on Switzerland's Social Security. *J. Int. Migr. Integr.* **2019**, *20*, 787–808. [\[CrossRef\]](#)
13. Hassani-Mahmooei, B.; Parris, B.W. Climate change and internal migration patterns in Bangladesh: An agent-based model. *Environ. Dev. Econ.* **2012**, *17*, 763–780. [\[CrossRef\]](#)
14. Tanaka, T.; Koga, A.; Yokomatsu, M. Agent-based intra-regional relocation model considering spatial local amenity for urban planning-based flood risk management: Assessing the impact of urban development on flood exposure. *J. Flood Risk Manag.* **2024**, *17*, e13035. [\[CrossRef\]](#)
15. Wood, E.; Sanders, M.; Frazier, T. The practical use of social vulnerability indicators in disaster management. *Int. J. Disaster Risk Reduct.* **2021**, *63*, 102464. [\[CrossRef\]](#)
16. Roderick, T.; Tremblay, M.C.; Kohli, R.; Castellanos, A.; Pengetnze, Y. When Vulnerability Drives Action: Designing Forward-Looking Frameworks for Disaster Preparedness. *Inf. Syst. Res.* **2026**; ahead of print. [\[CrossRef\]](#)
17. Steigenberger, C.; Flatscher-Thoeni, M.; Siebert, U.; Leiter, A.M. Determinants of willingness to pay for health services: A systematic review of contingent valuation studies. *Eur. J. Health Econ.* **2022**, *23*, 1455–1482. [\[CrossRef\]](#) [\[PubMed\]](#)
18. He, S.; Anderson, E.; Rucker, D. Measuring Willingness to Pay: A Comparative Method of Valuation. *J. Mark.* **2023**, *88*, 50–68. [\[CrossRef\]](#)
19. Sullivan, M.; Collins, M.T.; Schellenberg, J.; Larsen, P.H. *Estimating Power System Interruption Costs: A Guidebook for Electric Utilities*; Lawrence Berkeley National Laboratory: Berkeley, CA, USA, 2018. [\[CrossRef\]](#) [\[PubMed\]](#)
20. Cheng, Y.; Fang, S.; Yin, J. The effects of community safety support on COVID-19 event strength perception, risk perception, and health tourism intention: The moderating role of risk communication. *Manag. Decis. Econ.* **2022**, *43*, 496–509. [\[CrossRef\]](#) [\[PubMed\]](#)
21. Wachinger, G.; Renn, O.; Begg, C.; Kuhlicke, C. The Risk Perception Paradox-Implications for Governance and Communication of Natural Hazards. *Risk Anal. Off. Publ. Soc. Risk Anal.* **2012**, *33*, 1049–1065. [\[CrossRef\]](#) [\[PubMed\]](#)
22. Siegrist, M.; Luchsinger, L.; Bearth, A. The Impact of Trust and Risk Perception on the Acceptance of Measures to Reduce COVID-19 Cases. *Risk Anal.* **2021**, *41*, 787–800. [\[CrossRef\]](#) [\[PubMed\]](#)

23. Cheli, S.; Lam, W.W.T.; Estapé, T.; Winterling, J.; Bahcivan, O.; Andritsch, E.; Weis, J.; Centeno, I.; Serpentine, S.; Farkas, C.; et al. Risk perception, treatment adherence, and personality during COVID-19 pandemic: An international study on cancer patients. *Psycho-Oncol.* **2022**, *31*, 46–53. [\[CrossRef\]](#) [\[PubMed\]](#)
24. Miller, M.; Chowdhury, S.; Alzarrad, A.; Hossain, N.U.I. Synergetic Decision-Making: Analyzing the Interplay of Human Behavior and Physical Infrastructure in Emergency Evacuations via an Analytical Approach. *Glob. J. Flex. Syst. Manag.* **2024**, *25*, 785–803. [\[CrossRef\]](#)
25. Eby, L.T.; Robertson, M.M.; Facticeau, D.B.; Anker, J.G. Risky business: Understanding the association between objective COVID-19 occupational risk features and worker subjective risk perceptions. *J. Occup. Organ. Psychol.* **2025**, *98*, e12431. [\[CrossRef\]](#)
26. Hatab, A.A.; Lagerkvist, C.J.; Esmat, A. Risk perception and determinants in small & medium-sized agri-food enterprises amidst the COVID-19 pandemic: Evidence from Egypt. *Agribusiness* **2021**, *37*, 187–212. [\[PubMed\]](#)
27. Yilmaz, O.F.; Guan, Y.; Yilmaz, B.G. A hybrid risk management framework for resilient medical supply chain design in the post-pandemic era. *J. Clean. Prod.* **2025**, *491*, 144854. [\[CrossRef\]](#)
28. Faiz, T.I.; Harrison, K.W. A risk-averse stochastic optimization model for community resilience planning. *Socio-Econ. Plan. Sci.* **2024**, *92*, 101835. [\[CrossRef\]](#)
29. Elçi, O.; Noyan, N.; Bulbul, K. Chance-constrained stochastic programming under variable reliability levels with an application to humanitarian relief network design. *Comput. Oper. Res.* **2018**, *96*, 91–107. [\[CrossRef\]](#)
30. Quddus, M.A.; Chowdhury, S.; Marufuzzaman, M.; Yu, F.; Bian, L. A two-stage chance-constrained stochastic programming model for a bio-fuel supply chain network. *Int. J. Prod. Econ.* **2018**, *195*, 27–44. [\[CrossRef\]](#)
31. Wang, W.; Yang, K.; Yang, L.; Gao, Z. Distributionally robust chance-constrained programming for multi-period emergency resource allocation and vehicle routing in disaster response operations. *Omega* **2023**, *120*, 102915. [\[CrossRef\]](#)
32. Gomez, J.; Prieto, J.; Leon, E.; Rodríguez, A. INFEKTA—An agent-based model for transmission of infectious diseases: The COVID-19 case in Bogotá, Colombia. *PLoS ONE* **2021**, *16*, e0245787. [\[CrossRef\]](#) [\[PubMed\]](#)
33. Jalayer, M.; Orsenigo, C.; Vercellis, C. CoV-ABM: A stochastic discrete-event agent-based framework to simulate spatiotemporal dynamics of COVID-19. *arXiv* **2020**, arXiv:2007.13231. [\[CrossRef\]](#)
34. Silva, P.C.L.; Batista, P.V.C.; Lima, H.S.; Alves, M.A.; Guimarães, F.G.; Silva, R.C.P. COVID-ABS: An agent-based model of COVID-19 epidemic to simulate health and economic effects of social distancing interventions. *Chaos Solitons Fractals* **2020**, *139*, 110088. [\[CrossRef\]](#) [\[PubMed\]](#)
35. Fong, S.J.; Li, G.; Dey, N.; Crespo, R.G.; Herrera-Viedma, E. Composite Monte Carlo decision making under high uncertainty of novel coronavirus epidemic using hybridized deep learning and fuzzy rule induction. *Appl. Soft Comput.* **2020**, *93*, 106282. [\[CrossRef\]](#) [\[PubMed\]](#)
36. Wakeland, W.; Homer, J. Addressing Parameter Uncertainty in a Health Policy Simulation Model Using Monte Carlo Sensitivity Methods. *Systems* **2022**, *10*, 225. [\[CrossRef\]](#)
37. Hinch, R.; Probert, W.J.M.; Nurtay, A.; Kendall, M.; Wymant, C.; Hall, M.; Lythgoe, K.; Cruz, A.B.; Zhao, L.; Stewart, A.; et al. OpenABM-Covid19—An agent-based model for non-pharmaceutical interventions against COVID-19 including contact tracing. *PLoS Comput. Biol.* **2021**, *17*, e1009146. [\[CrossRef\]](#) [\[PubMed\]](#)
38. Blakely, T.; Thompson, J.; Bablani, L.; Andersen, P.; Ouakrim, D.A.; Carvalho, N.; Abraham, P.; Boujaoude, M.A.; Katar, A.; Akpan, E.; et al. Determining the optimal COVID-19 policy response using agent-based modelling linked to health and cost modelling: Case study for Victoria, Australia. *JAMA Health Forum* **2021**, *2*, e211749. [\[CrossRef\]](#) [\[PubMed\]](#)
39. Wood, A.D.; Berry, K. COVID-19 transmission in a resource dependent community with heterogeneous populations: An agent-based modeling approach. *Econ. Hum. Biol.* **2024**, *52*, 101314. [\[CrossRef\]](#) [\[PubMed\]](#)
40. Allen, J.E.; Davis, A.F.; Hu, W.; Owusu-Amankwah, E. Residents' willingness-to-pay for attributes of rural health care facilities. *J. Rural Heal.* **2015**, *31*, 7–18. [\[CrossRef\]](#) [\[PubMed\]](#)
41. An, L.; Linderman, M.; Qi, J.; Shortridge, A.; Liu, J. Exploring Complexity in a Human–Environment System: An Agent-Based Spatial Model for Multidisciplinary and Multiscale Integration. *Ann. Assoc. Am. Geogr.* **2005**, *95*, 54–79. [\[CrossRef\]](#)
42. Anatolitis, V.; Welisch, M. Putting renewable energy auctions into action – An agent-based model of onshore wind power auctions in Germany. *Energy Policy* **2017**, *110*, 394–402. [\[CrossRef\]](#)
43. Beaudoin, A.; Isaac, A.G. Direct and indirect transmission of avian influenza: Results from a calibrated agent-based model. *J. Econ. Interact. Coord.* **2023**, *18*, 191–212. [\[CrossRef\]](#)
44. Dislich, C.; Hettig, E.; Salecker, J.; Heinonen, J.; Lay, J.; Meyer, K.M.; Wiegand, K.; Tarigan, S. Land-use change in oil palm dominated tropical landscapes—An agent-based model to explore ecological and socio-economic trade-offs. *PLoS ONE* **2018**, *13*, e0190506. [\[CrossRef\]](#) [\[PubMed\]](#)
45. Haer, T.; Botzen, W.J.W.; de Moel, H.; Aerts, J.C.J.H. Integrating Household Risk Mitigation Behavior in Flood Risk Analysis: An Agent-Based Model Approach. *Risk Anal.* **2017**, *37*, 1977–1992. [\[CrossRef\]](#) [\[PubMed\]](#)
46. Kadinski, L.; Salcedo, C.; Boccelli, D.L.; Berglund, E.; Ostfeld, A. A Hybrid Data-Driven-Agent-Based Modelling Framework for Water Distribution Systems Contamination Response during COVID-19. *Water* **2022**, *14*, 1088. [\[CrossRef\]](#)

47. Kapeller, M.L.; Jäger, G. Threat and Anxiety in the Climate Debate—An Agent-Based Model to Investigate Climate Scepticism and Pro-Environmental Behaviour. *Sustainability* **2020**, *12*, 1823. [\[CrossRef\]](#)
48. MacGee, J.; Pugh, T.M.; See, K. The heterogeneous effects of COVID-19 on Canadian household consumption, debt and savings. *Can. J. Econ. Rev. Can. D'Econ.* **2022**, *55*, 54–87. [\[CrossRef\]](#) [\[PubMed\]](#)
49. Mastroeni, L.; Naldi, M.; Vellucci, P. Personal Finance Decisions with Untruthful Advisors: An Agent-Based Model. *Comput. Econ.* **2023**, *61*, 1477–1522. [\[CrossRef\]](#)
50. Miller, B.W.; Breckheimer, I.; McCleary, A.L.; Guzmán-Ramirez, L.; Caplow, S.C.; Jones-Smith, J.C.; Walsh, S.J. Using stylized agent-based models for population-environment research: A case study from the Galápagos Islands. *Popul. Environ.* **2010**, *75*, 279–287. [\[CrossRef\]](#) [\[PubMed\]](#)
51. Murphy, K.J.; Ciuti, S.; Kane, A. An introduction to agent-based models as an accessible surrogate to field-based research and teaching. *Ecol. Evol.* **2020**, *10*, 12482–12498. [\[CrossRef\]](#) [\[PubMed\]](#)
52. Nguyen, L.H.; Tran, B.X.; Nguyen, N.P.; Phan, H.T.T.; Bui, T.T.; Latkin, C.A. Mobilization for HIV Voluntary Counseling and Testing Services in Vietnam: Clients' Risk Behaviors, Attitudes and Willingness to Pay. *AIDS Behav.* **2016**, *20*, 848–858. [\[CrossRef\]](#) [\[PubMed\]](#)
53. Palmer, S.; Ciubotariu, I.; Ofori, R.; Saenz, M.; Ellison, B.; Prescott, M.P. School Nutrition Stakeholders Find Utility in MealSim: An Agent-Based Model. *J. Nutr. Educ. Behav.* **2024**, *56*, 361–369. [\[CrossRef\]](#) [\[PubMed\]](#)
54. Park, D.; Ryu, D. Supply Chain Ethics and Transparency: An agent-based model approach with Q-learning agents. *Manag. Decis. Econ.* **2022**, *43*, 3331–3337. [\[CrossRef\]](#)
55. Ramezani-Doroh, V.; Najafi-Ghobadi, S.; Karimi, F.; Rangchian, M.; Hamidi, O. Prediction of COVID-19 patients' participation in financing informal care using machine learning methods: Willingness to pay and willingness to accept approaches. *BMC Health Serv. Res.* **2024**, *24*, 796. [\[CrossRef\]](#) [\[PubMed\]](#)
56. Sengupta, R.; Vankeerberghen, G.; Wen, R.; Rao, J.; Chen, Y. GIS-enabled historiography to determine travel routes during the Western Han period via agent-based models and least-cost path analysis. *Trans. GIS* **2023**, *27*, 1090–1103. [\[CrossRef\]](#)
57. Um, S.; Adhikari, S. Considerations in Bayesian agent-based modeling for the analysis of COVID-19 data. *arXiv* **2022**, arXiv:2212.04045. [\[CrossRef\]](#)
58. Williams, B.C.; Criddle, K.R.; Kruse, G.H. An agent-based model to optimize transboundary management for the walleye pollock (*Gadus chalcogrammus*) fishery in the Gulf of Alaska. *Nat. Resour. Model.* **2021**, *34*, e12305. [\[CrossRef\]](#)
59. Zhang, W.; Liu, S.; Osgood, N.; Zhu, H.; Qian, Y.; Jia, P. Using simulation modelling and systems science to help contain COVID-19: A systematic review. *Syst. Res. Behav. Sci.* **2022**, *40*, 207–234. [\[CrossRef\]](#) [\[PubMed\]](#)
60. Pellegrino, M.; Lombardo, G.; Cagnoni, S.; Poggi, A. High-Performance Computing and ABMS for High-Resolution COVID-19 Spreading Simulation. *Future Internet* **2022**, *14*, 83. [\[CrossRef\]](#)
61. Pullano, G.; Alvarez-Zuzek, L.G.; Colizza, V.; Bansal, S. Characterizing US Spatial Connectivity and Implications for Geographical Disease Dynamics and Metapopulation Modeling: Longitudinal Observational Study. *JMIR Public Health Surveill.* **2025**, *11*, e64914. [\[CrossRef\]](#) [\[PubMed\]](#)
62. Balcan, D.; Colizza, V.; Gonçalves, B.; Hu, H.; Ramasco, J.J.; Vespignani, A. Multiscale mobility networks and the spatial spreading of infectious diseases. *Proc. Natl. Acad. Sci. USA* **2009**, *106*, 21484–21489. [\[CrossRef\]](#) [\[PubMed\]](#)
63. Kuchler, T.; Russel, D.; Stroebel, J. JUE Insight: The geographic spread of COVID-19 correlates with the structure of social networks as measured by Facebook. *J. Urban Econ.* **2022**, *127*, 103314. [\[CrossRef\]](#) [\[PubMed\]](#)
64. Allcott, H.; Boxell, L.; Conway, J.; Gentzkow, M.; Thaler, M.; Yang, D. Polarization and public health: Partisan differences in social distancing during the coronavirus pandemic. *J. Public Econ.* **2020**, *191*, 104254. [\[CrossRef\]](#) [\[PubMed\]](#)
65. Zou, H. Metropolitan Social Networks and Migration: Insights From Facebook Connections. *J. Reg. Sci.* **2025**, *65*, 1321–1337. [\[CrossRef\]](#)
66. Blumenstock, J.E.; Chi, G.; Tan, X. Migration and the Value of Social Networks. *Rev. Econ. Stud.* **2025**, *92*, 97–128. [\[CrossRef\]](#)
67. U.S. Census Bureau. American Community Survey (ACS). 2026. Available online: <https://www.census.gov/programs-surveys/acs.html> (accessed on 12 June 2026).
68. Meta Data for Good. Social Connectedness Index, 2026. Dataset Hosted by the Humanitarian Data Exchange. Available online: <https://data.humdata.org/dataset/social-connectedness-index> (accessed on 12 June 2026).
69. USAFacts. How Did COVID-19 Affect People in Illinois? 2026. Data Source. Available online: <https://usafacts.org/answers/how-did-covid-19-affect-people-in-the-us/state/illinois/> (accessed on 12 June 2026).
70. Khan, S.S.; Krefman, A.E.; McCabe, M.E.; Petito, L.C.; Yang, X.; Kershaw, K.N.; Pool, L.R.; Allen, N.B. Association between county-level risk groups and COVID-19 outcomes in the United States: A socioecological study. *BMC Public Health* **2022**, *22*, 81. [\[CrossRef\]](#) [\[PubMed\]](#)
71. Fielding-Miller, R.K.; Sundaram, M.E.; Brouwer, K. Social determinants of COVID-19 mortality at the county level. *PLoS ONE* **2020**, *15*, e0240151. [\[CrossRef\]](#) [\[PubMed\]](#)

72. Flanagan, B.E.; Gregory, E.W.; Hallisey, E.J.; Heitgerd, J.L.; Lewis, B. A Social Vulnerability Index for Disaster Management. *J. Homel. Secur. Emerg. Manag.* **2011**, *8*, 0000102202154773551792. [CrossRef]
73. Cutter, S.L.; Boruff, B.J.; Shirley, W.L. Social Vulnerability to Environmental Hazards. *Soc. Sci. Q.* **2003**, *84*, 242–261. [CrossRef]
74. Hazards Vulnerability and Resilience Institute. Baseline Resilience Indicators for Communities (BRIC) Index. 2025. Available online: <https://doi.org/10.3886/E209669V2> (accessed on 12 June 2026).
75. Adepoju, O.E.; Kiaghadi, A. Measuring Historic and Longitudinal Social Vulnerability in Disaster-Prone Communities: A Modification to the Centers for Disease Control and Prevention Social Vulnerability Index (CDC-SVI). *Disaster Med. Public Health Prep.* **2023**, *17*, e368. [CrossRef] [PubMed]
76. Patel, J.A.; Lor, M.A.; Jeph, P.; Jiang, R.; Luis, S.; Shibasaki, R.; Shyu, M.L.; Chen, S.C. Mobility vulnerability index for community vulnerability assessment in disaster response. *npj Nat. Hazards* **2026**, *early access*. [CrossRef]
77. USAFacts. Accessible Government Data and Statistics. 2026. Available online: <https://usafacts.org/> (accessed on 12 June 2026).
78. U.S. Census Bureau. Census Bureau Data. 2026. Available online: <https://data.census.gov/> (accessed on 12 June 2026).
79. Emami, S.G.; Lorenzoni, V.; Turchetti, G. Towards Resilient Healthcare Systems: A Framework for Crisis Management. *Int. J. Environ. Res. Public Health* **2024**, *21*, 286. [CrossRef] [PubMed]

Disclaimer/Publisher’s Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.