Reliability Boundaries under Synthetic Post-Earthquake Road Disruption:
A Reproducible Graph-Based Monte Carlo Study of Emergency-Response Accessibility in Beykoz, Istanbul

Faramarz Kowsari

Independent Researcher, Istanbul, Türkiye

ORCID: 0000-0003-1692-0453 | Correspondence: faramarzkowsari@gmail.com

Preprint manuscript based on frozen v1.0.0 computational evidence

Persistent identifier	Record
OSF retrospective Open-Ended Registration DOI	10.17605/OSF.IO/ZTJXK
Zenodo software version DOI	10.5281/zenodo.21903851
Zenodo software concept DOI	10.5281/zenodo.21903850
Frozen scientific commit	8078af69eff047d7da71beac3b5b3b125a121d1a

Registration note: the OSF record is retrospective and was created after the principal analyses; it is not a preregistration.

Abstract

This study investigates how responder availability and resource-allocation strategy affect incident accessibility when portions of a road network become unavailable following severe disruption. The road network of Beykoz, Istanbul, was represented as a computational model, and road disruptions were simulated across multiple stochastic scenarios. The principal reliability measure was the 80/80 boundary, defined as the disruption level at which the probability of maintaining access to at least 80% of incidents reaches 80%.

The results showed that the estimated 80/80 boundary increased from 0.1425 with 12 responders to 0.1625 with 32 responders. The study also compared a sequential Greedy assignment strategy with a Global Minimum-Cost assignment approach. Across all 102 tested cases, the Global Minimum-Cost method produced lower total response time than the Greedy method after Holm correction. However, the two approaches yielded identical accessibility-boundary point estimates at all six tested resource levels.

These findings suggest that improved allocation can make more efficient use of responders and road connections that remain available, but optimization alone cannot restore accessibility once the underlying road network loses connectivity. The disruption-severity parameter used in this study is a synthetic network-stress variable and should not be interpreted as earthquake magnitude, peak ground acceleration (PGA), a physical damage measure, or an operational safety threshold.

Keywords: emergency response; road network resilience; graph reliability; Monte Carlo simulation; resource allocation; min-cost flow; earthquake accessibility; reproducible research; OpenStreetMap; Istanbul

1. Introduction

When an earthquake or another severe disruption occurs, emergency-response capacity cannot be evaluated solely by counting ambulances, responders, or available equipment. A system may possess sufficient personnel and resources, yet still fail to deliver timely assistance if critical roads and transportation links become unavailable. What ultimately matters is whether a viable connection remains between locations where incidents occur and the resources intended to serve them. When one route is lost, alternative paths become essential if assistance is to reach affected locations within a useful response time. This creates an important distinction between resource availability and resource accessibility: increasing the number of responders alone may provide little benefit when the transportation network can no longer connect those responders to the incidents that require assistance.

A second problem arises once available responders must be distributed among multiple incidents occurring across different locations and times. Emergency assignments may involve incidents that occur simultaneously as well as those that arise sequentially, while the available response resources remain limited. One possible strategy is to make assignments incrementally, selecting an apparently suitable responder for each incident as the decision is encountered. Another possibility is to consider the set of assignments more globally and seek a configuration that minimizes the overall cost of response. In this study, these two perspectives are represented by a sequential Greedy assignment strategy and a Global Minimum-Cost assignment strategy. The comparison was motivated by a practical question: under degraded network conditions, how much difference can the assignment policy itself make to the performance of the response system?

Previous research has examined post-earthquake medical accessibility in Istanbul [1], emergency road-network accessibility and repair planning [2], probabilistic traffic-management decisions under uncertain post-earthquake damage [3], coordinated disaster-management optimization [4], and ambulance response optimization under uncertain demand and travel conditions [5]. More recent work has expanded this landscape through digital-twin and common-operating-picture approaches to seismic disaster management in Kocaeli, Türkiye [13], GIS-based assessment of post-earthquake road blockage and emergency routing after the 2023 Türkiye earthquakes [14], probabilistic road-blockage and resource-based restoration optimization [15,16], transportation-resilience models for emergency medical response under disrupted infrastructure [17,18], and integrated seismic vulnerability assessment of urban road networks [19]. Together, this literature shows that physical accessibility, network degradation, resource availability, and allocation policy are distinct but interacting determinants of emergency performance.

The question examined here is narrower. The study asks what happens when three constraints interact: road connectivity deteriorates, response resources are limited, and disruption increases stochastically. In particular, it tests whether improved allocation can compensate for progressive loss of network accessibility, or whether there is a point at which the topology of the disrupted road network becomes the dominant limitation regardless of assignment strategy. This distinction is important because an allocation algorithm can determine how existing resources are used, but it cannot create a physical connection between two locations once the network no longer provides one.

To examine this question, the Beykoz road network in Istanbul was represented computationally and subjected to controlled stochastic disruption scenarios. Responder availability was varied, and the resulting response system was evaluated under both Greedy and Global Minimum-Cost assignment policies. Rather than seeking a prescriptive answer about how emergency agencies should operate during an actual earthquake, the study uses this framework to examine the interaction between network degradation, resource availability, and allocation policy. Its contribution is methodological and computational: to identify what optimization can improve when viable network connections remain, and where its benefits become constrained by loss of connectivity itself.

2. Materials and Methods

2.1 Study Area and Road-Network Representation

Beykoz, a district of Istanbul, Türkiye, was selected as the study area. Its drivable road network was constructed from OpenStreetMap data. OSMnx-based procedures were then used to derive edge-level speed and travel-time attributes, allowing the road system to be represented as a computational network on which responder movement and accessibility could be evaluated [6,20].

The resulting network should be understood as a computational representation of the physical road system rather than a reconstruction of an operational emergency-response platform. Travel-time information was not obtained from live ambulance, police, fire-service, or emergency-dispatch telemetry. The purpose of the study was to examine the behavior of a response network under controlled disruption, rather than to reproduce the exact behavior of an existing emergency service.

The broader question was whether a response system confronted with an unanticipated network disruption could still make effective use of the resources available to it, and whether a more deliberate assignment strategy could improve the amount of assistance delivered through the surviving transportation network. The analysis therefore focuses on the interaction between network accessibility and allocation policy rather than on forecasting the behavior of a specific operational agency.

2.2 Synthetic Road-Network Disruption

Road disruption was modeled probabilistically. For each directed edge e, the probability of failure at disruption severity s was defined as:

$$p_e(s) = s \times \exp(-d_e / \lambda)$$

where d_e denotes the distance between the edge and the disruption center and λ controls the spatial decay of disruption probability. The default value of λ was 8 km. The disruption center was defined from the mean coordinate of the graph nodes. Under this formulation, road segments closer to the defined center have a higher probability of failure at a given value of s, while the probability decreases with distance.

The parameter s is deliberately synthetic. It is a network-stress parameter used to control the severity of simulated disruption and must not be interpreted as earthquake magnitude, peak ground acceleration, an engineering damage index, or an operational emergency threshold.

2.3 Synthetic Incident and Responder Generation

Incidents were generated synthetically. A uniform spatial distribution was avoided because real emergencies do not normally occur at evenly spaced locations or according to a conveniently balanced pattern. Incident generation was therefore weighted toward the central part of the study region, with demand decreasing with distance at a characteristic scale of 6 km.

Each incident was assigned a synthetic severity value from 1 to 5. These values provided relative weights for differences in incident demand within the computational experiments; they were not intended to represent a validated clinical or engineering severity scale. Responder locations were sampled from graph nodes not selected as incidents. Incidents and response resources were therefore represented within the same graph structure, allowing accessibility and assignment decisions to depend on the surviving road topology.

2.4 Coupling of Simulation Scenarios

A central design consideration was to avoid comparing completely unrelated random worlds when disruption severity or responder availability changed. Common random numbers were therefore used to couple scenarios.

Failed-edge sets were constructed in a nested manner. If an edge failed at a given disruption level, that failure remained present at higher severity levels, while additional failures could accumulate as severity increased. This produced monotone degradation within a simulated world rather than replacing one random network realization with an unrelated one at each severity.

Responder sets were coupled in a similar way. Smaller responder configurations were constructed as prefixes of a common maximum responder pool. Consequently, comparisons across responder counts were based on nested resource sets rather than independently resampled populations. This coupling reduced avoidable Monte Carlo variation and preserved meaningful within-world comparisons.

2.5 Responder-Assignment Strategies

2.5.1 Sequential Greedy Assignment

Under the Greedy strategy, incidents were processed sequentially. For each incident, the nearest currently unused responder that retained a feasible path to that incident was selected. Once assigned, the responder was no longer available for subsequent assignments within that allocation step. The method therefore makes decisions locally and sequentially; an assignment that is attractive for the current incident may restrict the choices available to later incidents.

2.5.2 Global Minimum-Cost Assignment

The Global Minimum-Cost approach considers the assignment problem jointly. Its first objective is to maximize the number of reachable incident-responder assignments. Conditional on achieving this maximum reachable cardinality, the method selects a one-to-one assignment that minimizes total travel cost using a min-cost-flow formulation. The method does not directly optimize fairness, worst-case response time, dynamic redeployment, or heterogeneous responder capacity.

2.6 Performance Metrics

System performance was evaluated using reachability, mean response time, median response time, 90th-percentile response time, total response time, and weighted unmet demand. Reachability represents the proportion of incidents successfully served under the available road network and responder configuration. Response-time statistics were calculated only for successful assignments. Unreachable incidents were not assigned arbitrary large travel times; their failure to receive service was represented through accessibility and unmet-demand measures.

2.7 The 80/80 Reliability Boundary

For each disruption severity, a simulated world was classified as a service success when reachability was at least 80%. Across repeated stochastic worlds, the probability of satisfying this service condition was then estimated. Wilson intervals were used for the estimated service-success probability [7].

Because service reliability should not increase as network disruption becomes more severe under the coupled design, the estimated probability curve was constrained to be non-increasing using a weighted pooled adjacent violators algorithm (PAVA) [8]. The 80/80 boundary was defined as the disruption-severity value at which the probability of maintaining at least 80% incident reachability crossed 80%. When the crossing occurred between adjacent tested severity values, its location was estimated using linear interpolation between the corresponding grid points.

2.8 Uncertainty Estimation

Uncertainty in the estimated reliability boundary was assessed using 5,000 complete-world cluster bootstrap resamples. The resampling unit was a complete simulated world rather than individual observations extracted independently from that world. This preserved the dependence structure created by coupled severity levels and resource configurations and follows the general bootstrap principle of resampling the empirical data-generating units [9].

2.9 Paired Statistical Comparison of Assignment Algorithms

Because Greedy and Global Minimum-Cost were evaluated on common stochastic worlds, algorithmic comparisons were conducted using paired differences. Uncertainty in the paired differences was estimated using 3,000 bootstrap resamples. Two-sided inference was additionally performed using 5,000 sign-flip permutations. Multiple comparisons were adjusted using the Holm procedure within each metric family [10].

2.10 Experimental Grids

The primary confirmatory experiment evaluated six responder counts (12, 14, 16, 20, 24, and 32) at 11 severity values from 0.130 to 0.165, with 300 stochastic worlds per severity-by-responder cell, 10 incidents per world, and base seed 7070. With two assignment algorithms, the primary run produced 39,600 algorithm-level outcome rows. An upper-bound extension evaluated responder counts 20, 24, and 32 over 12 severity values from 0.145 to 0.180 using the same seed, world count, incident count, coupling logic, and statistical pipeline, producing 21,600 algorithm-level rows.

Table 1. Confirmatory design and inferential parameters.

Element	Primary run	Upper-bound extension
Pilot area	Beykoz, Istanbul, Türkiye	Beykoz, Istanbul, Türkiye
Responder counts	12, 14, 16, 20, 24, 32	20, 24, 32
Severity grid	0.130-0.165 (11 levels)	0.145-0.180 (12 levels)
Worlds per grid cell	300	300
Incidents per world	10	10
Base seed	7070	7070
Boundary bootstrap	5,000 world-cluster resamples	5,000 world-cluster resamples
Paired bootstrap	3,000 resamples	3,000 resamples

Element	Primary run	Upper-bound extension
Permutation samples	5,000	5,000
Reliability criterion	P(reachability >=80%) = 80% crossing	Same
GitHub Actions run	31582171618	31585895864

2.11 Reproducibility and Research Provenance

The numerical evidence was frozen at scientific commit 8078af69eff047d7da71beac3b5b3b125a121d1a. The primary workflow run was 31582171618 with artifact 9136667673; the extension run was 31585895864 with artifact 9138183927. The archived ZIP SHA-256 digests were 0517161be1fb49657202df79b9ec7b7b9c0d41ee6bfbc553d6692cf52f195659 and 7a0c9b9e87576c12a19c233e226d629c1bb096baaa87098d081c29984e418caf, respectively, and internal checksum manifests were verified. The exact software release is archived at Zenodo [11], while the completed methodology and provenance are preserved in an accepted retrospective OSF Open-Ended Registration [12].

3. Results

3.1 Reliability Boundary and Responder Availability

The first pattern in the simulation results was the behavior of the 80/80 reliability boundary. As responder availability increased from 12 to 32, the boundary point estimates shifted toward higher levels of synthetic disruption. Within the tested model, larger responder pools therefore allowed the system to satisfy the 80/80 service criterion under greater disruption.

This pattern should not be interpreted as evidence that every incremental increase in responder count produced an independently significant improvement. The bootstrap intervals for several neighboring resource levels overlapped. The defensible conclusion is narrower: the boundary point estimates showed an overall upward trend across the tested responder levels, while the individual stepwise gains were not separately demonstrated to be statistically significant.

Table 2. Frozen Global Minimum-Cost 80/80 boundary estimates.

Responders	Boundary	95% CI low	95% CI high	Frozen source
12	0.142500	0.134000	0.148462	Primary
14	0.149167	0.142500	0.156000	Primary
16	0.152500	0.145333	0.160000	Primary
20	0.154500	0.147000	0.163214	Extension
24	0.158333	0.148462	0.168214	Extension
32	0.162500	0.154167	0.172727	Extension

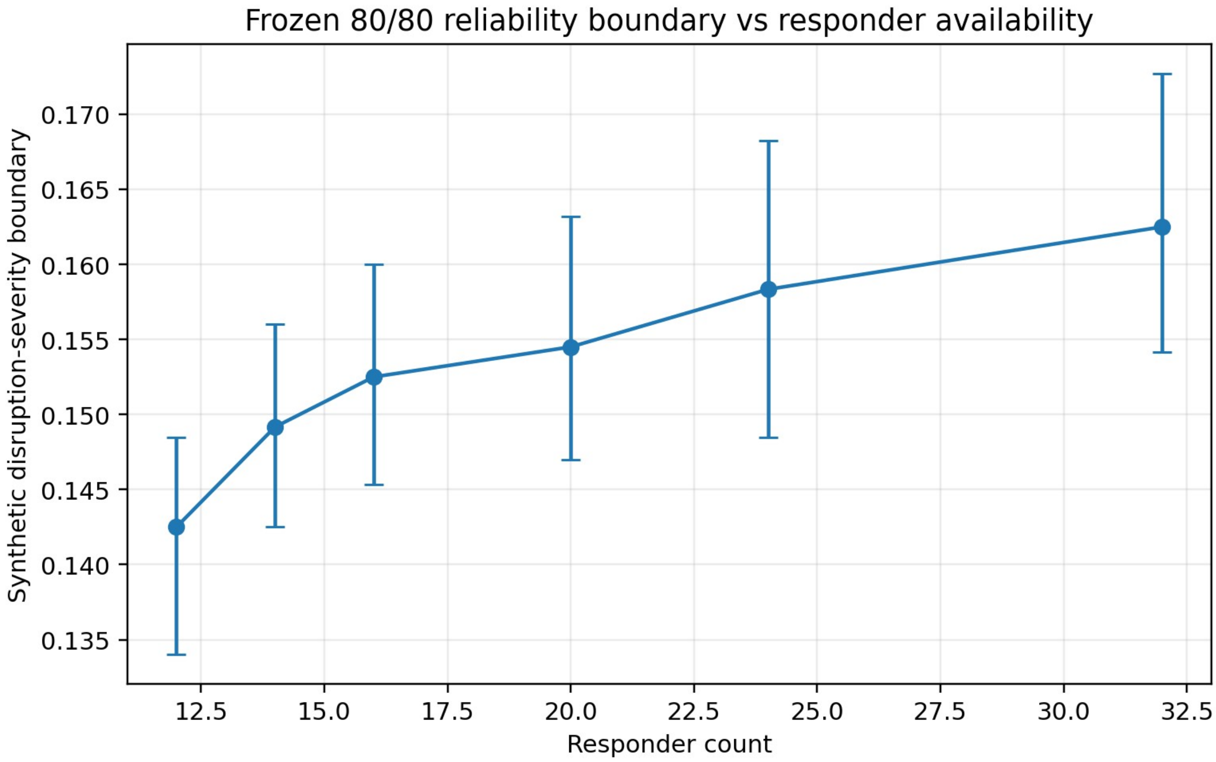

Figure 1. Frozen Global Minimum-Cost 80/80 boundary estimates as a function of responder availability. Error bars are 95% world-cluster bootstrap intervals. Severity is a synthetic disruption-control parameter, not a physical earthquake intensity.

3.2 Why the Upper-Bound Extension Was Necessary

The primary severity grid revealed a limitation in the original experimental range. For 20, 24, and 32 responders, the fractions of bootstrap samples producing finite boundary estimates were 0.9870, 0.8498, and 0.5154, respectively. The 32-responder result indicated that a substantial part of the boundary uncertainty distribution extended beyond the upper edge of the primary grid. After extending the severity range, the corresponding finite-bootstrap fractions increased to 0.9962, 0.9996, and 0.9998. The extension was therefore necessary to characterize the reliability boundary adequately at higher resource levels and to reduce upper-grid censoring.

3.3 The 32-Responder Transition

For 32 responders, the estimated probability of retaining at least 80% incident reachability was 81.33% at synthetic severity 0.160, exactly 80.00% at 0.1625, and 79.67% at 0.165. These neighboring points provide a direct view of the transition around the 80/80 criterion.

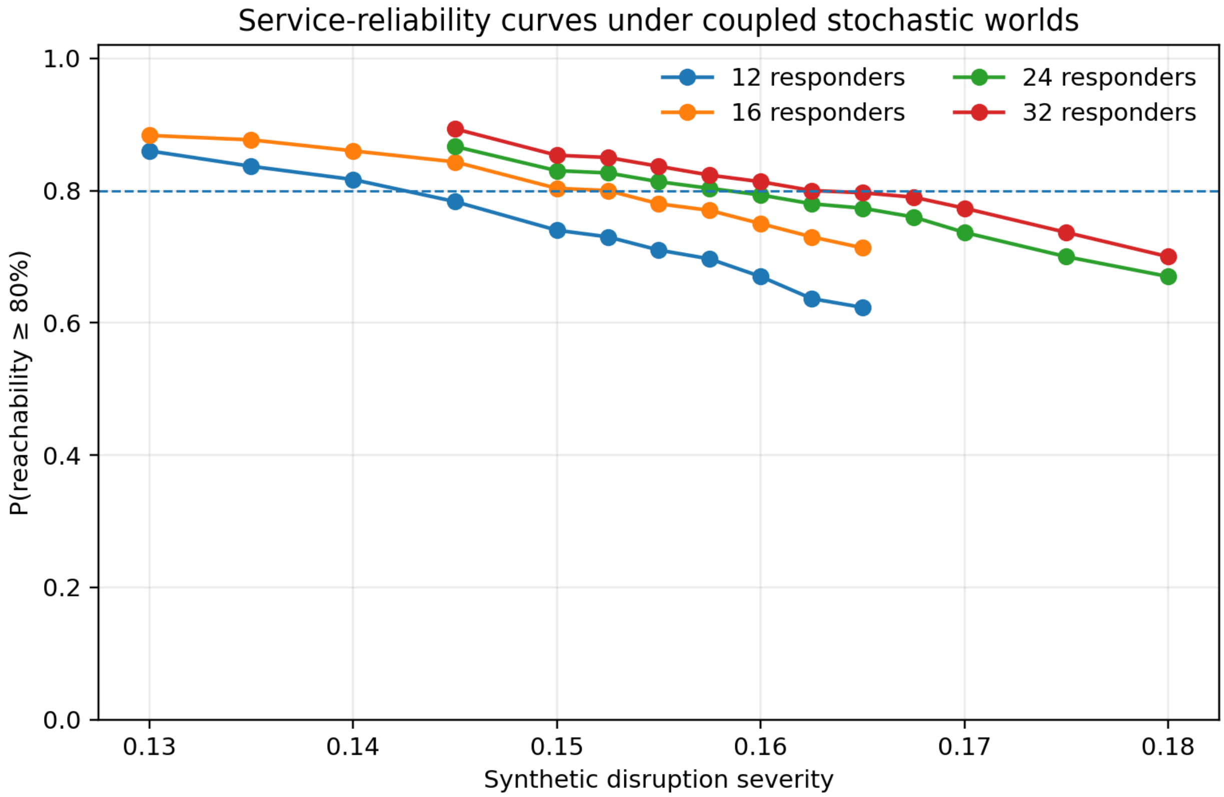

Figure 2. Representative Global service-reliability curves. Each point is the proportion of coupled stochastic worlds retaining at least 80% incident reachability. The dashed line marks the 80% reliability target. Curves for 24 and 32 responders use the upper-bound extension.

3.4 Resource Frontier

Viewed from the resource side, the tested Global 80/80 frontier required 24 responders at severity 0.155 and 32 responders at 0.160 and 0.1625. At severity 0.165, the criterion was not achieved by any tested resource level up to 32 responders. This grid-based frontier illustrates that additional resources can compensate for disruption while sufficient network connectivity remains, but the benefit is not unlimited.

3.5 Global Minimum-Cost versus Greedy Assignment

The clearest algorithmic difference appeared in total response time. In the primary confirmatory grid, Global Minimum-Cost produced lower total response time than Greedy in all 66 severity-by-responder comparisons. All 66 remained significant after Holm correction, with the largest adjusted p-value equal to 0.013197. The same pattern persisted in the upper-bound extension: all 36 tested cells favored Global, all remained significant after Holm correction, and the largest adjusted p-value was 0.007199.

Taken together, the two experiments produced the same directional total-response result in 102 of 102 confirmatory cells. This should be interpreted specifically as an allocation-efficiency finding under the tested model.

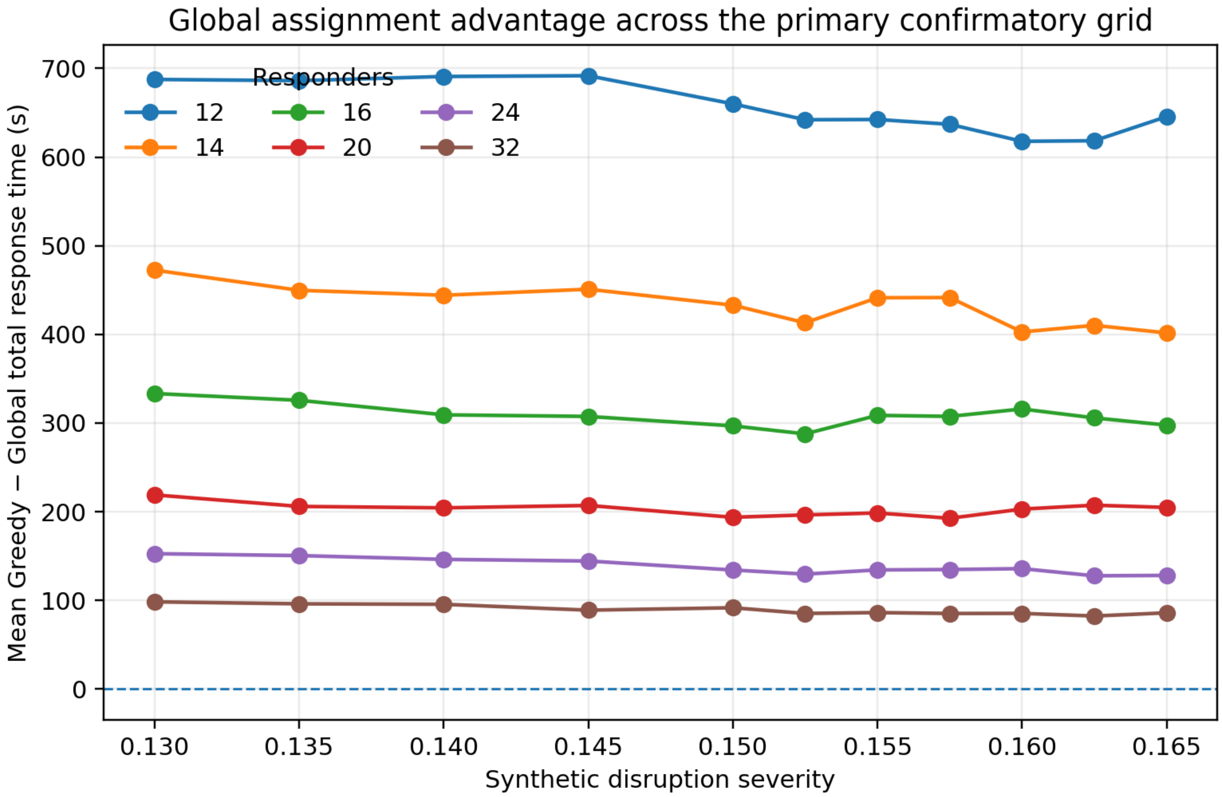

Figure 3. Mean paired total-response difference (Greedy minus Global) across the primary confirmatory grid. Positive values indicate lower total response time under Global Minimum-Cost Assignment. All 66 primary cells remained significant after Holm correction.

Table 3. Confirmatory algorithmic inference summary.

Evidence set	Metric	Cells	Holm-significant cells	Interpretation
Primary	Total response time	66	66	Global lower in every tested cell
Extension	Total response time	36	36	Global lower in every tested cell
Primary	Weighted unmet demand	66	0	No Holm-significant difference
Extension	Weighted unmet demand	36	0	No Holm-significant difference

3.6 Weighted Unmet Demand

The total-response advantage did not extend uniformly to all outcomes. For weighted unmet demand, none of the 66 primary comparisons and none of the 36 extension comparisons were Holm-significant. This absence of significance should not be interpreted as proof that the algorithms are equivalent on this metric. Rather, the experiments provided no evidence of a statistically detectable systematic difference in weighted unmet demand under the tested conditions.

3.7 Monotonicity Audit

Under Global Minimum-Cost Assignment, no increase-with-severity reachability violations were observed across 18,000 audited within-world transitions in the primary experiment. Greedy exhibited 12 such violations across the same 18,000 transitions. Twelve violations constitute a very small proportion and were not themselves subjected to a separate significance test. The result is therefore best treated as a diagnostic observation rather than evidence of a general stability difference.

3.8 Overall Pattern of Results

Viewed together, the experiments reveal two mechanisms. Increasing responder availability shifted the model-derived reliability boundary toward higher disruption levels, while Global Minimum-Cost consistently reduced total response time relative to Greedy. Neither mechanism eliminates the structural constraint imposed by the road network itself: resource availability determines how much response capacity enters the system, assignment optimization determines how efficiently reachable capacity is used, and network connectivity ultimately determines whether that capacity can reach an incident at all.

4. Discussion and Limitations

4.1 Resource Availability and Network Connectivity Are Different Constraints

The results distinguish two fundamentally different limitations in disrupted emergency-response systems: insufficient response resources and loss of network connectivity. These mechanisms should not be treated as interchangeable.

Increasing responder availability shifted the model-derived 80/80 reliability boundary toward higher levels of synthetic disruption. Within the tested system, larger responder pools therefore allowed the response network to maintain the specified service criterion under more severe disruption. However, the benefit of additional responders was not unlimited. Once disruption removed sufficient connectivity from the road graph, increasing the number of available responders could no longer compensate for the absence of feasible paths. A responder is useful only if the network permits that responder to reach an incident.

4.2 What Global Optimization Improved - and What It Did Not

Global Minimum-Cost consistently reduced total response time relative to the sequential Greedy method across the tested confirmatory conditions. This indicates that, when feasible incident-responder connections remain available, considering the allocation problem globally can make more efficient use of those connections than making assignments sequentially.

At the same time, the Global strategy did not shift the frozen 80/80 boundary point estimates relative to Greedy. The algorithm improved the cost of allocating resources across surviving feasible paths, but it did not restore connectivity that had already been lost from the network. The results therefore should not be summarized as evidence that Global Minimum-Cost generally increases resilience. The supported conclusion is metric-specific: within the tested computational model, Global Minimum-Cost reduced total response time relative to Greedy while producing the same frozen 80/80 accessibility-boundary point estimates.

4.3 From an Allocation-Limited to a Connectivity-Limited Regime

At lower levels of disruption, performance can be constrained by the amount of available response capacity. In this regime, additional responders and better allocation can improve system performance because sufficient road connectivity still exists for those resources to be deployed effectively. As disruption increases, the nature of the limitation changes: the system progressively enters a connectivity-limited regime in which the road graph itself restricts what any assignment policy can accomplish. Emergency-resource optimization and transportation-network accessibility should therefore be evaluated separately before they are combined into broader conclusions about system performance.

4.4 Generalizability

The study was conducted as a district-scale computational pilot in Beykoz, Istanbul. The results should not be interpreted as validated operational estimates for Istanbul as a whole or for Türkiye. Road topology, bridge dependence, street density, urban form, responder locations, demand distributions, and emergency infrastructure can vary substantially between districts and cities. External validation is necessary before the findings can support operational conclusions at a larger geographic scale.

4.5 Limitations

4.5.1 Synthetic Disruption Model

The disruption process is synthetic. The severity parameter controls edge-failure probability within the experimental model but is not calibrated to earthquake magnitude, peak ground acceleration, roadway fragility, bridge damage, debris blockage, or another physical measure of earthquake impact. The numerical boundary values therefore cannot be translated directly into real earthquake-intensity thresholds.

4.5.2 Scope of Bootstrap Uncertainty

The reported bootstrap intervals quantify uncertainty associated with sampling stochastic worlds under the specified computational model. They do not quantify structural uncertainty concerning the disruption law, the underlying road graph, demand-generation assumptions, responder deployment, or the validity of the model itself. They should therefore be interpreted as within-model stochastic uncertainty, not as a complete representation of real-world uncertainty.

4.5.3 OpenStreetMap Input Reproducibility

The road graph was retrieved from OpenStreetMap at execution time. Although the numerical results, workflow artifacts, source-code state, and integrity checks were frozen, the exact graph input was not preserved as a dedicated GraphML snapshot in the v1.0 research package. Because OpenStreetMap is continuously updated, a future researcher running the same acquisition procedure may obtain a graph that differs from the one used in the confirmatory experiments. This does not invalidate the frozen numerical evidence, but it limits exact future reconstruction of the original input state.

4.5.4 Simplified Assignment Model

Assignments are static and one-to-one. The model does not incorporate dynamic redeployment, queues, service times, real-time congestion, heterogeneous responder capacities, multiple sequential missions per unit, hospital capacity, clinical triage, priority-dependent objectives, or human dispatcher behavior. The Global Minimum-Cost objective first maximizes reachable assignment cardinality and then minimizes total travel cost; it does not directly optimize fairness, worst-case response time, tail-risk measures, or priority-weighted clinical outcomes. It should therefore be interpreted as an experimental optimization benchmark, not as a deployment-ready emergency-dispatch system.

4.6 Relation to Prior Work

The findings are compatible with prior work that treats post-earthquake accessibility, routing, infrastructure disruption, and emergency-resource allocation as coupled but analytically distinguishable problems [1-5,14-19]. They also complement recent digital-twin work in Türkiye that emphasizes integrated situational awareness and resource coordination for seismic disaster management [13]. The present study differs by deliberately using a transparent synthetic stress layer so that statistical coupling, resource-response behavior, and algorithmic attribution can be examined without assigning a physical earthquake meaning to the severity control. This narrower design makes the model less operationally realistic at the hazard layer but more explicit about which mechanism is responsible for each reported result.

4.7 Future Validation

A path toward operational interpretation would require replacing or augmenting the synthetic disruption layer with physically calibrated inputs. Recent studies demonstrate relevant directions, including satellite/GIS-based road-blockage assessment after the 2023 Türkiye earthquakes [14], probabilistic debris- and blockage-based road-network models [15,16], dynamic emergency-medical traffic and resource modeling [17,18], and integrated seismic vulnerability assessment that links infrastructure damage with accessibility [19]. Future work could therefore incorporate ground-motion fields, road and bridge fragility functions, correlated infrastructure failures, debris-induced blockage, and empirically validated responder and incident locations. The precise road-network snapshot should also be archived alongside the computational results. A more operational model could additionally represent congestion, service duration, responder redeployment, heterogeneous vehicle types, hospital capacity, incident priority, and dynamic dispatch decisions. The same 80/80 reliability-boundary framework could then be tested on top of a physically calibrated hazard and response model.

4.8 Overall Interpretation

The experiments help separate two classes of problems. When the principal limitation is resource scarcity, increasing responder availability and improving assignment can improve performance. When the principal limitation becomes loss of network connectivity, the range of actions available to an allocation algorithm becomes fundamentally constrained. Optimization can improve the use of connectivity that survives disruption; it cannot substitute for connectivity that no longer exists.

5. Conclusions

The central lesson of this study is that two constraints in disrupted emergency-response networks must be distinguished: how effectively available responders are allocated and whether the underlying network still provides feasible connections through which those responders can reach incidents. Treating these as a single problem can obscure an important difference between resource-allocation efficiency and network accessibility.

The reliability-boundary framework provides a quantitative way to examine that distinction. Rather than describing a disrupted network simply as performing well or poorly, the 80/80 boundary identifies a model-derived transition in service reliability under progressively increasing synthetic disruption.

Across the tested responder levels, increasing resource availability shifted this boundary upward. The frozen Global Minimum-Cost boundary increased from 0.1425 with 12 responders to 0.1625 with 32 responders. Within the computational model, this indicates that larger responder pools can maintain the specified accessibility criterion under greater levels of synthetic disruption. The result does not imply that additional resources can compensate indefinitely for deterioration in network connectivity.

Across all 102 confirmatory severity-by-responder cells, Global Minimum-Cost produced lower total response time than Greedy, with all comparisons remaining significant after Holm correction. Yet the frozen 80/80 boundary point estimates were identical between Greedy and Global Minimum-Cost at all six tested responder levels. This is the central conceptual result: improved allocation can make more efficient use of connectivity that survives disruption, but it cannot replace connectivity that the network no longer provides.

An additional contribution is auditability. The principal numerical findings are connected to identifiable computational runs, archived artifacts, a frozen source-code state, checksum verification, and persistent open-science records. The computational provenance of the findings can therefore be inspected independently rather than accepted solely from the narrative of the manuscript.

These conclusions must remain within the boundaries of the model. The disruption-severity parameter is a synthetic stress-control variable. It is not earthquake magnitude, peak ground acceleration, an engineering damage measure, or a prediction of physical road failure. The study does not predict how much of Beykoz or Istanbul would become inaccessible during a future earthquake, nor does it provide a deployment-ready emergency-dispatch protocol.

Ultimately, the study leads to a simple distinction: as long as feasible paths remain, better allocation can improve how effectively response resources use them. Once the network no longer provides those paths, however, optimization cannot substitute for physical connectivity.

Declarations and Open-Science Statements

Data and code availability. Source code, frozen research artifacts, reproducibility documentation, and the project interface are publicly available through the Türkiye Disaster Intelligence Digital Twin repository. The exact v1.0.0 software release is archived at Zenodo under DOI 10.5281/zenodo.21903851; the concept DOI is 10.5281/zenodo.21903850.

Registration. The completed study design and provenance are preserved in an accepted retrospective OSF Open-Ended Registration under DOI 10.17605/OSF.IO/ZTJXK. The registration was created after the principal software development, simulations, and confirmatory analyses and is not a preregistration.

Author contributions. Faramarz Kowsari: conceptualization, software, methodology, investigation, data curation, formal analysis, visualization, validation, writing - original draft, writing - review and editing, and project administration.

Ethics. The study is a computational simulation using public road-network data and synthetic incidents/responders; it does not involve human participants, identifiable personal data, or animal subjects.

Language and editing assistance. English translation and language copy-editing were assisted by a generative AI system using author-written Persian drafts. The author reviewed and verified the scientific content, numerical results, citations, interpretations, and final manuscript wording.

Appendix A. Evidence Provenance for Principal Manuscript Claims

Table A1. Principal claim-to-evidence map.

Claim	Evidence file	Run / artifact	Code locus	Interpretive boundary
80/80 boundary at six responder levels	boundary_estimates.csv; frozen final tables	31582171618 / 9136667673; 31585895864 / 9138183927	boundary_estimates()	Synthetic severity only
Global lower total response time in 66/66 + 36/36 cells	paired_grid_inference.csv	same two runs	paired_grid_inference(); holm_adjust()	Metric-specific; not all outcomes
No Holm-significant weighted unmet-demand differences	paired_grid_inference.csv	same two runs	paired_grid_inference()	Non-significance is not proof of equivalence
Identical algorithmic boundary point estimates	final_boundary_all_algorithms.csv	frozen v0.8 synthesis	boundary_estimates()	Point estimates identical; bootstrap intervals need not be
Global monotonicity audit	monotonicity_audit.csv	31582171618 / 9136667673	monotonicity_audit()	Audit of coupled computational worlds
Frozen provenance and checksums	FROZEN_MANIFEST.json ; SHA256SUMS	same two runs	artifact pipeline	Bit-level integrity, not external validity

References

[1] Bakioğlu G, Karaman H. Accessibility of medical services following an earthquake: A case study of traffic and economic aspects affecting the Istanbul roadway. International Journal of Disaster Risk Reduction. 2018;31:403-418. doi:10.1016/j.ijdrr.2018.05.018.

[2] Sakuraba CS, Santos AC, Prins C, Bouillot L, Durand A, Allenbach B. Road network emergency accessibility planning after a major earthquake. EURO Journal on Computational Optimization. 2016;4(3-4):381-402. doi:10.1007/s13675-016-0070-2.

[3] Gehl P, Auclair S, Fayjaloun R, Meresse P. Decision support for emergency road traffic management in post-earthquake conditions. International Journal of Disaster Risk Reduction. 2022;77:103098. doi:10.1016/j.ijdrr.2022.103098.

[4] Edrissi A, Poorzahedy H, Nassiri H, Nourinejad M. A multi-agent optimization formulation of earthquake disaster prevention and management. European Journal of Operational Research. 2013;229(1):261-275. doi:10.1016/j.ejor.2013.03.008.

[5] Boutilier JJ, Chan TCY. Ambulance Emergency Response Optimization in Developing Countries. Operations Research. 2020;68(5):1315-1334. doi:10.1287/opre.2019.1969.

[6] Boeing G. OSMnx: New methods for acquiring, constructing, analyzing, and visualizing complex street networks. Computers, Environment and Urban Systems. 2017;65:126-139. doi:10.1016/j.compenvurbsys.2017.05.004.

[7] Wilson EB. Probable inference, the law of succession, and statistical inference. Journal of the American Statistical Association. 1927;22(158):209-212. doi:10.1080/01621459.1927.10502953.

[8] Barlow RE, Brunk HD. The isotonic regression problem and its dual. Journal of the American Statistical Association. 1972;67(337):140-147. doi:10.1080/01621459.1972.10481216.

[9] Efron B. Bootstrap methods: Another look at the jackknife. The Annals of Statistics. 1979;7(1):1-26. doi:10.1214/aos/1176344552.

[10] Holm S. A simple sequentially rejective multiple test procedure. Scandinavian Journal of Statistics. 1979;6(2):65-70. doi:10.2307/4615733.

[11] Kowsari F. Türkiye Disaster Intelligence Digital Twin v1.0.0 [software]. Zenodo. 2026. doi:10.5281/zenodo.21903851.

[12] Kowsari F. Türkiye Disaster Intelligence Digital Twin: Open-Ended Registration of a Reproducible Computational Study of Post-Earthquake Emergency-Response Reliability. OSF Registries. 2026. doi:10.17605/OSF.IO/ZTJXK.

[13] Korkmaz M, Zulfikar AC, Demirkesen S. Leveraging Digital Twins as a Common Operating Picture for Disaster Management: Case of Seismic Hazards. ISPRS International Journal of Geo-Information. 2024;13(12):430. doi:10.3390/ijgi13120430.

[14] Awan HH, Kabdrakhmanova M, Varol HA, Karaca F. A GIS-Based emergency response and management support framework for Earthquake crisis: A case study of Antakya and Kahramanmaras earthquake in Turkey. International Journal of Disaster Risk Reduction. 2024;111:104706. doi:10.1016/j.ijdrr.2024.104706.

[15] Hosseini Y, Karami Mohammadi R, Yang TY. A comprehensive approach in post-earthquake blockage prediction of urban road network and emergency resilience optimization. Reliability Engineering & System Safety. 2024;244:109887. doi:10.1016/j.ress.2023.109887.

[16] Hosseini Y, Karami Mohammadi R, Yang TY. Resource-based seismic resilience optimization of the blocked urban road network in emergency response phase considering uncertainties. International Journal of Disaster Risk Reduction. 2023;85:103496. doi:10.1016/j.ijdrr.2022.103496.

[17] Wu Y, Chen S. Traffic resilience modeling for post-earthquake emergency medical response and planning considering disrupted infrastructure and dislocated residents. International Journal of Disaster Risk Reduction. 2023;93:103754. doi:10.1016/j.ijdrr.2023.103754.

[18] Wu Y, Chen S. Resilience modeling and pre-hazard mitigation planning of transportation network to support post-earthquake emergency medical response. Reliability Engineering & System Safety. 2023;230:108918. doi:10.1016/j.ress.2022.108918.

[19] El-Maissi AM, Argyroudis SA, Kassem MM, Nazri FM. Integrated seismic vulnerability assessment of road network in complex built environment toward more resilient cities. Sustainable Cities and Society. 2023;89:104363. doi:10.1016/j.scs.2022.104363.

[20] Boeing G. Modeling and Analyzing Urban Networks and Amenities With OSMnx. Geographical Analysis. 2025;57(4):567-577. doi:10.1111/gean.70009.