

1. This is not a quantitative catastrophe-dynamics model. No transition rate is estimated. No probability is calculated. No model is calibrated. The analysis uses only a binary edge list.

Author response:

We thank the reviewer for careful reading and candid criticism. We concur with the reviewer's judgement: we do not estimate transition rates, compute or calibrate probabilities in the original manuscript. All empirical metrics are derived from the binary edge list. We recognise that labelling our analysis a "quantitative catastrophe-dynamics model" was an overclaim. The revised title will read: "Network Modelling of Marine Seismic Hazard Chains: Topology, Core Amplification Nodes, and Implications". In the abstract and introduction, "quantitative model" will be changed to "directed-network model" or "topological analysis". Remaining mentions of "quantitative" will be adjusted as needed.

2. Equations 1 – 13 are disconnected from the results. The authors introduce τ , ρ , and μ . None of them is estimated from observations. These equations add complexity but no real analysis.

Author response:

We thank the reviewer for the comments and agree. All results from the original manuscript are calculated directly by graph theory on the binary directed graph built from documented couplings. Degree, density and centralisation stem from edge counts; average path length is calculated using shortest-path algorithms; cascade reachability is determined by searching for reachable nodes on the directed graph. We perform no numerical estimation or calculation for generator Q , transition matrix A , or cumulative matrix M . The revisions are as follows:

(1) Remove derivations unrelated to our results.

(2) Delete the propagation horizon $\tau = 10$, because binary analysis requires no truncation steps.

(3) Keep and revise the definitions of edges and reachability to clarify their links with our calculated outputs. Edge $i \rightarrow j$ denotes that hazard i can directly trigger hazard j according to literature records.

3. The definition of A is unclear. It is first defined using ρ . It is later treated as an empirical adjacency matrix. These are not automatically the same thing.

Author response:

We thank the reviewer for identifying this confusion. We will adopt separate symbols in the revised manuscript to prevent ambiguity.

4. The propagation horizon of ten steps has no clear role. The authors later use an infinite series. The final analysis uses only reachability. The choice of ten steps appears meaningless.

Author response:

We thank the reviewer for the comments. All statements regarding $\tau = 10$ will be removed in the revised manuscript.

5. The dataset construction is poorly explained. The paper does not provide a clear event-selection procedure. It does not explain how the 366 chains were identified.

Author response:

We thank the reviewer for the comments. The original manuscript provides inadequate documentation for dataset construction. We will supplement event sources and selection steps in the revision.

6. The edge-coding process is subjective. The authors mention expert encoding. They provide no information about the experts, coding rules, agreement, or conflict resolution.

Author response:

We thank the reviewer for the comments. The original manuscript only mentioned “expert encoding” without describing coders, coding rules, consistency tests or disagreement management. This prevents readers from evaluating the reliability of the network edge set.

We add the following details in revision:

- (1) The second author carried out duplicate coding and verification.
- (2) Coding rules.
- (3) Coding consistency.
- (4) Disagreement resolution.

7. The network may reflect reporting bias. Recent and famous disasters have better documentation. Older events have fewer available records. This can strongly affect the number of reported links.

Author response:

We thank the reviewer for raising this issue. We acknowledge that recent, prominent disasters have richer documentation, which inevitably shapes the number of couplings captured in our network.

Incomplete records omit genuine couplings and do not create spurious ones. Each network edge is tied to at least one literature-documented real event. Whether a coupling is reported once or twenty times in one event, it counts as a single edge. Well-documented events can introduce new couplings but cannot strengthen existing ones; recent records mainly add granular details. Reporting bias thus yields downward-biased edge counts and node degrees. These absolute quantities should be treated as lower bounds. Still, the core nodes and inter-layer propagation directions supporting our conclusions are already present in poorly documented historical events. We will add this discussion to the revised manuscript.

8. Every edge receives the same weight. A link reported once is treated like a link reported many times. This discards important information about frequency and reliability.

Author response:

We thank the reviewer for the comments. We agree that our network treats couplings documented only once the same as those repeatedly reported, which discards frequency information. We adopted a binary representation in the original manuscript for two reasons. First, frequencies largely reflect the completeness of documentation rather than merely the propensity for coupling occurrence (see Response 7). Second, we intended the network structure to be independent of any specific weight definition.

9. The taxonomy is highly subjective. It combines physical hazards, infrastructure failures, social outcomes, and broad consequence categories in one network.

Author response:

We thank the reviewer for the comments. We acknowledge that taxonomy construction relies on subjective judgements, which was poorly explained in the original manuscript. Grouping different node types in a single network is necessary to address our research question and is supported by prior literature. Our core research question examines how hazards transfer from natural systems to engineered systems and further turn into socio-economic losses, i.e. the locations of Natech cross-system couplings. Combining natural hazards, infrastructure failures and socio-economic outcomes within one interaction network is widely adopted in cascade hazard studies.

(1) Mignan and Wang (2020) constructed a network with 19 categories of “loss-generating events”, covering natural hazards, network failure and business interruption. They demonstrated that interactions are mainly amplified via these two sets of nodes.

(2) Pescaroli and Alexander (2016) emphasise the core role of critical infrastructure along cascade propagation paths.

(3) Our classification system follows the hazard framework from Mignan (2024) and is refined for marine earthquake scenarios.

10. Critical infrastructure failure is an extremely broad node. Its high degree is therefore expected. The paper presents this result as a discovery. It is largely created by the node definition.

Author response:

We thank the reviewer for the comments. We acknowledge that critical infrastructure failure (CF) has a broad definition that increases its degree. Calling this high degree a “finding” in the original text was inappropriate and will be rephrased in revision. However, a broadly defined node does not guarantee high in-degree or out-degree:

(1) A wide definition does not inherently lead to high degree. If definition scope dominated degree, categories merging more sub-nodes would always have larger degree values. Degree is determined by the number of couplings recorded for that hazard class, not only how broadly the category is defined.

(2) Broad scope accounts for convergence, not relay transmission. Other network nodes also have wide definitions: human casualties include death, drowning, missing people and psychological trauma; business interruption covers industry, fisheries, tourism and cross-border trade; critical network failure comprises power, transport, communication, water and gas networks. These nodes have high in-degree but low out-degree. They aggregate many incoming links and rarely propagate cascades onward.

11. Several aggregation choices are difficult to defend. Liquefaction, uplift, subsidence, methane blowout, and snow avalanche are grouped as mass slides. These are different processes.

Author response:

We thank the reviewer for the comments. We acknowledge that liquefaction, ground uplift, subsidence, methane seepage, avalanches and landslides are separate physical processes. We clarify that the grouping in Appendix Table A1 is not based on physical mechanisms.

Relevant wording will be modified in the revision to avoid conceptual confusion.

12. Volcanic eruptions are grouped with earthquakes. Tsunamis are grouped with floods. Insufficient rescue is grouped with healthcare degradation. These choices appear arbitrary.

Author response:

We thank the reviewer for the comments. Consistent with our preceding reply, these aggregations do not follow physical mechanisms but are constrained by the onshore-offshore comparison. The onshore dataset was coded into 15 predefined categories, and the 39 marine nodes need to be mapped onto these existing classes. The original manuscript omitted the justification for each mapping, which makes the mapping look arbitrary. Details are listed below:

- (1) Volcanic eruptions fall into the “Earthquake / trigger” category. This class represents the “trigger layer” in the onshore classification, the initial perturbation starting cascades, and is not limited only to earthquakes. All volcanic eruptions in the marine network are triggered by submarine or near-coast earthquakes and subsequently generate tsunamis, positioned right next to the trigger source in cascades. We rename this category “Seismic trigger (earthquake, aftershock, and seismically triggered eruption)” in revision to improve clarity.
- (2) Tsunamis are grouped under “Flood”. The onshore “Flood” class includes hazards causing losses through water inundation. Tsunamis produce damage via coastal seawater inundation and form continuous inundation pathways with coastal and river floods in our marine network (tsunami → coastal flood).
- (3) Insufficient rescue capacity is assigned to “degraded health and medical capacity”. Both terms describe weak post-disaster emergency response: insufficient search-and-rescue resources in the former, and limited medical and public health capacity in the latter. Both sit in the consequence layer and represent failures of the social response system.

13. The mainland network is not described properly. Its event selection, data sources, and coding process are missing. A fair comparison is therefore impossible.

Author response:

We thank the reviewer for the comments. The original manuscript contains inadequate descriptions of the onshore network, leaving readers unable to judge whether the comparison is fair. Relevant content will be supplemented in the revised manuscript.

14. The phrase “energy waterfall” is incorrect. The study counts edges. It does not calculate energy, intensity, probability, or loss.

Author response:

We thank the reviewer for the correction. Our inter-layer net flow quantifies the count of directed cross-layer couplings: the difference between outgoing and incoming edges per layer. It involves no energy, intensity, probability or loss calculations. “Energy waterfall” and comparable phrasing can mislead readers to assume we compute energy or loss propagation along cascades. This metaphor is inaccurate and all related mentions will be deleted in the revised manuscript.

15. Degree is repeatedly interpreted as amplification. Degree only measures connectivity. It

does not show that a node increases hazard intensity or damage.

Author response:

We thank the reviewer for the comments and agree. Degree merely quantifies connection count and cannot prove that a node raises hazard intensity or worsens losses. The original manuscript did not define “amplification” precisely, so readers may misinterpret it as physical amplification. Here, “amplification” describes structural cascade expansion in the network: via one node, a single trigger can reach more downstream hazards and outcomes. It does not mean growth in hazard intensity or losses. This clarification will be added in the revision.

16. The intervention claims are unsupported. Removing a node from a graph is not equivalent to preventing a real infrastructure failure.

Author response:

We thank the reviewer for the comments. Our node-removal analysis does not aim to quantify loss reduction from mitigation measures; instead, it identifies structurally critical network components. We selectively remove nodes and monitor changes in network connectivity to find components that dominate overall network structure. The output gives a structural priority ranking for nodes that best support cascade propagation paths. It does not assess intervention performance or investment priorities (see Response 35).

17. No intervention cost is considered. No failure probability is included. No avoided damage is calculated. The term “optimal intervention” is unjustified.

Author response:

We thank the reviewer for the correction and fully agree. The term “optimal intervention” requires balancing intervention cost, reduced failure probability and avoided losses under defined constraints to minimise expense or maximise benefit. Our study includes none of these terms, so statements involving “optimal” are unsupported. All mentions of “optimal” and similar expressions will be removed in the revision.

18. The node-removal analysis is biased. Removing a node also removes every pair involving that node. Part of the reported reduction is therefore automatic.

Author response:

We thank the reviewer for the comments. The effect noted by the reviewer is real: removing a node eliminates all hazard pairs where the node acts as source or target. We maintain that this removal is not just a computational artefact and has clear physical interpretation, and it does not alter our conclusions.

For critical infrastructure failure: inbound pairs “tsunami → critical-infrastructure-failure” and “liquefaction → critical-infrastructure-failure” describe facility damage caused by secondary hazards. Outbound pairs “critical-infrastructure-failure → business interruption” and “critical-infrastructure-failure → human casualties” represent losses cascading from failed infrastructure. If infrastructure failure is prevented, these couplings will not occur in practice, and infrastructure failure itself is a major component of disaster loss. Network vulnerability analysis routinely uses this convention when assessing node removal effects:

links attached to the removed node are included when calculating connectivity loss.

19. The so-called hindcast is not genuinely independent. The three events are famous historical disasters. Palu is even discussed in the Introduction.

Author response:

We thank the reviewer for the comments. We agree “hindcast” is unsuitable, and Section 3.5 will be renamed “out-of-sample coverage check” in revision. We also recognise these three events are well known and lack full knowledge independence. Still, well-documented prominent events provide a more stringent test. Choosing famous events does not lower the bar for passing the test; it makes the test stricter. The core question of this check does not depend on event notoriety. The Palu event cited in the introduction is used merely to exemplify how submarine landslides can dominate bay-scale tsunami generation.

20. The paper provides no evidence that the three validation events were selected before examining the network. The analysis appears post hoc.

Author response:

We thank the reviewer for the comments. The reviewer’s concern regarding post-hoc testing is valid. We confirm the three test events were not pre-registered. We selected and coded them after finishing network construction, so this test is post-hoc. This information will be added in the revision.

21. The reported coverage is only 74 percent. More than one-quarter of the links are missing. This is not strong validation.

Author response:

We thank the reviewer for the comments. We acknowledge that 74% coverage implies our network misses some couplings.

(1) Relative to the random benchmark, 74% greatly exceeds chance levels.

(2) This 74% aligns with expectations for our sample size; omitted couplings form the unsaturated long tail.

(3) Missing couplings cluster on the network periphery, and coverage of the core structure is much higher.

(4) The key of this test is the stability of conclusions, not full coverage.

22. The authors dismiss the missing links as a rare tail. This is not justified. Rare pathways can produce the most unexpected and serious consequences.

Author response:

We thank the reviewer for the comments and fully agree. “Rare” only means a coupling has few recorded occurrences and does not indicate low importance. Our degree-derived structural metrics primarily reflect the backbone built from frequently reported couplings. Rare couplings barely affect these metrics, yet we cannot conclude their risk impact is small. Still, our backbone analysis does not conflict with the rare high-consequence pathways highlighted by the reviewer. Data demonstrate most such rare paths branch directly from backbone hubs. Rare high-consequence paths are branches of the backbone instead of standalone chains separate from it. Detecting backbone hubs identifies the most probable

origins of these high-consequence branches. Our topological analysis thus supplies a starting point for scenario studies of rare pathways. Assessing rare severe outcomes triggered by critical infrastructure failure or tsunami is essentially examining outward extensions from backbone hubs.

Revisions:

1. Delete wording that minimises rare couplings. In Section 3.4, sentences like “only affect the long tail” and “have no effect on network topology” are replaced by the precise sentence: “...rarely documented couplings contribute little to these degree-based metrics but may carry severe consequences.”

2. Add to limitations: rare couplings have the poorest documentation, and our present network likely misses some high-consequence paths that have not yet happened or been documented.

23. The robustness analysis tests very broad conclusions. Earthquakes remain source nodes because this is an earthquake-triggered network. Consequences remain sinks because they were defined as consequences.

Author response:

We thank the reviewers for their comments. The conclusions in this paper that are genuinely not determined by definition but require empirical testing concern the intermediate layers and specific nodes: which nodes act as transit hubs, and at which layer natural hazards are converted into technological disasters.

Identification of transit hubs: Critical infrastructure failure exhibits both high in-degree and high out-degree, a feature that holds robustly across the vast majority of resampled samples. Tsunami ranks among the highest in out-degree among all non-triggering nodes; indeed, the two nodes with the highest out-degree among non-triggering nodes are precisely critical infrastructure failure and tsunami.

Location of Natech conversion: The net flow level of the engineering layer is lower than that of the natural process layer; the engineering layer consistently behaves as a net sink, while the natural process layer is close to balance, with net flow accounting for only a very small proportion of the total number of edges.

Sink nodes: Critical network failure nodes have in-degree greater than out-degree, a feature that holds across all resampled samples.

None of these conclusions can be derived from the way the network was constructed or from the layered definitions of the nodes: the classification scheme does not specify which engineering node would become a hub, nor does it specify whether the net flow of the natural process layer should be higher or lower than that of the engineering layer. The fact that these findings hold across the vast majority of resamples also indicates that the core results of this paper do not depend on a small number of events.

24. Bootstrapping cannot validate assumptions built into the taxonomy. It only resamples the same subjective coding decisions.

Author response:

We thank the reviewer for the comments and agree. The object of the Bootstrap resampling is the events themselves; it assesses how sensitive the conclusions are to which events are

included — that is, sampling variability. It cannot test whether the classification scheme and the coding decisions are themselves sound, because every resample is based on the same classification and the same set of codings. The original manuscript framed the Bootstrap results as a validation of the overall conclusions, which was too broad in scope. In the revised manuscript, we will present the tests of the classification scheme and the coding as a separate section.

25. The main uncertainty comes from human judgment. The bootstrap does not include coding disagreement, missing documents, or alternative causal interpretations.

Author response:

We thank the reviewer for the comments. Human judgment is the primary source of uncertainty in this paper; however, event-level Bootstrap retains the same coding and thus cannot capture coding disagreements, missing literature, or divergent causal interpretations. The revised manuscript will add simulations of coding uncertainty, covering coding, literature, and causal interpretation.

26. The high-frequency coverage analysis may contain data leakage. The paper does not explain whether the held-out events were used to identify high-frequency links.

Author response:

We thank the reviewer for pointing this out. The revised manuscript will add a chronological split test: using a given year as the cutoff, we construct the network from events occurring before that year only, identify high-frequency couplings within these early events, and then examine whether these couplings recur in subsequent events.

27. The mean path length excludes unreachable pairs. This can make the network appear artificially shallow.

Author response:

We thank the reviewer for the comments. We agree that average path length excludes unreachable disaster pairs. The reviewer's comment also prompted us to reconsider what "shallow" precisely means. What the network reveals is not a "shallow cascade" but rather that the same disaster often has multiple reachable paths, the shortest of which takes only one or two steps. The revised manuscript makes the following changes: we remove the phrase "shallow cascade" and replace it with "short shortest-path distances from the seismic triggers," clarifying that it means disasters can be triggered within a few steps at the fastest, rather than the depth of the cascade itself. We also report the actual depth distribution of recorded chains and the distribution of shortest-path distances in the network, so readers can distinguish the two.

28. A short topological path does not mean a short response time. The edges have no physical duration. The claim about a short intervention window is unsupported.

Author response:

We thank the reviewer for the correction and fully agree that the edges in our network only indicate the existence of a coupling and carry no temporal information. The short paths do have a legitimate structural meaning, but it is unrelated to time: it concerns the number of intervention points, not the duration of intervention. A shortest path of one or two steps

from the trigger to a hazard means that very few intermediate couplings exist along that route at which the cascade could be interrupted; if a hazard can be triggered directly by the earthquake, there is no intermediate coupling to block, leaving only the option of improving the resilience of the hazard object itself. The revised manuscript rewrites the relevant passage as follows:

Original (lines 265–266): Because the cascade remains shallow, the window for interrupting it is short.

Revised: Because most hazards can be reached from the seismic triggers within one or two steps, few intermediate couplings are available at which a cascade could be interrupted along these shortest routes. This is a statement about the number of intervention points, not about the time available; the network carries no information on the duration of each coupling.

29. The paper calls the network a single-source cascade. It also reports two pure sources. This is inconsistent.

Author response:

We thank the reviewer for pointing out this inconsistency. "Single-source" originally referred to a single category of trigger: submarine and near-shore earthquakes both belong to the same trigger category, earthquakes. In the 15-class aggregated network, the two are merged into the single "earthquake/trigger" class, and the network indeed has only one pure source at that level. The revised manuscript will adjust the wording to avoid the contradiction.

30. Causality is assumed. A documented sequence does not necessarily prove that one event caused another.

Author response:

We thank the reviewer for the comments. We agree that the chronological order of documented events alone does not prove that the former caused the latter within any single event. Our edges do not claim that hazard i caused hazard j in a specific event; rather, they claim that hazard i can trigger hazard j , i.e., that a possible triggering relationship exists between the two. Gill and Malamud (2014) systematically compiled triggering relationships among natural hazards, based on documented instances in the literature and known physical mechanisms, rather than on causal inference for each instance. We follow that framework and extend it to the engineering and societal consequence layers.

The revised manuscript will clarify the meaning of an edge: "an edge $i \rightarrow j$ indicates that hazard i has been documented to trigger hazard j at least once and that a physical or functional mechanism for this triggering exists; it does not imply that i caused j in every event, nor does it quantify the causal effect."

31. Many events can happen simultaneously. Infrastructure failure, network failure, transport disruption, and business interruption may not have a unique causal order.

Author response:

We thank the reviewer for the comments and agree. Infrastructure failure, network failure, transport disruption, and business disruption often occur nearly simultaneously and reinforce one another, without a unique temporal order. Our network is directed, but it does

not require a single direction between hazards: whenever the literature documents both directions in different events or accounts, both edges are retained. The data already contain several bidirectional couplings, e.g., critical infrastructure failure $\leftrightarrow$ critical network failure, critical infrastructure failure $\leftrightarrow$ fire, critical infrastructure failure $\leftrightarrow$ port failure, and explosion $\leftrightarrow$ fire, which act as each other's upstream and downstream, and thus feedback loops exist in the network. Moreover, our coding rules encode simultaneously occurring consequences as parallel branches under a common cause (e.g., $A > B + C$ is coded as $A \rightarrow B$ and $A \rightarrow C$), rather than forcing them into a $B \rightarrow C$ chain.

32. The paper repeatedly uses the words probability, amplification, prediction, and optimization. None of these quantities is actually estimated.

Author response:

We thank the reviewer for the comments. In the revised manuscript, we have conducted a word-by-word check of the full text; these terms are retained only in the following three cases: (a) when describing definitions within the theoretical framework; (b) when reporting existing literature; and (c) when referring to quantities that we actually computed and for which we have given explicit definitions. Any usage falling outside these three cases has been deleted or rewritten.

33. The claim that the network is fundamentally complete is too strong. The accumulation curve does not fully saturate. The external events also add 14 new links.

Author response:

We thank the reviewer for the correction. The phrases "fundamentally complete" and "exerts no influence on network topology" will be removed from the revised manuscript.

34. The claim that missing data have no influence is unsupported. The authors only show that four broad conclusions remain unchanged.

Author response:

We thank the reviewer for the comments and agree. The original manuscript did not provide sufficient support for the claim that the missing data had "no influence," and this statement will be removed from the revised manuscript.

35. The mitigation recommendations go far beyond the evidence. A binary network cannot determine the most effective real-world investment strategy.

Author response:

We thank the reviewer for the comments. Our network can only identify elements that are structurally central and thus worthy of prioritized inclusion in site-specific assessments; it cannot determine the most cost-effective investment plan in reality, which requires data on empirical probabilities, costs, and avoidable losses.

The revised manuscript rewrites the mitigation-related conclusions as follows: "Under a range of plausible triggering probabilities, local reinforcement of engineering nodes such as critical infrastructure, ports, and critical networks consistently ranks among the more effective means of suppressing cascade propagation in the network. This result merely

identifies robust candidate targets for site-specific assessment and cannot determine the optimal investment plan in practice; determining the optimal plan would additionally require data on empirical probabilities, costs, and avoidable losses."

In addition, the revised manuscript removes all statements about investment priorities and intervention effectiveness from the Discussion and Conclusions, retaining only the ranking of structural importance.

Mignan, A. (2024). Introduction to Catastrophe Risk Modelling: A Physics-based Approach. Cambridge University Press. <https://doi.org/10.1017/9781009437370>

Mignan, A., Wang, Z.: Exploring the Space of Possibilities in Cascading Disasters with Catastrophe Dynamics. International Journal of Environmental Research and Public Health, 17, <https://doi.org/10.3390/ijerph17197317>, 2020.

Pescaroli, G., Alexander, D.: Critical infrastructure, panarchies and the vulnerability paths of cascading disasters. Natural Hazards, 82, 175–192, <https://doi.org/10.1007/s11069-016-2186-3>, 2016.