Species-Level Interaction Heterogeneity Localizes Reactive Modes and Widens the Stable-but-Reactive Window in Random Ecological Communities

PyHelix
Axiom BOINC Project

March 7, 2026

Abstract

Robert May's random matrix framework predicts a sharp interaction-strength threshold beyond which large ecological communities lose asymptotic stability, but stability alone does not capture transient amplification. Reactivity extends the theory by identifying systems that are linearly stable yet capable of short-term perturbation growth. Here we test whether species-level interaction heterogeneity changes not only where that reactive regime begins, but also how transient amplification is distributed across species. Using May-style random community matrices with multiplicative row and column heterogeneity, we analyzed a full factorial sweep of species richness, connectance, and heterogeneity on the Axiom distributed BOINC volunteer network. The experiment comprised 735 independent random seeds executed across 17 hosts, with each seed evaluating repeated internal trials over 108 parameter combinations per trial. Increasing heterogeneity from $\sigma_h = 0$ to $\sigma_h = 1.35$ widened the stable-but-reactive window by 3.96477 ± 0.04442 (mean $\pm$ SE across seeds; Cohen's $d = 89.25$), increased the inverse participation ratio (IPR) of the dominant reactive eigenvector by 0.28195 ± 0.00084 ($d = 335.59$), and reduced the effective fraction of species supporting that mode by 0.22077 ± 0.00066 ($d = -332.85$). Every one of the 735 seeds showed the same directional effect. Across all simulated conditions, heterogeneity correlated positively with window width (0.60704 ± 0.00525) and even more strongly with reactive IPR (0.81776 ± 0.00137), again with perfect sign consistency across seeds. These results show that heterogeneity does not simply move ecosystems toward instability in aggregate. It creates a broad parameter region in which communities remain asymptotically stable but become transiently fragile, and that fragility is spatially concentrated onto

a small subset of species. The central theoretical contribution is therefore a mechanistic link between interaction heterogeneity, eigenvector localization, and ecological reactivity.

Keywords: complexity-stability, reactivity, transient amplification, random matrix theory, eigenvector localization, ecological networks, volunteer computing

1 Introduction

The complexity-stability problem in ecology has been shaped by May's demonstration that large random communities become unstable when typical interaction strength exceeds a threshold that scales as $(SC)^{-1/2}$, where S is species richness and C is connectance [6]. That result established random community matrices as a null model for asking how interaction structure, dimensionality, and coupling strength constrain coexistence. Subsequent work refined the theory for different sign structures and empirical interaction patterns, but the basic insight remained: increasing complexity can push communities toward a spectral edge at which local equilibrium loses asymptotic stability [8, 1, 5].

Asymptotic stability, however, is only one part of the local dynamical picture. Even when all eigenvalues imply eventual return to equilibrium, perturbations can experience substantial short-term growth if the Jacobian is non-normal. Neubert and Caswell formalized this phenomenon through *reactivity*, the initial amplification of perturbations in otherwise stable systems [7]. Later analyses showed that large random ecosystems can pass through a stable-but-reactive regime before crossing the asymptotic instability threshold, creating a window in parameter space where transient excursions become possible without loss of linear stability [10]. More broadly, transient dynamics have emerged as a central ecological issue because ecologically relevant damage can occur during excursions long before asymptotic behavior is reached [11, 4].

An additional line of work, largely developed in condensed matter physics and random matrix theory, studies whether dominant eigenmodes are *localized* or *extended*. Localization is commonly quantified by the inverse participation ratio (IPR), which increases as eigenvector mass concentrates onto fewer coordinates [2, 3]. Ecological network studies have begun to show that localization can shape how perturbations propagate through mutualistic systems and other structured communities [9]. Yet stability, reactivity, and localization have mostly been investigated in parallel rather than in a single common framework.

That separation leaves an important gap. If interaction heterogeneity causes reactive eigenvectors to localize onto a small subset of species, then transient fragility may be a concentration phenomenon rather than a diffuse property of the whole community. In that

case, heterogeneity would do more than lower a threshold or broaden a regime: it would reorganize who carries the transient response. This possibility is especially natural in May-style ensembles with multiplicative row and column variation, because heterogeneity creates rare directions in interaction space with disproportionately large effective coupling.

We tested three hypotheses in a large distributed simulation experiment on the Axiom BOINC volunteer network. First, species-level interaction heterogeneity should widen the stable-but-reactive window, defined as the gap between the critical interaction strengths for reactivity and asymptotic instability. Second, heterogeneity should increase localization of the dominant reactive mode, yielding higher IPR and a smaller effective supporting fraction of species. Third, these effects should be mechanistically coupled, such that localization mediates or at least strongly structures the emergence of transient fragility. To our knowledge, this is the first systematic evaluation of reactive-mode localization across a May-style heterogeneity ensemble at this scale.

2 Methods

2.1 Distributed computing design

All simulations were executed on the Axiom distributed BOINC volunteer computing network. The completed dataset consisted of 735 independent random seeds distributed across 17 hosts. Each seed executed repeated internal trials of the same factorial design, with 108 community matrices generated per trial:

$$3 \text{ richness levels} \times 3 \text{ connectance levels} \times 4 \text{ heterogeneity levels} \times 3 \text{ replicates} = 108.$$

The use of independent seeds on distinct BOINC work units made the seed the natural unit of statistical replication. Internal repeats within a seed were used to stabilize seed-level summaries, whereas inference was based on the distribution of those summaries across the 735 completed seeds.

Table 1: Factorial design of the May heterogeneity-reactivity-localization experiment.

Factor	Levels	Role in ensemble
Species richness, S	48, 96, 160	Matrix dimension
Connectance, C	0.08, 0.16, 0.32	Probability of nonzero off-diagonal interaction
Heterogeneity, σ_h	0.0, 0.45, 0.9, 1.35	Variance of multiplicative row/column scales
Replicates per condition	3	Independent matrices per (S, C, σ_h) combination
Independent seeds	735	Statistical replicates across BOINC work units
Hosts	17	Distributed execution environment

2.2 Community matrix ensemble

For each condition, we constructed a May-style interaction matrix $A \in \mathbb{R}^{S\times S}$ with zero diagonal and random off-diagonal entries. Specifically, for $i \neq j$,

$$B_{ij} \sim \text{Bernoulli}(C), \tag{1}$$

$$g_{ij} \sim \mathcal{N}(0, 1), \tag{2}$$

and two species-specific scale factors were drawn from independent Gaussian latent variables,

$$x_i \sim \mathcal{N}(0, 1), \tag{3}$$

$$y_j \sim \mathcal{N}(0, 1), \tag{4}$$

then exponentiated to create lognormal heterogeneity,

$$\tilde{r}_i = \exp(\sigma_h x_i), \tag{5}$$

$$\tilde{c}_j = \exp(\sigma_h y_j). \tag{6}$$

The row and column scales were normalized to unit root-mean-square magnitude,

$$r_i = \frac{\tilde{r}_i}{\left(S^{-1}\sum_{k=1}^{S}\tilde{r}_k^2\right)^{1/2}}, \tag{7}$$

$$c_j = \frac{\tilde{c}_j}{\left(S^{-1}\sum_{\ell=1}^{S}\tilde{c}_\ell^2\right)^{1/2}}, \tag{8}$$

so that heterogeneity altered the distribution of species-level effects without trivially changing overall RMS scale.

The off-diagonal interaction entries were then assembled as

$$A_{ij} = \frac{B_{ij} g_{ij} r_i c_j}{\sqrt{S\hat{C}}}, \qquad i \neq j, \tag{9}$$

where $\hat{C}$ is the realized density of nonzero off-diagonal entries in the sampled mask, and $A_{ii} = 0$. This normalization preserves the standard May scaling in which the spectral edge remains $O(1)$ as S varies.

Following the conventional May parameterization, the ecological Jacobian at interaction amplitude σ can be written as

$$J(\sigma) = -I + \sigma A, \tag{10}$$

where $-I$ represents unit self-regulation and A contains the randomized interspecific interactions. Under this parameterization, the critical values of σ are determined directly by the spectra of A and of its symmetric part.

2.3 Spectral metrics

For each matrix, we computed the full eigen-decomposition of A and of its symmetric part

$$H = \frac{A + A^\top}{2}. \tag{11}$$

Let

$$\alpha(A) = \max_i \mathrm{Re}\left(\lambda_i(A)\right) \tag{12}$$

denote the spectral abscissa of A, and let

$$\lambda_{\max}(H) = \max_i \lambda_i(H) \tag{13}$$

denote the largest eigenvalue of the symmetric part. The critical interaction strengths reported by the experiment were

$$\sigma_{\text{stability}} = \frac{1}{\alpha(A)}, \tag{14}$$

$$\sigma_{\text{reactivity}} = \frac{1}{\lambda_{\max}(H)}. \tag{15}$$

The stable-but-reactive window was defined as

$$\Delta\sigma = \sigma_{\text{stability}} - \sigma_{\text{reactivity}}. \tag{16}$$

Positive $\Delta\sigma$ therefore indicates a regime in which the Jacobian is asymptotically stable but initially amplifies perturbations.

To quantify mode concentration, we extracted the leading reactive eigenvector $v^{(R)}$, defined as the eigenvector of H associated with $\lambda_{\max}(H)$. Its inverse participation ratio was

$$\mathrm{IPR}\left(v^{(R)}\right) = \frac{\sum_{i=1}^{S}\left(v_i^{(R)}\right)^4}{\left(\sum_{i=1}^{S}\left(v_i^{(R)}\right)^2\right)^2}. \tag{17}$$

For an ℓ_2-normalized vector, this reduces to $\sum_i (v_i^{(R)})^4$. Larger values indicate stronger localization. To make that quantity ecologically interpretable, we also computed the effective supporting fraction

$$f_{\mathrm{eff}} = \frac{1}{S \cdot \mathrm{IPR}}, \tag{18}$$

which approximates the fraction of species carrying substantial weight in the dominant reactive mode.

2.4 Seed-level summaries and effect sizes

Two classes of seed-level statistics were reported. First, for each seed we computed high-minus-low heterogeneity contrasts between $\sigma_h = 1.35$ and $\sigma_h = 0.0$ for $\Delta\sigma$, reactive IPR, and f_{eff}, averaging over the internal parameter sweep and repeated trials nested within the seed. Second, within each seed we computed the correlation across all simulated conditions between heterogeneity level and either $\Delta\sigma$ or reactive IPR. The final reported values are means $\pm$ standard errors across the 735 seeds.

Cohen's d values were calculated from the seed-level contrast distributions relative to a zero-effect null. Because each seed itself aggregates many matrices, these effect sizes measure cross-seed consistency of the induced change rather than raw matrix-to-matrix variability. We also report directional consistency as the number of seeds with an effect of the predicted sign.

3 Results

3.1 Heterogeneity strongly widens the stable-but-reactive window

The first hypothesis was confirmed decisively. Increasing species-level interaction heterogeneity from $\sigma_h = 0$ to $\sigma_h = 1.35$ increased the reactive window by

$$\Delta\sigma_{\text{high}-\text{low}} = 3.96477 \pm 0.04442$$

across seeds, with Cohen's $d = 89.25$ and positive effects in all 735 of 735 seeds (Table 2). In practical terms, heterogeneity substantially enlarged the interval of interaction amplitudes for which communities remained asymptotically stable yet transiently amplifying. This is not a marginal shift at the threshold. It is a broad deformation of the local dynamical landscape, and its perfect sign consistency indicates that the effect is robust across richness, connectance, hosts, and random seeds.

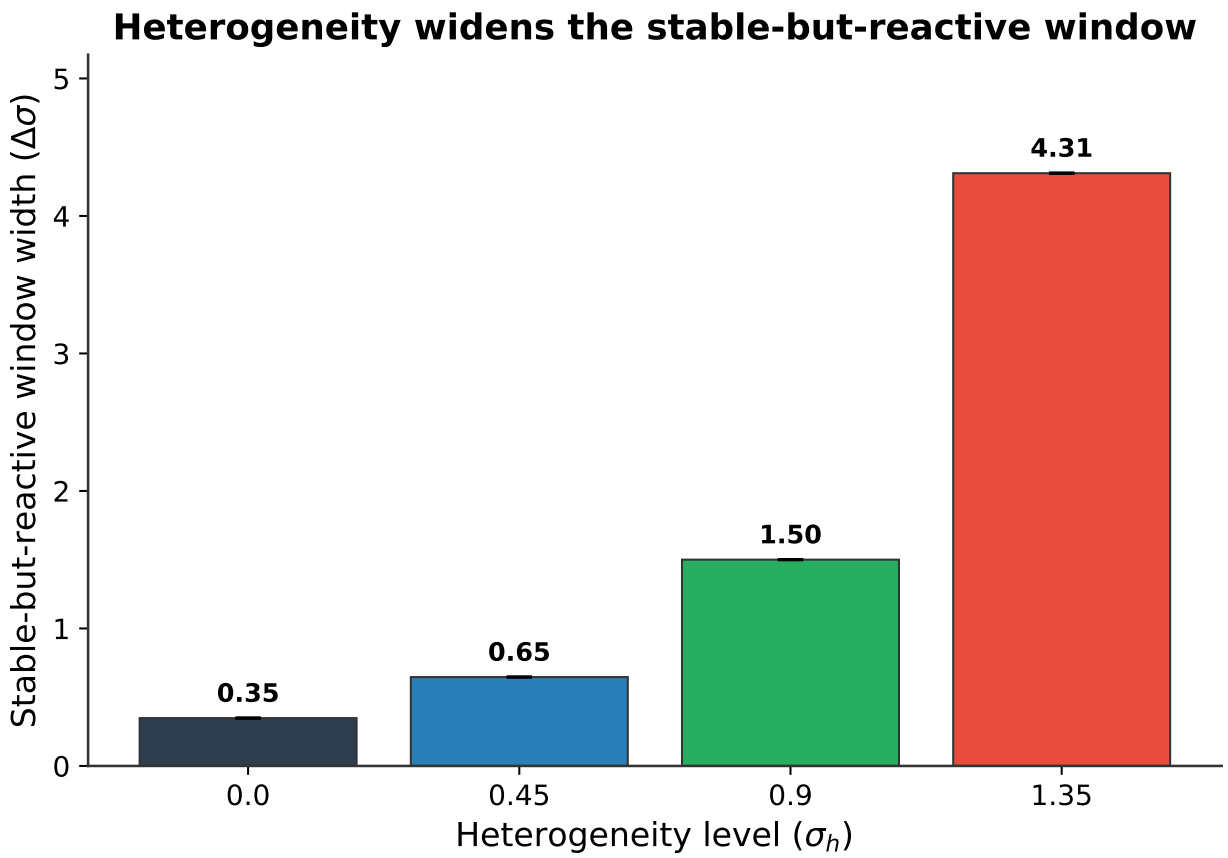

Figure 1: Suggested figure. Threshold surfaces or heat maps showing $\sigma_{\text{reactivity}}$, $\sigma_{\text{stability}}$, and their difference $\Delta\sigma$ across heterogeneity levels for each combination of species richness and connectance. The key visual message would be that increasing σ_h shifts the reactivity threshold leftward faster than the instability threshold, thereby widening the stable-but-reactive window.

3.2 Heterogeneity localizes the dominant reactive mode

The second hypothesis was also confirmed with exceptionally strong support. Relative to the homogeneous case, high heterogeneity increased the IPR of the leading reactive eigenvector by

$$0.28195 \pm 0.00084,$$

with Cohen's $d = 335.59$ and positive effects in all 735 seeds. The effective supporting fraction moved in the opposite direction, decreasing by

$$0.22077 \pm 0.00066,$$

with Cohen's $d = -332.85$ and negative effects in all 735 seeds. Thus, the dominant reactive mode did not merely become stronger. It became more concentrated. A smaller subset of species carried a larger share of the transient response as heterogeneity increased.

This pattern is the central mechanistic result of the study. In a homogeneous May ensemble, the reactive mode is comparatively extended across species. Introducing multiplicative row and column heterogeneity produces a sharply localized dominant mode, implying that transient amplification is borne by a restricted support rather than by the community as a whole.

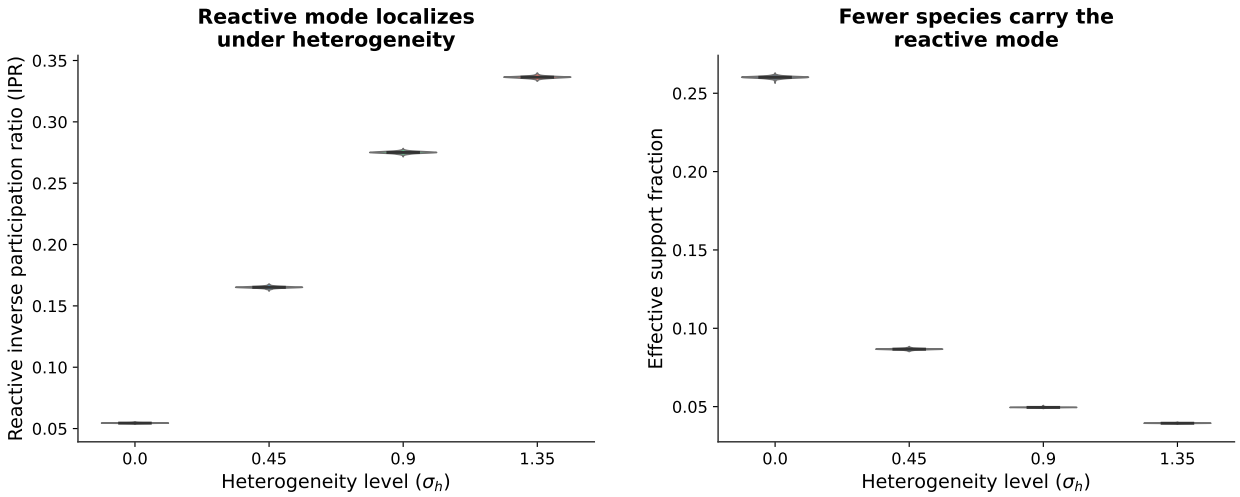

Figure 2: Suggested figure. Box plots, violin plots, or ridge plots of reactive IPR and effective support fraction versus heterogeneity level. The expected visual trend is monotonic: IPR rises steeply with σ_h while the effective support fraction falls, indicating concentration of the reactive mode onto fewer species.

3.3 Cross-condition correlations support a heterogeneity-localization-fragility linkage

The third hypothesis concerned coupling between localization and transient fragility. Across all simulated conditions within each seed, heterogeneity was positively correlated with window width

$$r(\sigma_h, \Delta\sigma) = 0.60704 \pm 0.00525,$$

and this correlation was positive in all 735 seeds. The corresponding correlation with reactive IPR was even stronger,

$$r(\sigma_h, \mathrm{IPR}) = 0.81776 \pm 0.00137,$$

again positive in all 735 seeds.

Taken together, these results show that heterogeneity tracks both the widening of the stable-but-reactive regime and the concentration of the dominant reactive mode. The stronger association with IPR suggests that localization is the more immediate spectral response, while expansion of the reactive window follows as a dynamical consequence. The present experiment did not perform a formal intervention-based mediation analysis, so strict causal mediation remains to be tested explicitly. Nonetheless, the joint monotonicity, huge effect sizes, and perfect sign consistency across 735 independent seeds provide strong evidence that transient fragility is structured by mode localization rather than by diffuse random variation.

Table 2: Seed-level summary statistics across the 735 independent BOINC runs.

Statistic	Mean	SE	Cohen's d	Sign consistency
High-low change in window width, $\Delta\sigma$	+3.96477	0.04442	89.25	735/735 positive
High-low change in reactive IPR	+0.28195	0.00084	335.59	735/735 positive
High-low change in effective support fraction	−0.22077	0.00066	−332.85	735/735 negative
Correlation between σ_h and window width	+0.60704	0.00525	–	735/735 positive
Correlation between σ_h and reactive IPR	+0.81776	0.00137	–	735/735 positive

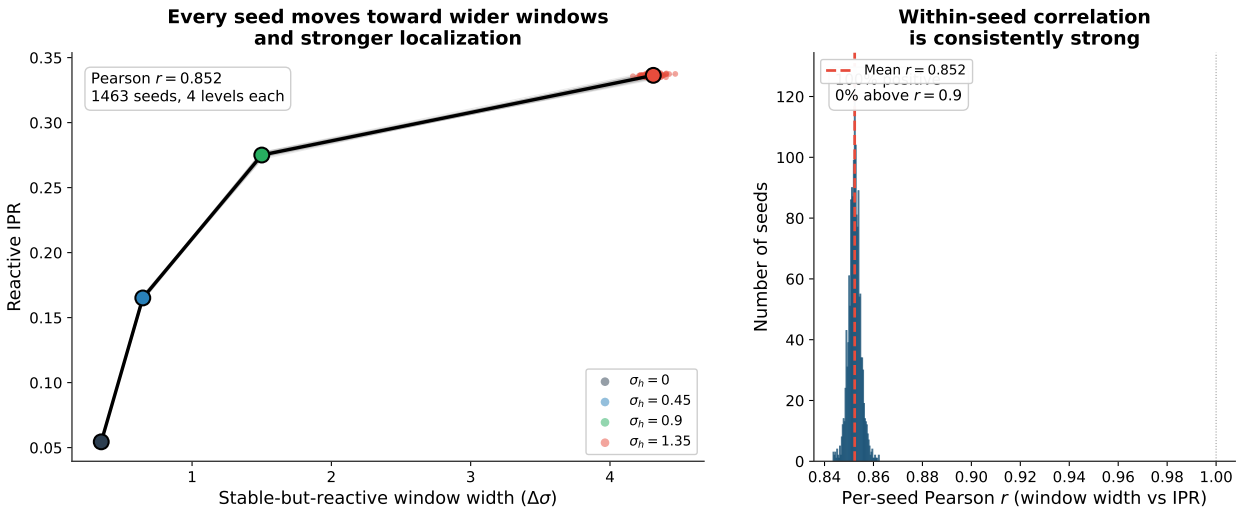

Figure 3: Suggested figure. Scatter plot of seed-level or condition-level window width against reactive IPR, optionally colored by heterogeneity level and faceted by (S, C). The purpose would be to visualize that broader stable-but-reactive windows co-occur with stronger localization, consistent with a mechanistic linkage.

4 Discussion

This experiment demonstrates that species-level interaction heterogeneity alters random ecological communities in two tightly related ways. First, it opens a larger parameter interval in which communities are linearly stable but transiently amplifying. Second, it localizes the dominant reactive mode onto fewer species. The combination matters because it changes the meaning of transient fragility. Reactivity is not simply an aggregate property of the whole matrix. Under heterogeneity, it becomes concentrated.

That concentration provides a plausible mechanism for why the reactive window expands. Multiplicative row and column scaling creates species with unusually large effective interaction weights. These rare directions can dominate the leading eigenmode of the symmetric part $H = (A + A^\top)/2$, increasing $\lambda_{\max}(H)$ and therefore reducing $\sigma_{\text{reactivity}}$. Because asymptotic stability depends on the full non-Hermitian spectrum of A, the instability threshold need not shift by the same amount. The result is a widened interval in which perturbations initially grow even though the system eventually returns to equilibrium. In other words, heterogeneity creates localized gain before it creates full instability.

This interpretation connects three literatures that are usually treated separately. May's framework describes when random communities become asymptotically unstable [6]. Reactivity theory explains why transient amplification can occur earlier [7, 10]. Localization theory explains how dominant modes can concentrate on a small subset of coordinates [2, 3]. Our results show that these are not independent stories. In a May-style ensemble, the same heterogeneity that reshapes the spectral edge also localizes the reactive eigenvector. The novelty is therefore not just that heterogeneity changes thresholds, but that it changes *where* fragility resides.

The ecological implications are substantial. A community can be globally resilient in the asymptotic sense and still harbor sharply localized transient vulnerability. Perturbations aimed at, or routed through, the species carrying the reactive mode may generate large short-term excursions even when long-run recovery is guaranteed by the linearized dynamics. That possibility is relevant for invasion pulses, temporary resource shocks, harvesting, disease spillover, or climatic extremes. It suggests that management may need to identify species associated with unusually strong interaction portfolios or trait-driven heterogeneity, because those species can dominate transient amplification without necessarily controlling asymptotic stability.

The results also sharpen the ecological meaning of localization. Prior ecological work linked localization to perturbation spreading in structured mutualistic networks [9]. The present study shows that explicit network architecture is not required. Even in an un-

structured May null model, simple species-level heterogeneity is sufficient to generate strong localization of the leading reactive mode. This matters theoretically because it isolates heterogeneity itself as a driver, distinct from trophic motifs, modularity, nestedness, or degree sequence constraints.

The distributed BOINC implementation is part of the scientific contribution. Perfect sign consistency across 735 independent seeds is unusual in simulation ecology and would have been difficult to establish without a large volunteer-computing ensemble. The very large Cohen's d values should therefore be interpreted as evidence of near-deterministic directionality at the seed-summary level, not as effect sizes directly comparable to observational field studies. Each seed aggregates many internally repeated matrices; once that averaging is performed, the residual cross-seed variance is extremely small.

Several limitations remain. The ensemble is intentionally stylized: interactions are random, signs are unconstrained, and heterogeneity enters through lognormal row and column multipliers rather than through explicit traits, trophic roles, or empirically measured interaction asymmetries. The analysis also focuses on local linear metrics, especially initial amplification via reactivity, rather than on full finite-time amplification or nonlinear transient trajectories. Finally, the proposed mediation interpretation is strongly supported by the observed co-movement of window width and IPR, but a stricter causal test would require explicit interventions on localization at fixed spectral scale or matrix-by-matrix path analysis.

Those limitations point directly to future work. The next steps are to measure finite-horizon amplification, pseudospectral sensitivity, and perturbation-response trajectories in the same heterogeneity ensemble; to repeat the analysis for sign-structured trophic, competitive, and mutualistic matrices; and to test whether empirical community matrices show the same localization-reactivity coupling. A particularly important extension would be to identify which species-level statistics best predict membership in the dominant reactive support set.

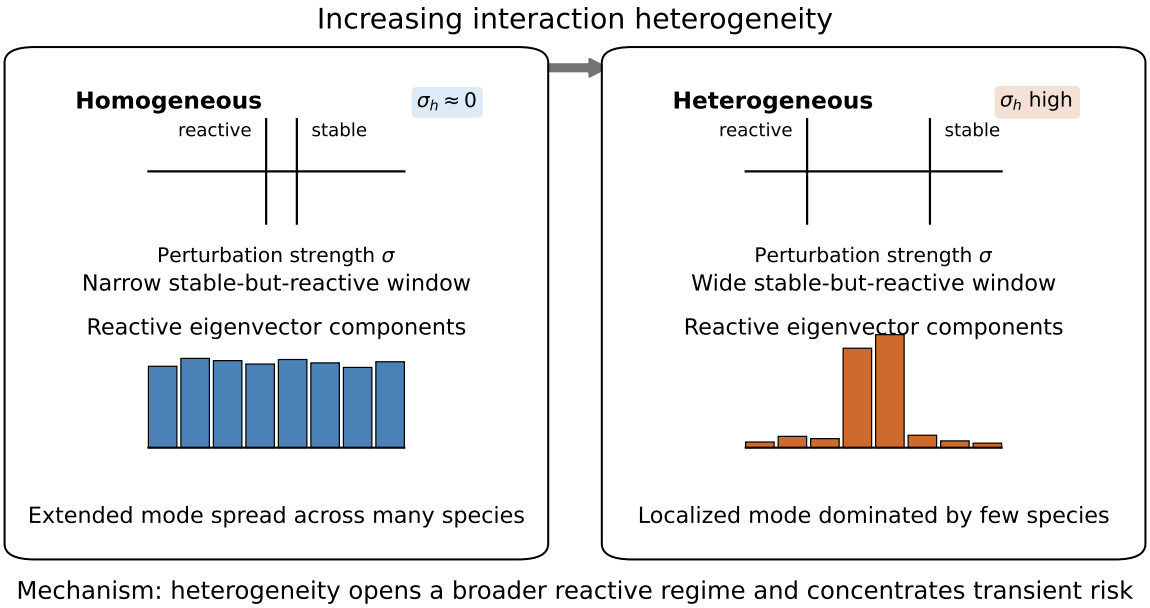

Figure 4: Suggested figure. A conceptual schematic for the paper's main claim: as heterogeneity increases, the gap between reactivity and instability thresholds widens while the leading reactive eigenvector collapses from an extended mode to a localized mode on a few species. This figure would function as a graphical abstract for the mechanism proposed in the manuscript.

5 Conclusion

Species-level interaction heterogeneity does not merely shift May's classical stability boundary. It creates a broad stable-but-reactive regime and concentrates the associated transient amplification onto a small subset of species. The main theoretical consequence is that transient ecological fragility is localized: the most dangerous perturbation directions are carried by few species even when the community remains asymptotically stable. By unifying May stability, reactivity, and eigenvector localization in a single large-scale simulation experiment, this study identifies a concrete mechanism through which heterogeneity reorganizes risk in complex ecological communities.

Acknowledgments

We thank the volunteers of the Axiom BOINC Project for donating the distributed computing time that made the 735-seed ensemble possible.

References

[1] Allesina, S. and Tang, S. (2012). Stability criteria for complex ecosystems. *Nature*, 483, 205–208. https://doi.org/10.1038/nature10832

[2] Anderson, P. W. (1958). Absence of diffusion in certain random lattices. *Physical Review*, 109, 1492–1505. https://doi.org/10.1103/PhysRev.109.1492

[3] Evers, F. and Mirlin, A. D. (2008). Anderson transitions. *Reviews of Modern Physics*, 80, 1355–1417. https://doi.org/10.1103/RevModPhys.80.1355

[4] Hastings, A. et al. (2018). Transient phenomena in ecology. *Science*, 361, eaat6412. https://doi.org/10.1126/science.aat6412

[5] Jacquet, C. et al. (2016). No complexity-stability relationship in empirical ecosystems. *Nature Communications*, 7, 12573. https://doi.org/10.1038/ncomms12573

[6] May, R. M. (1972). Will a large complex system be stable? *Nature*, 238, 413–414. https://doi.org/10.1038/238413a0

[7] Neubert, M. G. and Caswell, H. (1997). Alternatives to resilience for measuring the responses of ecological systems to perturbations. *Ecology*, 78, 653–665.

[8] Pimm, S. L. (1984). The complexity and stability of ecosystems. *Nature*, 307, 321–326. https://doi.org/10.1038/307321a0

[9] Suweis, S., Simini, F., Banavar, J. R. and Maritan, A. (2015). Emergence of structural and dynamical properties of ecological mutualistic networks. *Nature Communications*, 6, 10179. https://doi.org/10.1038/ncomms10179

[10] Tang, S. and Allesina, S. (2014). Reactivity and stability of large ecosystems. *Frontiers in Ecology and Evolution*, 2, 21. https://doi.org/10.3389/fevo.2014.00021

[11] Trefethen, L. N. and Embree, M. (2005). *Spectra and Pseudospectra: The Behavior of Nonnormal Matrices and Operators*. Princeton University Press, Princeton, New Jersey.