

RESEARCH ARTICLE

Collective Ecological Protection: A Systems-Level Theory of Microbial Community Stability

Estefan Monteiro da Fonseca^{1*}, Christine Gaylarde²

¹Department of Geology and Geophysics/LAGEMAR – Marine Geology Laboratory, Institute of Geosciences, Universidade Federal Fluminense (UFF), Avenida Litorânea s/n, 24210-340 Niterói, RJ, Brazil

²Department of Microbiology and Plant Biology, University of Oklahoma, 770 Van Vleet Oval, Norman, OK 73019, USA

*Corresponding author: Estefan Monteiro da Fonseca; oceano25@hotmail.com

ABSTRACT

Microbial communities persist under recurrent abiotic and biotic disturbance, yet the organizational basis of this persistence remains incompletely explained. Existing frameworks describe resistance, resilience, functional redundancy, social evolution, and ecological networks, but do not explicitly define how demonstrated protective effects are distributed among community members or whether this distribution predicts stability beyond diversity, composition, and functional potential. Here, we propose the Collective Ecological Protection Theory (CEPT), a systems-level framework in which stability is influenced by a protective architecture: the directed, weighted, disturbance- and response-specific network through which the activity of a provider increases the persistence or functional performance of a recipient relative to a matched counterfactual. CEPT separates protective mechanisms from their organization and identifies six candidate descriptors: protective coverage, protective redundancy, mechanistic complementarity, spatial reach, modularity, and dependency concentration. Protection is context dependent and may be weakened by spatial separation, correlated failure, costly cooperation, exploitation, or dependence on a small number of providers. We derive four discriminating predictions based on matched-community rewiring, targeted provider loss, independent protection, and architectural reconstruction during recovery. CEPT receives support to the extent that independently measured protective architecture improves causal or out-of-sample prediction of a specified stability response beyond established compositional, functional, and environmental factors. It therefore offers a testable organizational explanation for how distinct protective processes combine to influence microbial resistance, persistence, recovery, and maintenance of function. By identifying distributed protection and single points of failure, CEPT may also provide a mechanistic basis for assessing ecological vulnerability and risk in disturbed microbial systems.

Keywords: microbial ecology; community stability; collective ecological protection; protective architecture; microbial interactions; ecological networks; ecological risk; ecological security

ARTICLE INFO

Received: 04 September 2026 | Accepted: 14 October 2026 | Available online: 27 October 2026

CITATION

Fonseca E. M. d, Gaylarde C. Collective Ecological Protection: A Systems-Level Theory of Microbial Community Stability. Ecological Risk and Security Research. 2026; 4(1): 15090 doi:10.59429/ersr.v4i1.15090

COPYRIGHT

Copyright © 2026 by author(s). *Ecological Risk and Security Research* is published by Arts and Science Press Pte. Ltd. This is an Open Access article distributed under the terms of the Creative Commons Attribution License (<https://creativecommons.org/licenses/by/4.0/>), permitting distribution and reproduction in any medium, provided the original work is cited.

1. Introduction

Microbial communities repeatedly encounter abiotic disturbances—including shifts in temperature,

moisture, salinity, nutrient supply, and chemical exposure—and biotic disturbances such as antimicrobial attack, contact-dependent antagonism, and invasion by non-resident populations (Allison and Martiny 2008; van Elsas et al. 2012; Shade et al. 2012; de Vos et al. 2017; Granato et al. 2023; Philippot et al. 2021). These disturbances can alter community composition, resource availability, physicochemical conditions, and ecological interactions. Explaining how microbial communities persist under such pressures requires distinguishing among several dimensions of stability, including resistance during disturbance, recovery after disturbance, temporal invariability, population persistence, and maintenance of ecosystem function (Grimm and Wissel 1997; Ives and Carpenter 2007; Shade et al. 2012; Philippot et al. 2021). These dimensions need not covary: substantial taxonomic turnover may occur with little functional change, while compositional, functional, and network-level responses may differ in both their resistance and recovery trajectories (Louca et al. 2016; Gao et al. 2022; Oprei et al. 2023). Identifying the processes that connect microbial organization to particular stability outcomes therefore remains a central challenge in microbial ecology.

Established frameworks provide complementary explanations for community stability. Biodiversity can buffer ecosystem functioning against environmental variation, functional redundancy can compensate for taxonomic loss, and resistance and resilience characterize community responses during and after disturbance (Yachi and Loreau 1999; Allison and Martiny 2008; Loreau 2010; Louca et al. 2018). Ecological network theory further describes how interaction topology, modularity, centrality, and the balance of positive and negative links can influence community dynamics (Faust and Raes 2012; Coyte et al. 2015; Hernandez et al. 2021; Kajihara and Hynson 2024). Together, these approaches establish that stability depends not only on which organisms and functions are present but also on how ecological interactions are organized.

Experimental microbial ecology has demonstrated that the activity of one community member can modify the exposure, susceptibility, or physiological response of another member to abiotic and biotic disturbances. Extracellular matrices, antibiotic degradation, metabolic exchange, environmental detoxification, nutrient provisioning, and spatial organization can generate benefits that extend beyond the organisms performing the underlying activity (Lee et al. 2014; de Vos et al. 2017; Pande and Kost 2017; Piccardi et al. 2019; Granato et al. 2023). Experimental studies have shown that mixed-species biofilms can exhibit enhanced antimicrobial tolerance through interspecies protection, while the spatial arrangement of microbial populations influences the accessibility and distribution of shared extracellular benefits (Lee et al. 2014; Nadell et al. 2016; Granato et al. 2023). Moreover, metabolic exchanges and environmentally mediated facilitative interactions demonstrate that the ecological consequences of microbial activities depend on recipient requirements and environmental context (Pande and Kost 2017; Piccardi et al. 2019). Further evidence from facilitative interaction networks, collective antibiotic tolerance, extracellular-polymer-mediated defense, and colonization resistance indicates that these interactions may extend beyond individual physiological responses to influence recipient persistence and community-level outcomes (Kelsic et al. 2015; de Vos et al. 2017; Fujita et al. 2023; Granato et al. 2023; Spragge et al. 2023).

Importantly, protective interactions may also depend on ecological relationships that sustain the

organisms responsible for providing protection. Metabolic exchanges and resource provisioning can influence the abundance and activity of potential protective providers before or during disturbance. For example, exudates released by *Prochlorococcus* have been shown to stimulate heterotrophic activity in *Alteromonas*, influencing competitive interactions involving other cyanobacteria under nitrogen limitation (Calfee et al. 2022). Although this finding does not establish that reciprocal resource exchange maintains protection against future disturbances, it illustrates how recipient-associated activities may influence the ecological conditions supporting potential protective providers. Consequently, the persistence of protective capacity may depend not only on the immediate effects of providers on recipients but also on the ecological interactions that sustain those providers over time.

These contributions establish important precedents, demonstrating that microbial interactions can influence community responses to disturbance through diverse physiological and ecological mechanisms. Previous studies have also advanced beyond individual interactions by examining the organization and dynamics of facilitative networks and identifying influential species within microbial communities (Fujita et al. 2023; Angulo et al. 2019). Nevertheless, these approaches address distinct aspects of microbial interaction networks and community dynamics, rather than explicitly integrating experimentally demonstrated, disturbance-specific protective effects into a common framework linking providers, recipients, and defined stability outcomes. In particular, it remains insufficiently resolved how demonstrated protective effects are distributed among providers and recipients across an entire community, whether multiple providers contribute independently or only in combination, and whether general properties of this organization predict a specified stability response after diversity, composition, and functional potential have been accounted for. The distinction between interactions that maintain protective capacity and those that provide effective protection during disturbance introduces an additional temporal dimension to this unresolved organizational problem. This unresolved relationship between the causal identification of protective interactions and their collective organization constitutes the central knowledge gap addressed by the Collective Ecological Protection Theory (CEPT).

To address this gap, we introduce the Collective Ecological Protection Theory (CEPT). CEPT defines protection as a causal and context-dependent effect in which the presence or activity of a provider improves the persistence or functional performance of a recipient under a specified disturbance relative to an appropriate counterfactual. Interactions that sustain a provider may contribute indirectly to future protection, but they do not themselves qualify as protective effects unless a benefit to a defined recipient is demonstrated under the focal disturbance. The directed and weighted organization of these effects constitutes the protective architecture. The distinctive proposition of CEPT is not merely that protective interactions exist, but that differences in their distribution, strength, independence, accessibility, and concentration can explain variation in microbial community stability beyond the presence of particular taxa, genes, or mechanisms. This proposition also accommodates temporal changes in protective architecture as providers and their supporting ecological relationships change. From an ecological-risk perspective, these architectural properties can be

interpreted as potential determinants of vulnerability: communities with sparse coverage, inaccessible protection, or highly concentrated dependencies may be more susceptible to population or functional loss when disturbance disrupts critical providers or pathways.

CEPT therefore shifts the analytical focus from asking only which organisms or functions are present to asking who protects whom, through which mechanism, under which disturbance, across what spatial range, and with what degree of redundancy or dependency. It further recognizes that protective capacity may depend on ecological relationships that sustain providers across changing environmental conditions. We first define the theory and derive its operational unit from established ecological and experimental precedents. We then formalize the protective architecture and its candidate descriptors, distinguish CEPT from neighboring frameworks, and derive predictions whose experimental failure would require revision or rejection of the theory (Figure 1).




Figure 1. Conceptual position of the Collective Ecological Protection Theory (CEPT). Established frameworks characterize dimensions and determinants of ecological stability, while experimental studies identify particular mechanisms of community-level protection. CEPT addresses the intervening organizational question by integrating demonstrated, stress-contingent provider–recipient effects into a directed and weighted protective architecture and asking whether this architecture predicts a specified stability response beyond diversity, composition, and functional potential. Ecological interactions that maintain potential protective providers may influence the availability and temporal development of this architecture, but are distinguished from protective edges unless their recipient-level effects under disturbance are causally demonstrated.

2. Collective Ecological Protection Theory

The Collective Ecological Protection Theory (CEPT) proposes that, for a specified disturbance and ecological response, variation in microbial community stability is influenced by the organization of stress-contingent protective effects among community members (Coyte et al. 2015; Fujita et al. 2023). CEPT does not infer protection from taxonomic identity, co-occurrence, or the mere presence of a potentially beneficial mechanism (Faust and Raes 2012). Evidence must demonstrate that the presence or activity of a provider improves the performance of a recipient under the focal disturbance relative to an appropriate counterfactual (Morris et al. 2011; Piccardi et al. 2019). Such evidence requires experimental designs capable of distinguishing the provider's protective contribution from concomitant changes in community composition, resource availability, or environmental conditions (Piccardi et al. 2019). The organization formed by these directional effects constitutes the protective architecture.

CEPT does not assume that protective interactions explain why community members coexist under baseline conditions. In the absence of the focal disturbance, coexistence may result from resource partitioning, spatial segregation, dispersal, weak competition, metabolic exchange, or other assembly processes (Goldford et al. 2018; Pande and Kost 2017). A relationship that is neutral or competitive under baseline conditions may become protective when a disturbance activates or reveals an effect that benefits another community member (Brooker et al. 2008; Maestre et al. 2009; Piccardi et al. 2019). Protection is therefore a disturbance- and response-specific classification rather than a permanent property of a taxon or relationship. Furthermore, interactions that sustain potential protective providers under baseline conditions need not themselves constitute protective effects. Experimental evidence that microbial exudates can stimulate the activity of heterotrophic partners supports the plausibility of such maintenance relationships (Calfee et al. 2022). However, whether these interactions preserve protective capacity against subsequent disturbances remains an open empirical question. Such interactions may influence the availability of protection during subsequent disturbances, introducing a distinction between the ecological maintenance of protective capacity and its effective expression under stress.

The *Alteromonas*–*Prochlorococcus* system provides a practical illustration of this conditionality. Under extracellular H₂O₂ stress, *Alteromonas* sp. EZ55 removes peroxide from the shared environment and improves the growth, survival, and photosynthetic performance of *Prochlorococcus* UH18301 relative to matched cultures lacking the helper (Morris et al. 2011). Within CEPT, this result supports a directional protective effect from *Alteromonas* to *Prochlorococcus* specifically under H₂O₂ stress and for the measured recipient responses. Related work has also demonstrated that organic compounds released by *Prochlorococcus* can stimulate the activity of *Alteromonas macleodii* and alter competitive interactions with *Synechococcus* under nitrogen limitation (Calfee et al. 2022). Although these studies involve different experimental systems and do not establish a reciprocal protective loop, they illustrate how interactions originating from a potential recipient may influence the ecological activity of heterotrophic organisms capable of providing benefits in other contexts. Repeating equivalent causal comparisons across community members would reveal whether

peroxide protection is broadly distributed among several independent providers or concentrated in a single helper. CEPT predicts that these contrasting architectures would differ in their capacity to maintain community stability, particularly following the loss or inhibition of a major provider. This prediction remains to be tested by comparing communities with similar initial protective capacity but different degrees of provider redundancy and dependency.

Several established research traditions provide the conceptual and experimental elements needed to recognize such effects. Social-evolutionary theory distinguishes actors and recipients according to the consequences of an activity for each participant, whereas facilitation ecology describes benefactor–beneficiary relationships and demonstrates that their magnitude and direction can vary along environmental stress gradients (West et al. 2007; Brooker et al. 2008; Maestre et al. 2009). Microbial disturbance ecology distinguishes among responses such as resistance, persistence, recovery, temporal invariability, and maintenance of function (Shade et al. 2012). Experimental microbial studies have further employed monoculture–coculture comparisons, provider removal, pathway inhibition, mutant comparison, and environmental manipulation to determine whether one organism improves another’s performance under stress (Morris et al. 2011; de Vos et al. 2017; Piccardi et al. 2019; Granato et al. 2023). Together, these approaches demonstrate that identifying a beneficial interaction requires attention to the organism receiving the benefit, the environmental conditions under which it occurs, and the causal consequences of the underlying activity.

These traditions therefore supply the principal elements required to recognize protective effects, but they generally examine them using different terminologies, mechanisms, and levels of organization. CEPT integrates these precedents into a single operational criterion that can be applied consistently across mechanisms and community members. Operationally, a protective interaction identifies five elements: a provider, a recipient, a specified disturbance, a measurable recipient-level response, and an appropriate counterfactual. The provider is the organism, functional group, or experimentally defined set whose activity may generate protection. The recipient is the community member whose response is evaluated. The disturbance specifies the abiotic or biotic condition under which protection is proposed to occur. The recipient-level response defines the measured outcome, such as survival, persistence, growth, or functional performance. The counterfactual is an otherwise comparable condition in which the provider or relevant activity is absent, removed, inhibited, or experimentally suppressed. Where provider removal changes multiple environmental or biological variables simultaneously, pathway-specific inhibition, complementation, or rescue experiments may be necessary to isolate the protective mechanism and reduce confounding.

Only when these elements are specified can an effect be classified as a protective interaction within CEPT. Quorum sensing, extracellular enzyme activity, biofilm formation, cross-feeding, detoxification, and nutrient provisioning are therefore candidate mechanisms rather than intrinsically protective traits. Their presence establishes protective potential, but their inclusion in the protective architecture requires evidence of a net benefit to a defined recipient under the focal disturbance. Importantly, the identification of a protective effect does not require the provider to benefit from the interaction, nor does it imply that the recipient contributes to

maintaining the provider. Conversely, a reciprocal metabolic interaction may support the persistence of a potential provider without qualifying as a protective edge unless the corresponding recipient-level benefit under disturbance is independently demonstrated.

Empirical studies provide examples of how this operational criterion can be applied at different levels of biological organization (Table 1).

These examples represent different degrees of correspondence with the CEPT operational unit. The *Alteromonas–Prochlorococcus* and EPS-producing–non-producing *E. coli* systems provide particularly clear directional provider–recipient effects supported by experimental counterfactuals. The metal-working-fluid consortium demonstrates that the sign and magnitude of microbial interactions can change with the disturbance context. The gut-community experiment demonstrates a collective provider-set effect at the level of colonization resistance, although it does not by itself resolve every provider–recipient edge within the resident community. Thus, these studies establish the existence of local protective effects and collectively generated protection, but they do not individually test the broader CEPT proposition that the organization of such effects predicts community stability beyond diversity, composition, and functional potential.

The provider in CEPT may consequently be a taxon, strain, functional guild, or experimentally defined set of community members. When protection emerges only through complementary activities of multiple providers, the relevant causal unit is the provider set rather than any individual member. Pairwise effects constitute the first-order representation of protection, whereas effects that cannot be adequately explained by the contributions of individual providers may require composite provider units or explicit higher-order interaction terms (Levine et al. 2017; Grilli et al. 2017). Importantly, collective protection is not necessarily synonymous with synergistic or non-additive protection: multiple providers may contribute additively, redundantly, or through interactions that modify one another's effects. Distinguishing these possibilities requires comparisons among individual-provider, combined-provider, and appropriate control treatments. This distinction prevents an emergent community effect from being incorrectly attributed to a taxon that is ineffective in isolation.

Table 1. Empirical examples of stress-contingent protective interactions and provider-set effects in microbial communities.

Provider or provider set	Recipient or protected response	Specified disturbance	Protective activity	Experimental evidence
<i>Alteromonas</i> sp. EZ55	<i>Prochlorococcus</i> UH18301	Oxidative stress caused by extracellular H ₂ O ₂	Removal of H ₂ O ₂ from extracellular environment	Dilute axenic <i>Prochlorococcus</i> cultures exhibited poor growth, whereas medium conditioned by EZ55 supported robust growth. Reintroducing H ₂ O ₂ into the conditioned medium eliminated the benefit, identifying peroxide removal as the protective effect (Morris et al. 2011).

<i>Agrobacterium tumefaciens</i>, <i>Comamonas testosteroni</i>, <i>Microbacterium saperdae</i>, <i>Ochrobactrum anthropi</i>	Other members of the four-species and consortium	Toxicity of metal-working fluid containing biocides	Collective modification and increasing nutrient availability shifted detoxification of interactions toward competition, the shared medium demonstrating that facilitation depended on the stressful environment (Piccardi et al. 2019).
EPS-producing <i>Escherichia coli</i>	Neighboring non-EPS-producing <i>E. coli</i>	Contact-dependent attack by the T6SS of <i>Acinetobacter baylyi</i>	EPS sharing and spatial shielding of susceptible cells
Diverse human gut bacterial consortia, frequently including key <i>E. coli</i> strains	Community-level colonization resistance	Invasion by <i>Klebsiella pneumoniae</i> or <i>Salmonella enterica</i> serovar Typhimurium	Individual species produced negligible suppression, whereas appropriately composed multispecies communities strongly restricted pathogen growth. Nutrient-use overlap of predicted protection, supporting a higher-order provider-set effect generated by the collective nutrient blocking (Spragge et al. 2023). This result establishes protection at the consortium level, without resolving every directed protective interaction among resident species.

The examples in Table 1 should be interpreted as mechanistic precedents rather than direct experimental validations of CEPT. None independently establishes that a quantified protective architecture predicts community stability beyond established compositional or functional descriptors.

A localized physiological response becomes ecologically protective when its effect extends beyond the responding organism (Figure 2). A microorganism may, for example, reduce the extracellular concentration of a toxic compound, release a limiting metabolite, degrade an antibiotic, produce a shared matrix, or modify a microhabitat in a manner that increases the performance of a neighboring population. Protection is demonstrated by comparing the recipient response in the presence and absence of the provider or protective activity while holding other relevant conditions as constant as possible. The causal interpretation of this comparison depends on whether alternative explanations, including changes in resource competition, population density, or abiotic conditions, have been adequately controlled. Repeating this causal comparison

across community members transforms individual protective effects into a directed map of who protects whom under the specified disturbance. CEPT begins at this transition from isolated protective effects to their organization across the community.




Figure 2. Causal identification of a protective interaction and its scaling to protective architecture. Protection is demonstrated when a provider, provider set, or protective activity improves a specified recipient response under disturbance relative to a matched counterfactual. Repetition of this comparison across community members produces a directed and weighted map of stress-contingent protective effects. Higher-order effects are represented by provider sets when protection cannot be attributed to any member acting in isolation. Interactions that maintain potential providers without demonstrating recipient-level protection under the focal disturbance remain conceptually distinct from the protective edges represented in the architecture.

Protective interactions can make physiological capabilities ecologically accessible beyond the organisms that express them. Environmental modification, metabolite exchange, extracellular catalysis, detoxification, spatial shielding, and shared matrices may expand the conditions under which recipient populations persist. CEPT refers to these demonstrable recipient-level benefits as distributed protective capacity, without assuming that every physiological trait expressed by one organism becomes a property of the entire community. Only effects demonstrated under the focal disturbance and oriented toward a specified response enter the protective architecture. The persistence of this capacity may nevertheless depend on additional ecological interactions that support provider survival or activity across time. These maintenance relationships can be investigated as potential determinants of the availability and temporal development of protective architecture, without automatically classifying them as protective interactions.

2.1. Scope and causal identification of protection

The operational unit of CEPT is a disturbance-conditional effect, not an intrinsic taxonomic property. The five required elements are an identifiable provider (or provider set), a recipient, a prespecified disturbance, an

oriented recipient response, and a matched counterfactual. This distinction builds on experimental evidence showing that the direction and magnitude of microbial interactions can change with environmental stress, including transitions between facilitation and competition (Piccardi et al. 2019). Baseline facilitation without the focal disturbance may help establish biological plausibility, but it does not by itself qualify an edge as protective under that disturbance. Similarly, the presence of a metabolic capability or a statistically inferred association does not establish a causal protective effect on a particular recipient (Faust and Raes 2012; Weiss et al. 2016; Carr et al. 2019).

The counterfactual must distinguish the contribution of the provider or its activity from other changes introduced by the experimental manipulation. Provider-removal experiments, for example, may simultaneously alter resource availability, competitive interactions, or population density. Where feasible, selective inhibition, functional complementation, or restoration of the proposed protective activity should therefore be used to strengthen causal identification. Because correlation-based network inference may be sensitive to methodological assumptions and data structure, inferred associations should be treated as candidates for experimental validation rather than direct evidence of protection (Weiss et al. 2016; Carr et al. 2019). The protective effect is defined for the specified recipient response, observation window, and experimental context, without assuming that the same effect persists across other disturbances or spatial scales.

For a prespecified response scale, an experimentally testable two-provider interaction contrast is $\Delta_{AC,j}(s) = Y_j(A,C,s) - Y_j(A,0,s) - Y_j(0,C,s) + Y_j(0,0,s)$. A nonzero contrast identifies departure from additivity on that scale, given suitable factorial controls; it does not automatically establish a beneficial synergy or universal higher-order effect. Higher-order terms for larger sets require correspondingly specified lower-order contrasts and uncertainty estimates (Sanchez 2019; Sanchez-Gorostiaga et al. 2019).

Where a benefit emerges exclusively from a provider combination, the combination is represented as a higher-order unit rather than assigning unsupported pairwise effects. Such effects must be distinguished from additive contributions by individual providers. Higher-order interactions are relevant when the combined effect cannot be adequately represented by independently estimated lower-order contributions (Levine et al. 2017; Grilli et al. 2017; Sanchez-Gorostiaga et al. 2019; Sanchez 2019). Experimental studies of microbial consortia have demonstrated that higher-order interactions can modify community-level functional outcomes in ways that cannot be adequately inferred from pairwise combinations alone (Sanchez-Gorostiaga et al. 2019). This operational restriction allows the theory to make predictions that can be rejected by targeted intervention.

A matched intervention should preserve the focal stressor, recipient definition, observation window and relevant environmental conditions as far as feasible. Provider deletion estimates a total removal effect, whereas activity-specific inhibition or complementation targets a mechanistic effect; these estimands should not be treated as interchangeable. When collateral changes cannot be controlled, their influence should be quantified and the causal claim qualified. This distinction also applies to interactions involving provider sets.

An additional distinction concerns interactions that maintain protective capacity before disturbance. Resource exchange, metabolic facilitation, or other ecological relationships may sustain the abundance or

activity of potential protective providers without directly constituting protective effects. These maintenance interactions may influence which protective pathways become available during subsequent disturbances, but their contribution to future protection must be evaluated separately rather than assumed from baseline coexistence (Goldford et al. 2018; Calfee et al. 2022).

2.2. Fundamental principles

The organizational perspective developed by CEPT rests on five fundamental principles. Two non-equivalent evidentiary routes are envisaged: prospective predictive validation, which tests whether independently estimated architecture improves held-out forecasts, and randomized intervention-based tests, which assess whether changing protective pathways produces a corresponding response. Prediction alone is not evidence of causal mediation, and an intervention effect alone does not demonstrate broad generality.

Disturbance and response specificity. A taxon, activity, or interaction is protective only with respect to a specified disturbance and response variable. The same relationship may be protective under one disturbance, neutral under another, and detrimental under different environmental or biological conditions (Shade et al. 2012; Piccardi et al. 2019). Consequently, a protective effect established for one stressor or recipient response cannot automatically be generalized to another without additional evidence.

Directionality and relationality. Protection links a provider, or an experimentally defined provider set, to a recipient through a measurable effect. CEPT therefore does not assign a fixed protective identity to a taxon, trait, or mechanism. Directionality specifies whose activity generates the benefit and which recipient receives it. Reciprocal interactions may occur, but protection in one direction does not establish protection in the opposite direction. This distinction is particularly relevant when metabolic exchanges sustain potential providers without demonstrating reciprocal protection under the focal disturbance (Pande and Kost 2017; Calfee et al. 2022).

Causal and independent characterization. Protective effects require an appropriate counterfactual and must be characterized before or independently of the community-level stability outcome they are subsequently used to explain. This separation is essential to avoid circular reasoning in which an interaction is classified as protective merely because it occurs in a stable community. Experimental manipulation and matched counterfactuals provide a basis for causal identification, whereas independently characterized protective effects can subsequently be evaluated as predictors of community-level responses. As emphasized by Carr et al. (2019), associations inferred from microbial community data cannot, by themselves, establish direct ecological interactions, reinforcing the need for independent experimental validation of protective effects. The predictive contribution of CEPT must therefore be tested against models incorporating established compositional, functional, and environmental covariates.

Dynamic and historical contingency. Succession, immigration, extinction, spatial reorganization, recurrent disturbance, phenotypic plasticity, horizontal gene transfer, and evolution may add, remove, strengthen, or redirect protective effects. CEPT therefore does not assume that a community must preserve a fixed set of taxa or interactions to remain stable (Renes et al. 2020; Canarini et al. 2021; Shade et al. 2012;

Fukami 2015). Historical contingency and priority effects can influence community assembly by modifying which species establish and how subsequent interactions develop (Fukami 2015). Furthermore, protective capacity may depend on ecological relationships that sustain potential providers before disturbance. Changes in these supporting relationships could alter the availability of protection even when the community initially retains similar functional capabilities. Accordingly, CEPT distinguishes the maintenance of potential protective capacity from the effective protection observed under a specified disturbance. The temporal development of protective architecture is therefore an empirical question rather than an assumption of increasing protection with community age or complexity.

The architecture effect. After established covariates have been controlled, variation in the coverage, distribution, independence, accessibility, and concentration of protective effects should explain variation in a specified stability response. This is the principal proposition that distinguishes CEPT from the simple observation that protective interactions exist. Although previous network-based approaches have demonstrated that ecological interaction structure can influence microbial community dynamics and identified influential species within complex communities (Fujita et al. 2023; Angulo et al. 2019), CEPT proposes a more specifically defined test: whether independently characterized, disturbance-conditional protective architecture provides additional causal or out-of-sample predictive information beyond established descriptors. The magnitude, direction, and generality of this architectural contribution remain open empirical questions.

These principles imply that ecological importance depends not only on abundance, metabolic output, or generic network centrality, but also on stress-specific contribution to recipient persistence. A low-abundance provider may be disproportionately important when it supplies a unique protective pathway to multiple vulnerable recipients. Conversely, a highly connected provider may generate broad short-term protection while simultaneously creating a single point of failure. CEPT therefore distinguishes protective centrality from abundance and from centrality inferred from undirected association networks (Power et al. 1996; Banerjee et al. 2018; Kajihara and Hynson 2024). This distinction does not imply that protective centrality is universally superior to established network metrics; rather, its additional explanatory value must be demonstrated through direct comparison with competing predictors.

The five operational elements define which effects may enter the protective architecture, whereas the principles establish the conditions under which their organization can be evaluated as an explanation of community stability. Interactions that sustain potential protective providers may influence the development of this architecture over time, but they remain analytically distinct from the disturbance-specific protective effects used to construct it. The following section formalizes this architecture and introduces the descriptors used to characterize its structure.

3. Protective Architecture: Operational Definition and Descriptors

Protective architecture is the directed, weighted, disturbance- and response-specific network of experimentally supported protective effects among community members. Nodes represent taxa, strains, functional guilds, or experimentally defined populations or provider sets. A directed edge from provider “*i*”

to recipient “*j*” is present when the activity of *i* improves the oriented response of *j* under disturbance *s* relative to a matched counterfactual. Because the same pair may have different effects under different disturbances or response definitions, CEPT assumes no context-free protective architecture.

Protective architecture is an organizational property rather than a physical object, although spatial structure determines whether an effect reaches its recipient. It is measured before or independently of the focal stability outcome. This separation prevents architecture from being defined as “robust” merely because it was observed in a stable community and permits a genuine test of whether organization predicts resistance, persistence, recovery, temporal invariability, or maintenance of function.

A community may contain several candidate protective mechanisms without having broad or accessible protection. Effects may reach only a small fraction of vulnerable recipients, be concentrated in one provider, share a common failure mode, or disappear under spatial separation. Conversely, distinct mechanisms may reach overlapping recipients through independent pathways. CEPT therefore distinguishes protective capacity—the presence and magnitude of demonstrated effects—from protective organization—their distribution, independence, accessibility, and concentration.

3.1. Operational representation

For a fixed observation window, let $E_{ij}(s)$ denote the signed change in an oriented recipient response attributable to provider *i* under disturbance *s* relative to the matched counterfactual. Recipient variables are oriented so that larger values always indicate improved performance; for example, recovery time and functional deviation are sign-reversed or otherwise transformed before comparison. Positive values therefore represent benefit, zero denotes no detectable effect, and negative values indicate increased vulnerability. The protective architecture is the positive subnetwork $P_{ij}(s) = \max[0, E_{ij}(s)]$; negative effects remain part of the broader signed effect matrix but are not protective edges.

In addition to the positive protective subnetwork, retain the signed effect matrix $E(s)$ as a companion representation. Baseline predictors should include preregistered summaries of detrimental effects, where estimable, so improvements attributed to $P(s)$ are not merely consequences of omitted antagonistic interactions. For prospective tests, calibration and outcome measurement must use separate experimental units or nonoverlapping time intervals.

The pairwise matrix $P(s)$ provides the first-order representation. Complementary provider sets that generate protection only in combination can be represented as composite nodes or higher-order terms. When temporal change is central, $P(s,t)$ denotes the longitudinal extension, in which nodes and protective links may change through succession, immigration, extinction, spatial rearrangement, phenotypic plasticity, horizontal gene transfer, or evolution (Renes et al. 2020; Canarini et al. 2021). This notation does not imply that stability requires preservation of a fixed network; protective effects may instead be maintained or reconstructed through different taxa or pathways.

For prospective tests, distinguish pre-perturbation protective capacity $P_{cap}(s)$, estimated in independent calibration assays under a standardized challenge, from effective protection $P_{eff}(s,t)$, measured

during a later focal disturbance. A prospective stability prediction uses $P_{\text{cap}}(s)$, or a prespecified early-time estimate that cannot incorporate the outcome interval. $P_{\text{eff}}(s,t)$ is subsequently measured to test whether provider mortality, inhibition, or spatial reorganization mediates deviation from the forecast. Inference from P_{eff} measured after the outcome is not a valid independent prediction. Protective coverage and dependency concentration are defined over positive P edges only; negative entries of $E(s)$ remain available as separately modeled antagonistic covariates, rather than being silently recast as protection.

3.2. Candidate architectural descriptors

Protective coverage is the fraction of vulnerable recipients, or of explicitly defined recipient-level functions, reached by at least one protective pathway. Coverage distinguishes a community in which protection is broadly accessible from one in which strong effects benefit only a small subset of members.

Protective redundancy is the number of independent pathways capable of protecting the same recipient or response. It differs from functional redundancy: two taxa may carry the same gene or perform the same process but fail together because they share spatial constraints, resource dependence, or disturbance sensitivity. Redundancy contributes to stability only when alternative pathways remain accessible and effective after selective loss.

Mechanistic complementarity describes the extent to which different protective processes address distinct components, phases, or consequences of a disturbance. Detoxification and nutrient provisioning, for example, may protect the same recipient through non-equivalent mechanisms. Complementarity may broaden protection but can also create serial dependency when one pathway requires another.

Spatial reach describes whether a protective effect can access its intended recipients across the relevant microhabitat. Cell proximity, diffusion, biofilm structure, and physicochemical gradients determine whether extracellular protection remains local or becomes community-wide (Flemming et al. 2016; Nadell et al. 2016).

Modularity describes the compartmentalization of protective pathways. Modules may limit the propagation of disturbance and preserve local effects, but excessive compartmentalization may prevent protection from reaching vulnerable recipients outside a module. Its predicted effect is therefore conditional rather than universally stabilizing (Stouffer and Bascompte 2011; Hernandez et al. 2021).

Dependency concentration quantifies how strongly total protective effect relies on a small number of providers or pathways. Concentration may generate efficient short-term protection but creates single points of failure when those providers are selectively affected. None of these descriptors is assumed to be universally monotonic: their meaning depends on the focal disturbance, response, spatial scale, and correlations among pathway failures.

Operational conventions and assumptions. Let V denote a preregistered set of recipient units and let $v_j \geq 0$ be prespecified recipient-importance weights with (Eq. 1). Estimate a signed provider effect $E_{ij}(s)$ using an independently replicated, disturbance-matched counterfactual and define (Eq. 2), where $r_j > 0$ is a fixed, prespecified recipient-response scale. All edge weights must use comparable response definitions and

observation windows; otherwise descriptors should be stratified by response. The potential protective weight w_{ij} is estimated in a standardized calibration environment. Independently measured accessibility $a_{ij} \in [0, 1]$ adjusts that potential to the target spatial configuration. Do not apply a_{ij} twice when E_{ij} was already measured under the actual spatial configuration. Define a biologically justified minimum-effect threshold τ before analysis and retain uncertainty intervals and detection limits rather than treating nonsignificant edges as confirmed zeros. These are proposed operational indices, not uniquely determined quantities.

$$\sum_{j \in V} v_j = 1 \quad (1)$$

$$w_{ij} = \max(0, E_{ij}(s)) / r_j \quad (2)$$

Coverage and concentration. Define (Eq. 3), which counts recipients reached by at least one individually sufficient validated pathway, without assuming that subthreshold effects add linearly. Where jointly sufficient provider sets are established experimentally, count each such set as a separate higher-order pathway; a sensitivity analysis may use the additive threshold criterion if additivity is justified. Define provider contributions q_i and concentration H (Eqs. 4–5) when $\sum_k q_k > 0$. This Herfindahl concentration index characterizes the distribution of measured provider contributions, not the probability of collapse. H depends on the number and definition of provider units; report provider count, effective provider number $1/H$, and results under relevant provider-set representations. If total contribution is zero, H is undefined, not zero. Shared resource dependence or simultaneous failure may make even a low- H system fragile; failure correlations therefore require separate analysis.

Provider protective contribution q_i in Eq. (4) serves as the default, response-specific weighted provider score for targeted-removal tests; it should not be described as a universal centrality measure. For vulnerability-sensitive prioritization, recipient weights v_j are prespecified, and a leave-one-provider-out protection-loss score $L_i = C(P) - C(P \text{ without provider } i)$ may be reported alongside q_i . L_i must be recalculated with the same thresholds and provider-set rules, and shared failure modes can make either score a poor proxy for actual collapse risk.

$$C = \sum_{j \in V} v_j \cdot \mathbb{1}[\max_i(a_{ij}w_{ij}) > \tau] \quad (3)$$

$$q_i = \sum_{j \in V} v_j a_{ij} w_{ij} \quad (4)$$

$$H = \sum_i [q_i / (\sum_k q_k)]^2 \quad (5)$$

Redundancy, mechanisms, and spatial structure. For recipient j , let A_j contain experimentally distinguishable protective pathways with accessible weights $b_k \geq 0$. A candidate effective-pathway index is R_j (Eq. 6), with $R_j = 0$ when no pathway is observed and R (Eq. 7). The prespecified $f_{kl} \in [0,1]$ represents failure-mode overlap, with $f_{kk} = 1$, and its matrix must be positive semidefinite (for example, an empirically estimated nonnegative failure-similarity kernel), so the denominator is mathematically coherent. The index equals two for two equal-weight independent pathways and one for two equal-weight fully

overlapping pathways; unequal independent pathways yield an effective number between one and two. Because overlap may vary across disturbances, report sensitivity intervals; unknown dependence must not be coded as independence. The expression D_m (Eq. 8), where h_{jm} denotes shares across prespecified mechanism classes, is a mechanism-diversity index, NOT a measure of causal complementarity. Mechanistic complementarity should instead be assessed by preregistered factorial provider-set contrasts, testing whether combined protection exceeds (synergy), equals (additivity), or falls below (antagonism) a justified null model on the chosen response scale. If these contrasts are unavailable, complementarity is not estimable. Spatial reach S (Eq. 9) is the fraction of standardized potential effect spatially accessible (undefined if the denominator is zero); it is meaningful only when w_{ij} is assessed without the same accessibility loss. Network modularity Q can be estimated from a specified directed, weighted null model and preregistered detection algorithm, but it is optional and should be recorded as not estimable in sparse or very small graphs. Composite provider sets must be retained as hyperedges or analyzed under explicit node-expansion sensitivity scenarios.

$$R_j = (\sum_k b_k)^2 / [\sum_k b_k^2 + 2\sum_{k<l} b_k b_l f_{kl}] \quad (6)$$

$$R = \sum_j v_j R_j \quad (7)$$

$$D_m = \sum_j v_j [1 - \sum_m h_{jm}^2] \quad (8)$$

$$S = [\sum_j v_j \sum_i a_{ij} w_{ij}] / [\sum_j v_j \sum_i w_{ij}] \quad (9)$$

Statistical identifiability and missing-edge uncertainty. Architectural descriptors are community-level predictors, not independent observations for each organism in the same microcosm. Replicated communities (or independently perturbed community states) are the sampling units for predictive validation; otherwise pseudoreplication inflates precision. Estimated edge effects, uncertain thresholds, and inferred links must be propagated through descriptor uncertainty by hierarchical resampling or posterior simulation. Cross-validation must hold out entire communities or experimental batches, not edges or technical replicates drawn from the same community. When predictors are deterministic or highly collinear in minimal networks (e.g., the one-provider case), coefficients for C , R , and H cannot be identified separately; report descriptive indices but do not fit the full multivariable model. An observed improvement in predictive performance does not by itself establish causal mediation; randomized rewiring and pathway rescue offer more direct tests.

Linking architecture to stability without circularity. Define the response Y from held-out community replicates or from an observation window that starts after calibration of the architecture. A baseline model is defined by (Eq. 10), where X includes prespecified compositional, functional, environmental, and conventional-network comparators. The primary architectural extension is given in Eq. (11). This parsimonious candidate should be chosen before analyzing outcomes; it is not a claim that all coefficients are nonzero or have universal signs. Spatial reach S and optional modularity Q and mechanism-diversity D_M can be tested as secondary predictors when independently measurable and when sample size supports their

estimation. Interactions (such as $C \times S$) and nonlinear modularity terms (Q^2) are exploratory unless independently motivated, prespecified, and adequately powered. Center or standardize predictors, check collinearity and sensitivity to threshold τ and recipient weights, and compare nested out-of-sample prediction, uncertainty, and calibration with identical data partitions. Mechanistic-diversity alone must not be interpreted as demonstrated complementarity. In causal rewiring studies, randomize the perturbation, measure the proposed mediators independently, and account for collateral ecological effects rather than attributing every outcome change to protective architecture.

$$g(E[Y]) = \alpha + \gamma^T X \quad (10)$$

$$g(E[Y]) = \alpha + \gamma^T X + \beta_c C + \beta_r R + \beta_h H \quad (11)$$

Worked illustration (hypothetical; not values reported by Morris et al. 2011). Consider three recipient units with equal $v_j = 1/3$ and two prospective helpers A (Alteromonas-like peroxide scavenger) and B (an invented alternative). Prespecify $\tau = 0.1$, $a_{ij} = 1$, and standardized potential effects $A \rightarrow P = 0.6$, $A \rightarrow R = 0.4$, and $B \rightarrow P = 0.4$; all other links have zero modeled effect for illustration only. The validated-pathway coverage is $C = 2/3$. Provider contributions are $q_A = 1/3$ and $q_B = 0.4/3$, yielding $H = (1/1.4)^2 + (0.4/1.4)^2 \approx 0.592$ and effective provider number $1/H \approx 1.69$. If failures of $A \rightarrow P$ and $B \rightarrow P$ are independent, $R_P = (0.6+0.4)^2/(0.6^2+0.4^2) \approx 1.923$, $R_R = 1$, $R_{\text{unprotected}} = 0$, and $R \approx 0.974$; under complete overlap of their failure modes, $R_P = 1$ and $R = 2/3$. Under the explicitly assumed complete accessibility, $S = 1$. If both effects are classified as peroxide detoxification, $D_M = 0$; this does NOT establish that mechanistic complementarity is absent, which requires factorial evidence. Q and causal complementarity are not estimable from this illustration. In the simpler validated Alteromonas→Prochlorococcus single-pathway illustration, $C = R = H = S = 1$ for the single recipient/provider, and Q remains not estimable. These are index calculations, not observed measurements or numeric stability forecasts; the latter require independently estimated regression coefficients and validation.

3.3. Mechanisms, accessibility, and failure

Cross-feeding, nutrient provisioning, environmental modification, extracellular matrices, catalytic detoxification, and spatial shielding are biological mechanisms through which protective effects may arise. They enter $P(s)$ only when their recipient-level benefit is demonstrated under the focal disturbance. A mechanism inventory alone is insufficient because it does not establish who receives the benefit, whether it remains accessible, or whether alternative pathways fail independently.

Protective effects may impose costs on providers and create opportunities for non-contributing organisms to exploit shared benefits. Competition, cheating, antagonism, diffusion loss, and spatial separation can therefore weaken or reverse protection. A pathway contributes to stability when its benefit remains accessible and maintainable under the disturbance; the same pathway can generate fragility when recipients depend on a single costly or vulnerable provider. CEPT consequently treats protection and dependency as possible outcomes of the same organization rather than as opposing categories.

Feasible reconstruction in species-rich communities. Exhaustive ordered-pair intervention scales approximately as $n(n-1)$ assays before replication and is not required for every CEPT application. A tiered protocol first identifies candidate provider–recipient effects through longitudinal perturbation, spatial profiling, metabolic models, and factorial or dilution-to-extinction/reassembly screens. Second, it validates a stratified subset of high-leverage positive, null, and negative edges using independent provider removal, inhibition, and activity-specific rescue under matched disturbances. Third, it estimates remaining edges with uncertainty-aware hierarchical or sparse inverse models, keeping inferred edges distinct from intervention-validated edges. Finally, it propagates edge-level uncertainty into C, R, M, S, Q, and H and compares predictions against held-out communities and disturbances. If causal identification or sampling coverage is inadequate, report a partial architecture and refrain from claims of fully measured community-wide protection.

Figure 3 depicts an ordered test rather than a retrospective description of a stable community. The disturbance, response and observation window are fixed first; protective effects are then independently characterized with counterfactuals, and architectural descriptors enter a comparison against established covariates. CEPT gains support only when intervention or held-out prediction establishes incremental value, with uncertainty and missing edges carried forward.




Figure 3. Operational test of CEPT. The disturbance, recipient response, scale, and observation window are specified before provider-recipient effects are measured against matched counterfactuals. Descriptors derived from $P(s)$ are then evaluated together with established covariates. CEPT requires incremental causal or predictive value for the specified stability response.

4. Relationship to Existing Ecological Frameworks

CEPT does not replace established ecological frameworks; it addresses a distinct organizational question. Resistance, resilience, persistence, functional deviation, and recovery time are distinct dimensions or operational measures of ecological stability, whose interpretation depends on the response variable, disturbance regime, and observation window (Grimm and Wissel 1997; Ives and Carpenter 2007; Shade et al. 2012). CEPT proposes one possible determinant of variation in those responses: the organization of demonstrated protective effects. It therefore does not redefine stability but links a specified stability dimension to a separately measured architecture. Importantly, CEPT does not assume that all stability dimensions respond similarly to protective organization. An architecture that increases resistance during disturbance may not necessarily accelerate subsequent recovery, and the direction and magnitude of these relationships must be established empirically.

Functional redundancy concerns overlap in ecological roles or potential functions (Yachi and Loreau 1999; Louca et al. 2018). Protective redundancy is narrower: alternative provider-recipient pathways must compensate for one another under the focal disturbance. Functional duplicates may not be protective, and protective alternatives may rely on different mechanisms. Moreover, species performing similar ecological functions may differ in their sensitivity to environmental disturbances, making functional overlap insufficient to guarantee compensatory protection (Elmqvist et al. 2003; Loreau and de Mazancourt 2013). Protective redundancy therefore requires evidence that alternative pathways remain effective and accessible when another pathway is disrupted, rather than merely demonstrating overlap in functional traits or metabolic capabilities. CEPT is supported only when this distinction improves explanation or prediction beyond functional redundancy, response diversity, and established biodiversity–stability mechanisms.

Public-goods, cooperation, facilitation, and social-evolution frameworks explain the production, sharing, costs, exploitation, and evolutionary maintenance of beneficial traits (West et al. 2007; Pande and Kost 2017; Zachar and Boza 2022). CEPT uses evidence generated by those frameworks but neither requires cooperation nor assumes mutual benefit. A provider may incur a cost, receive a benefit, or be unaffected; classification depends on the causal effect on the recipient under disturbance. Furthermore, beneficial interactions may depend on ecological relationships that maintain provider populations or their metabolic activities before a disturbance occurs. Such relationships can influence the availability of protective capacity without themselves constituting demonstrated protective effects (Goldford et al. 2018; Calfee et al. 2022). CEPT therefore distinguishes the ecological maintenance of potential providers from the disturbance-specific protection those providers subsequently deliver. Similarly, colonization-resistance theory describes the suppression of invaders, whereas CEPT asks whether the distribution of resident protective effects explains variation in that response beyond community richness, resource competition, and functional composition (van Elsas et al. 2012; Spragge et al. 2023).

Ecological network theory supplies tools for representing direction, weights, centrality, modularity, and dependency (Faust and Raes 2012; Coyte et al. 2015). Classical theoretical approaches have further

demonstrated that the stability of ecological systems depends on interaction structure, interaction strength, and the balance among different interaction types (Allesina and Tang 2012). CEPT restricts the focal network to effects supported by provider-recipient counterfactuals under a defined disturbance and response. Co-occurrence edges do not by themselves establish direction, mechanism, or protection (Weiss et al. 2016; Carr et al. 2019). The distinctive claim is thus not that interactions matter, but that the independently characterized protective subnetwork may contain additional explanatory information not captured by composition, functional inventories, or undirected associations. This claim must be evaluated against established network-based predictors, including conventional measures of centrality, interaction strength, and modularity, rather than assuming that protective architecture necessarily offers superior explanatory power.

Table 2. Comparison of CEPT with neighboring ecological frameworks. Differences indicate proposed operational contrasts, not empirical superiority.

Framework	Main unit / evidence		Disturbance-specific?	Directional counterfactual?	/	CEPT-specific test beyond comparator		
Facilitative networks (Fujita et al. 2023)	Inferred community interactions	directed	Context dependent	Directed; validation inherent	causal not	Causally recipient-level edges forecast a response beyond inferred network structure	validated protective specified	
Control / driver species (Angulo et al. 2019)	Structural accessibility controlling sets	and	Not necessarily tied to focal stress	Directional network; recipient benefit not required	not	Stress-specific effects controllability / driver metrics	provider outperform	
Keystone and centrality (Power et al. 1996; Banerjee et al. 2018)	Disproportionate taxon influence		Not necessarily	Removal possible; edges not required	tests protective	q_i and replaceability matched-loss predictions	pathway improve	
Functional insurance / response diversity (Yachi and Loreau 1999;	Functional and responses	overlap differential	Often environmental variation	Recipient-specific counterfactual required	not	Accessible redundancy prediction	recipient-directed adds held-out	

Elmqvist et al. 2003)

Generic	network	Interaction Jacobian /	Model-dependent	Signed networks; not	Protective	P(s) contributes
stability	(Allesina	local stability		necessarily protective	beyond conventional network	
and Tang 2012)				intervention	predictors	

Relationship to the closest theoretical alternatives. Fujita et al. (2023) reconstructed potential facilitative interaction networks using metabolic inference across a longitudinal microbial experiment, identifying network features, including positive-feedback structures, associated with subsequent compositional transitions. Their work provides an important precedent for linking facilitative network organization to microbial community dynamics, rather than evidence that the use of directed or weighted networks is itself novel. CEPT adds the narrower requirement for a matched, disturbance-specific causal benefit to an identified recipient, and tests whether the organization of those effects explains a preregistered response beyond existing network predictions. Unlike networks inferred primarily from metabolic dependencies, CEPT distinguishes experimentally validated protective edges from hypothesized or indirectly inferred interactions, particularly when evaluating the causal contribution of architecture to stability.

Angulo et al. (2019) developed a control-theoretic framework based on structural accessibility, identifying sets of driver species whose manipulation may allow microbial communities to be directed toward desired states. This approach addresses community controllability, whereas CEPT addresses the organization of protective effects contributing to a specified response under disturbance. A CEPT provider instead confers demonstrated stress-conditional recipient protection and need not be a topological driver. Conversely, a driver species identified through structural accessibility need not provide protective benefits to a particular recipient. The two approaches may identify overlapping influential taxa, but they define ecological importance through different operational criteria and make distinct predictions.

Keystone and node-removal theories predict outsized consequences of removing influential taxa (Power et al. 1996; Banerjee et al. 2018); CEPT additionally predicts response-specific rescue by substituting the protective activity, even while taxon identity and conventional centrality differ. The distinction becomes particularly informative when two communities share similar abundance distributions and conventional network topology but differ in the accessibility, replaceability, or failure dependence of experimentally identified protective pathways. Under these conditions, CEPT predicts that protective architecture may explain differences in resistance or recovery that are not captured by conventional keystone descriptors alone.

Insurance and response-diversity theories explain buffering through functional overlap and different environmental responses (Yachi and Loreau 1999; Elmqvist et al. 2003; Loreau and de Mazancourt 2013); CEPT asks whether the independently validated provider–recipient pathways and their spatial accessibility add information after that buffering is measured. Accordingly, two communities with equivalent functional

inventories and comparable response diversity may still differ in their ability to protect vulnerable recipients if protective pathways are spatially inaccessible, concentrated in a small number of providers, or subject to correlated failure. Whether such differences contribute independent predictive information constitutes a direct empirical test of CEPT rather than an established ecological result.

General network stability criteria (Allesina and Tang 2012) address interaction structure at a different level of generality. These criteria provide mathematical insights into how interaction types and network organization influence local dynamical stability, whereas CEPT focuses on disturbance-conditional, recipient-specific protective effects and their relationship to explicitly defined resistance, persistence, or recovery outcomes. The approaches are complementary, but CEPT must demonstrate incremental predictive or causal value beyond suitable network-stability comparators. The proposed novelty is therefore incremental and falsifiable, not a claim to have discovered beneficial networks or redundancy. Its central contribution lies in proposing a common operational framework for distinguishing protective effects from potential mechanisms, representing their distribution among providers and recipients, and testing whether that organization improves explanations of microbial community stability beyond established ecological alternatives.

5. Testable Predictions and Falsifiability

Ecological theories gain explanatory power when they generate discriminating predictions and specify observations that would challenge their central claims. CEPT therefore focuses on comparisons in which the disturbance and response are fixed, established predictors are controlled, and protective architecture is manipulated or characterized independently. Four predictions isolate the proposed architecture effect. Importantly, the proposed tests do not require competing ecological frameworks to predict the complete absence of an effect. Instead, CEPT must demonstrate an additional explanatory or predictive contribution after relevant compositional, functional, and network-based mechanisms have been accounted for. Such comparisons are especially important because previous research has already established relationships between ecological interaction structure, stability, and community dynamics (Allesina and Tang 2012; Fujita et al. 2023). The corresponding discriminating tests and failure criteria are summarized in Table 3.

5.1. Matched-community rewiring

Discriminating contrast (P1). With taxon identities, functional inventory, abiotic context and global degree or conventional network metrics held as constant as feasible, CEPT predicts that selectively redirecting experimentally verified protective edges toward previously unprotected recipients changes their response to the specified challenge. Composition-only and functional-inventory models predict no difference if those inputs remain fixed. Generic interaction-network models must also be included as comparators; a difference explained equally well by their topology is not distinctive CEPT support.

Communities matched as closely as feasible in richness, composition, abundance, functional potential, and abiotic context will differ in resistance, persistence, or recovery when provider-recipient protective effects are experimentally reorganized. Spatial separation, selective pathway inhibition, or substitution of functionally

similar providers can alter protection while minimizing changes to the broader community. CEPT is supported if independently measured coverage, pathway independence, or dependency concentration predicts the resulting response beyond the matched controls.

A particularly informative experiment would redistribute a comparable total magnitude of validated protective effects among recipients while preserving, as far as experimentally feasible, community composition, provider abundance, and conventional network descriptors. Under this design, communities with broader coverage of vulnerable recipients may exhibit greater recipient-level resistance than communities in which equivalent protective capacity is concentrated in a subset of recipients. However, this prediction is conditional on the identity and vulnerability of the recipients and should not be interpreted as a universal positive effect of coverage.

Because spatial rearrangement or pathway inhibition may alter non-protective interactions simultaneously, accompanying changes in competition, resource exchange, and environmental exposure should be measured. Evidence for CEPT requires that independently quantified protective architecture explain the response beyond these collateral effects. A null effect of rewiring alone would not reject CEPT unless the manipulation demonstrably changed the targeted architectural descriptor and the experiment had sufficient precision to detect the prespecified effect.

5.2. Targeted provider loss

Discriminating contrast (P2). In the *Alteromonas*–*Prochlorococcus* H₂O₂ system, compare helper removal with matched non-helper removal, catalase-pathway inhibition, and rescue with peroxide removal supplied independently of *Alteromonas* (for example, a separately controlled detoxification intervention). Measure prospective H₂O₂ scavenging and *Prochlorococcus* response under the same challenge. CEPT specifically predicts that restored protective activity can rescue the recipient without restoring helper identity, while a generic taxon-identity or abundance-only prediction fails. This proposed factorial extension has not been performed here, and a generic keystone framework that includes the same mechanism could predict the result; incremental explanatory value must therefore be demonstrated rather than presumed.

Removing a provider that supplies a large and weakly redundant share of stress-specific protection will reduce recipient persistence or ecosystem function more than removing a taxon matched for abundance, growth rate, or metabolic contribution but with replaceable protective effects. Rescue by restoring the missing activity through a taxonomically different provider would further separate protective centrality from taxonomic identity or generic keystone status.

The crucial discriminatory variable is the independently estimated contribution of each provider to recipient protection, rather than provider removal alone. Keystone-species theory already predicts that the loss of disproportionately influential organisms can cause substantial community changes (Power et al. 1996; Banerjee et al. 2018). CEPT therefore gains distinctive support only if stress-specific protective contribution, including dependency concentration and the availability of alternative pathways, improves prediction of recipient responses beyond abundance, conventional centrality, and other established measures of ecological

importance.

In the proposed *Alteromonas*–*Prochlorococcus* experiment, an activity-specific rescue would support the causal role of peroxide detoxification but would not independently validate the entire CEPT framework. A stronger architectural test would compare communities containing one dominant peroxide-scavenging provider with communities containing multiple independently functioning providers, while controlling their aggregate detoxification capacity and the relevant environmental conditions. Under selective helper loss, CEPT predicts that the resulting response should depend on the independently characterized distribution, accessibility, and failure dependence of the alternative protective pathways.

5.3. Protective redundancy beyond functional redundancy

Discriminating contrast (P3). Match function inventories and provider-specific response diversity while experimentally altering co-failure dependence or access to the recipient. CEPT predicts differential recipient performance if the same functions remain present but recipient-directed compensation differs. If a response-diversity model already captures the change, CEPT has no demonstrated additional contribution. Overlap-adjusted *R* and independently measured accessibility *S* must improve the preregistered comparison.

Communities with comparable functional redundancy will respond differently when equivalent protective functions are supplied through independent versus correlated-failure pathways. Selective loss of one provider should be compensated only when alternative protection remains spatially accessible and effective under the same disturbance. Gene-level or taxonomic redundancy without compensatory recipient protection does not satisfy this prediction.

This prediction builds on, rather than replaces, the insurance hypothesis and response-diversity theory, which already recognize the stabilizing potential of functional overlap and differences in species responses to environmental change (Yachi and Loreau 1999; Elmqvist et al. 2003; Loreau and de Mazancourt 2013). CEPT introduces the additional requirement that putatively redundant functions produce demonstrable benefits for defined recipients and remain accessible when alternative protective pathways fail.

An experimental comparison could therefore combine selective provider inhibition with spatial rearrangement or manipulation of shared failure modes. Communities matched in functional inventory and provider response diversity could still differ in recipient persistence if protection does not reach vulnerable recipients or if nominally alternative pathways depend on the same limiting resource. However, if established response-diversity or spatial-interaction models already explain these differences, CEPT cannot claim an independent architectural contribution.

The prediction is supported only to the extent that independently estimated recipient-directed protective redundancy and accessibility improve causal explanation or out-of-sample prediction beyond the relevant competing models.

5.4. Reconstruction during recovery

Discriminating contrast (P4). Select communities matched for taxonomic similarity, conventional functional recovery indicators and broad interaction-network recovery, but differing in prospectively measured

restoration of causal protective edges. CEPT predicts that targeted restoration of those edges will change later resistance or functional recovery. Resilience theory already allows taxonomic–functional decoupling; therefore taxonomic mismatch alone does not distinguish CEPT. Only an architectural intervention or incremental out-of-sample prediction after these alternatives are modeled does so.

Following disturbance, recovery of richness or taxonomic similarity will not reliably predict functional recovery when protective pathways remain absent, inaccessible, or concentrated. CEPT predicts that longitudinal reconstruction of provider-recipient effects will improve prediction of recovery time and post-disturbance resistance beyond taxonomic return alone. Experimental restoration of a missing protective pathway may improve subsequent recovery or resistance when that pathway limits the specified response, without requiring re-establishment of the original composition.

This prediction is consistent with the established distinction between microbial community composition and function, including evidence that functional processes may persist despite taxonomic turnover (Louca et al. 2016; Louca et al. 2018). Its distinctive CEPT component is therefore not the general observation that taxonomic and functional recovery may be decoupled, but the prediction that independently measured reconstruction of recipient-directed protective pathways provides additional explanatory information.

A prospective test would characterize protective architecture at predefined recovery stages and evaluate its ability to predict subsequent functional performance or resistance to a standardized secondary challenge. Where feasible, selective restoration of a missing protective activity should be compared with matched interventions that restore biomass, taxonomic representation, or non-protective metabolic functions without restoring the focal protective pathway.

Such experiments would distinguish recovery associated with the reappearance of taxa from recovery associated with re-establishment of causal protection. Nevertheless, restoration of a protective pathway need not accelerate every recovery response, particularly when other ecological constraints remain limiting.

5.5. Criteria for evaluating and challenging CEPT

The four predictions should be evaluated using prospectively characterized protective architecture, prespecified disturbances and response variables, and explicit comparisons with competing ecological models. A positive result is informative only when the experimental design or predictive analysis distinguishes the proposed architectural contribution from established explanations.

CEPT would be weakened if protective effects could not be measured reproducibly, if descriptors derived from P(s) added no causal or out-of-sample predictive value after established covariates were controlled, or if rewiring, provider removal, pathway rescue, and architectural reconstruction produced no effects consistent with the corresponding predictions. These criteria must be interpreted in relation to measurement uncertainty, manipulation fidelity, statistical power, and the ecological domain for which the theory is being evaluated. Failure to detect an architectural effect in a poorly resolved network is not equivalent to evidence that the effect is absent. Conversely, repeated high-precision null results in systems where the relevant protective pathways have been independently established would substantially weaken the theory's proposed generality.

No single positive association is decisive. Replication across disturbances and systems, comparison with simpler models, and preregistered or out-of-sample tests are required to determine whether protective architecture is a general ecological principle or a context-limited description (Popper 1959).

Comparator requirement for P1–P4. Each prediction is evaluated against at least one predeclared competing model (composition/functional inventory; conventional network-centrality or keystone measures; insurance/response diversity; and established resistance–resilience metrics, respectively). Support requires incremental effects or prospective prediction beyond the relevant comparator, not merely consistency with CEPT. Where competing frameworks include the same protective mechanism or network information, the relevant test concerns the additional value of the CEPT-specific operationalization, rather than an assumed contradiction between theoretical perspectives.

Table 3. Principal discriminating predictions generated by CEPT.

Prediction	Discriminating test	Competing explanation and CEPT-specific increment	Result that challenges CEPT
P1 — Matched-community rewiring	Reorganize protective edges compositionally functionally communities while controlling conventional network descriptors.	validated in functional-inventory models and similar outcomes when those variables remain fixed. CEPT must add information beyond generic network topology.	and Verified architectural changes yield no detectable effect on the prespecified response and no incremental predictive value under adequately powered tests.
P2 — Targeted provider loss	Remove high-contribution providers, compare with matched removals, and restore the protective activity of original provider identity.	Keystone and centrality theories already predict effects of influential-taxon loss. CEPT must improve prediction using independently measured stress-specific protective contribution and pathway replaceability.	Protective contribution and dependency descriptors fail to improve prediction beyond keystone measures; targeted functional rescue provides no benefit when the proposed mechanism is verified as limiting.
P3 — Protective redundancy	Compare communities matched for functional redundancy and response diversity while varying pathway accessibility or co-failure dependence.	Insurance and response-diversity theory explain buffering through functional overlap and differentiated responses. CEPT must add information beyond established accessibility.	Validated redundancy and accessibility descriptors provide no additional causal or predictive information beyond established models.
P4 — Architectural reconstruction	Measure protective pathways prospectively during recovery and selectively restore missing protective activities.	Established resilience theory accommodates taxonomic–functional decoupling. CEPT must show additional information from reconstruction of causal protective pathways.	Independently measured pathway reconstruction adds no predictive value and targeted restoration produces no response when protection is experimentally established as a limiting mechanism.

6. Implications and Conclusions

An operational risk endpoint would be the preregistered probability of recipient or functional loss after a specified provider-failure intervention. Whether dependency concentration H or the leave-one-provider-out score L_i improves prediction of that endpoint beyond conventional network and environmental indicators is an empirical test, not an assumed property (Chen et al. 2013).

CEPT provides a basis for experiments that manipulate protective organization rather than merely cataloguing taxa or genes. Synthetic communities, microfluidic spatial manipulations, selective provider removal, pathway inhibition and rescue, and longitudinal reconstruction can test whether architecture contributes to resistance and recovery (Nadell et al. 2016; Jiang et al. 2023; Mehlferber et al. 2024). Defined synthetic microbial communities offer particularly useful experimental systems because their composition and environmental conditions can be controlled while specific interactions are manipulated and their consequences evaluated. Microfluidic approaches additionally permit spatial constraints and interaction accessibility to be investigated under controlled conditions. In engineered consortia, the framework suggests selecting not only metabolic functions but also independent, accessible protective pathways while avoiding unbuffered reliance on a single provider. Such applications remain prospective, but they are consistent with broader efforts to improve consortium stability through the deliberate management of microbial interactions (Jiang et al. 2023).

A practical validation roadmap is to reconstruct experimentally supported partial architectures in small synthetic communities, calibrate the proposed descriptors and intervention contrasts across independent replicates, and then test transferability to other disturbances and more species-rich natural and engineered assemblages.

The theory also extends to biological invasion when invasion is defined as the focal biotic disturbance. Resident diversity can restrict invader establishment through resource use, niche occupation, antagonism, environmental modification, and collective nutrient blocking (van Elsas et al. 2012; Spragge et al. 2023). CEPT does not treat richness as protection by itself; it predicts that colonization resistance will depend on which residents protect which recipients or functions, through which mechanisms, and whether those effects remain accessible during invasion. Importantly, the protective contribution of resident communities should be distinguished from the general suppression of invading organisms. Within CEPT, the relevant causal unit may be a resident provider, a recipient-specific protective interaction, or an experimentally demonstrated provider-set effect. The extent to which these protective arrangements improve prediction beyond diversity and established invasion-resistance mechanisms remains to be determined.

Protective architectures are historically contingent. Recurrent disturbance, succession, and immigration may add complementary pathways or fill gaps in coverage, but they can also remove providers, concentrate dependency, favor exploiters, or improve responses to familiar stresses while reducing performance under novel ones (Renes et al. 2020; Canarini et al. 2021). Historical contingency and priority effects may further influence which organisms establish and which ecological interactions develop (Fukami 2015). As proposed in the preceding sections, ecological interactions that sustain potential protective providers before disturbance

may influence the subsequent availability of protection without themselves constituting validated protective edges. Community age or diversity therefore creates opportunities for architectural development but does not guarantee greater protection. The relevant empirical question is how history changes $P(s,t)$, including whether protective capacity is maintained, lost, or reconstructed through changes in provider populations and their supporting ecological relationships.

Microbial community stability is multidimensional and cannot be inferred from taxonomic composition or mechanism inventories alone. CEPT proposes an additional explanatory level: the disturbance- and response-specific organization of protective effects through which community members alter one another's persistence or functional performance. Coverage, protective redundancy, mechanistic complementarity, spatial reach, modularity, and dependency concentration provide initial descriptors, but their effects remain hypotheses rather than universal rules. Moreover, these descriptors differ in their operational requirements. Coverage, dependency concentration, and suitably characterized redundancy may be quantified using independently estimated protective effects, whereas causal mechanistic complementarity requires appropriate combined-provider experiments. Spatial reach and modularity require additional spatial or network information and may not be identifiable in every experimental system. These descriptors could eventually complement ecological-risk and ecological-security assessment by identifying vulnerable recipients, single points of protective failure, and pathways whose loss is most likely to propagate functional decline under disturbance. Such applications would require validation against established ecological-risk indicators and should not yet be interpreted as operational assessment tools.

The current evidential scope of CEPT remains limited. The studies summarized in Table 1 provide experimental precedents for stress-conditional protection, facilitation, and collective provider-set effects, but they do not establish that the full protective architecture or all six descriptors predict community stability across ecosystems. The numerical illustration presented in Section 3 is hypothetical, and the proposed mathematical indices require independent empirical calibration. Consequently, CEPT should presently be regarded as a falsifiable theoretical and methodological framework rather than a demonstrated general law of microbial community stability. Adequately replicated perturbation experiments are required to estimate architectural effects, characterize pathway failure dependence, evaluate uncertainty, and determine whether the proposed descriptors improve causal or out-of-sample prediction beyond established compositional, functional, and network-based models. Where network-wide identification is infeasible, conclusions should be restricted to experimentally validated or explicitly uncertainty-qualified partial architectures. Generalization to species-rich soil, freshwater, wastewater, and other natural or engineered communities remains an open empirical question.

The central significance of CEPT therefore lies not in proposing previously unknown protective mechanisms, but in asking whether their causal, disturbance-specific organization constitutes an independently informative property of microbial communities. By distinguishing protective potential from demonstrated protection, identifying the distribution of benefits among providers and recipients, and examining how

redundancy, accessibility, and dependency influence specified stability responses, CEPT offers a framework for translating local microbial interactions into testable community-level hypotheses. Its broader ecological relevance will depend on whether protective architecture can be measured reproducibly and shown to contribute explanatory or predictive value beyond established alternatives. Comparative experiments across disturbances, community types, and organizational scales can now determine whether collective ecological protection represents a generalizable organizing principle or a context-dependent feature of particular microbial systems.

Author Contributions

All authors contributed equally to the conception and design of the study, development of the theoretical framework, interpretation of ecological concepts, manuscript writing, figure design, critical revision of the manuscript, and approval of the final version.

Acknowledgements

The authors gratefully acknowledge Prof. Juan Nogales for his valuable discussions, scientific support, and for providing an inspiring academic environment that contributed to the development of the ideas presented in this work. The authors also thank the Department of Geology and Geophysics at Universidade Federal Fluminense (UFF), Brazil, for granting the sabbatical leave that enabled the development of this research.

References

1. Allesina S, Tang S (2012) Stability criteria for complex ecosystems. *Nature* 483:205–208. <https://doi.org/10.1038/nature10832>
2. Allison SD, Martiny JBH (2008) Resistance, resilience, and redundancy in microbial communities. *Proc Natl Acad Sci USA* 105:11512–11519. <https://doi.org/10.1073/pnas.0801925105>
3. Angulo MT, Moog CH, Liu YY (2019) A theoretical framework for controlling complex microbial communities. *Nat Commun* 10:1045. <https://doi.org/10.1038/s41467-019-08890-y>
4. Banerjee S, Schlaeppi K, van der Heijden MGA (2018) Keystone taxa as drivers of microbiome structure and functioning. *Nat Rev Microbiol* 16:567–576. <https://doi.org/10.1038/s41579-018-0024-1>
5. Brooker RW, Maestre FT, Callaway RM, Lortie CL, Cavieres LA, Kunstler G, Liancourt P, Tielbörger K, Travis MJJ, Anthelme F, Armas C, Coll L, Corcket E, Delzon S, Forey E, Kikvidze Z, Olofsson J, Pugnaire F, Quiroz CL, Saccone P, Schiffrers K, Seifan M, Touzard B, Michalet R (2008) Facilitation in plant communities: the past, the present, and the future. *J Ecol* 96:18–34. <https://doi.org/10.1111/j.1365-2745.2007.01295.x>
6. Calfee BC, Glasgo LD, Zinser ER (2022) *Prochlorococcus* exudate stimulates heterotrophic bacterial competition with rival phytoplankton for available nitrogen. *mBio* 13:e02571-21. <https://doi.org/10.1128/mbio.02571-21>
7. Canarini A, Schmidt H, Fuchslueger L, Martin V, Herbold CW, Zezula D, Gündler P, Hasibeder R, Jecmenica M, Bahn M, Richter A (2021) Ecological memory of recurrent drought modifies soil processes via changes in soil microbial community. *Nat Commun* 12:5308. <https://doi.org/10.1038/s41467-021-25675-4>
8. Carr A, Diener C, Baliga NS, Gibbons SM (2019). Use and abuse of correlation analyses in microbial ecology. *The ISME Journal* 13:2647–2655. <https://doi.org/10.1038/s41396-019-0459-z>
9. Chen S, Chen B, Fath BD (2013) Ecological risk assessment on the system scale: A review of state-of-the-art

- models and future perspectives. *Ecol Model* 250:25–33. <https://doi.org/10.1016/j.ecolmodel.2012.10.015>
10. Coyte KZ, Schluter J, Foster KR (2015) The ecology of the microbiome: networks, competition, and stability. *Science* 350:663–666. <https://doi.org/10.1126/science.aad2602>
11. de Vos MGJ, Zagorski M, McNally A, Bollenbach T (2017) Interaction networks, ecological stability, and collective antibiotic tolerance in polymicrobial infections. *Proc Natl Acad Sci USA* 114:10666–10671. <https://doi.org/10.1073/pnas.1713372114>
12. Elmqvist T, Folke C, Nyström M, Peterson G, Bengtsson J, Walker B, Norberg J (2003) Response diversity, ecosystem change, and resilience. *Front Ecol Environ* 1:488–494. [https://doi.org/10.1890/1540-9295\(2003\)001\[0488:RDECAR\]2.0.CO;2](https://doi.org/10.1890/1540-9295(2003)001[0488:RDECAR]2.0.CO;2)
13. Faust K, Raes J (2012) Microbial interactions: from networks to models. *Nat Rev Microbiol* 10:538–550. <https://doi.org/10.1038/nrmicro2832>
14. Flemming HC, Wingender J, Szewzyk U, Steinberg P, Rice SA, Kjelleberg S (2016) Biofilms: an emergent form of bacterial life. *Nat Rev Microbiol* 14:563–575. <https://doi.org/10.1038/nrmicro.2016.94>
15. Fujita H, Ushio M, Suzuki K, Abe MS, Yamamichi M, Okazaki Y, Canarini A, Hayashi I, Fukushima K, Fukuda S, Kiers ET, Toju H (2023) Facilitative interaction networks in experimental microbial community dynamics. *Front Microbiol* 14:1153952. <https://doi.org/10.3389/fmicb.2023.1153952>
16. Fukami T (2015). Historical contingency in community assembly: integrating niches, species pools, and priority effects. *Annual Review of Ecology, Evolution, and Systematics* 46:1–23. <https://doi.org/10.1146/annurev-ecolsys-110411-160340>
17. Gao C, Xu L, Montoya L, Madera M, Hollingsworth J, Chen L, Purdom E, Singan V, Vogel J, Hutmacher RB, Dahlberg JA, Coleman-Derr D, Lemaux PG, Taylor JW (2022) Co-occurrence networks reveal more complexity than community composition in resistance and resilience of microbial communities. *Nat Commun* 13:3867. <https://doi.org/10.1038/s41467-022-31343-y>
18. Goldford JE, Lu N, Bajić D, Estrela S, Tikhonov M, Sanchez-Gorostiaga A, Segrè D, Mehta P, Sanchez A (2018). Emergent simplicity in microbial community assembly. *Science* 361:469–474. DOI: <https://doi.org/10.1126/science.aat1168>
19. Granato ET, Smith WPJ, Foster KR (2023) Collective protection against the type VI secretion system in bacteria. *ISME J* 17:1052–1062. <https://doi.org/10.1038/s41396-023-01401-4>
20. Grilli J, Barabás G, Michalska-Smith MJ, Allesina S (2017) Higher-order interactions stabilize dynamics in competitive network models. *Nature* 548:210–213. <https://doi.org/10.1038/nature23273>
21. Grimm V, Wissel C (1997) Babel, or the ecological stability discussions: an inventory and analysis of terminology and a guide for avoiding confusion. *Oecologia* 109:323–334. <https://doi.org/10.1007/s004420050090>
22. Hernandez DJ, David AS, Menges ES, Searcy CA, Afkhami ME (2021) Environmental stress destabilizes microbial networks. *ISME J* 15:1722–1734. <https://doi.org/10.1038/s41396-020-00882-x>
23. Ives AR, Carpenter SR (2007) Stability and diversity of ecosystems. *Science* 317:58–62. <https://doi.org/10.1126/science.1133258>
24. Jiang Y, Wu R, Zhang W, Xin F, Jiang M (2023) Construction of stable microbial consortia for effective biochemical synthesis. *Trends Biotechnol* 41:1430–1441. <https://doi.org/10.1016/j.tibtech.2023.05.008>
25. Kajihara KT, Hynson NA (2024) Networks as tools for defining emergent properties of microbiomes and their stability. *Microbiome* 12:184. <https://doi.org/10.1186/s40168-024-01868-z>
26. Kelsic ED, Zhao J, Vetsigian K, Kishony R (2015) Counteraction of antibiotic production and degradation stabilizes microbial communities. *Nature* 521:516–519. <https://doi.org/10.1038/nature14485>
27. Lee KWK, Periasamy S, Mukherjee M, Xie C, Kjelleberg S, Rice SA (2014) Biofilm development and enhanced stress resistance of a model, mixed-species community biofilm. *ISME J* 8:894–907.

<https://doi.org/10.1038/ismej.2013.194>

28. Levine JM, Bascompte J, Adler PB, Allesina S (2017) Beyond pairwise mechanisms of species coexistence in complex communities. *Nature* 546:56–64. <https://doi.org/10.1038/nature22898>
29. Loreau M (2010) *From Populations to Ecosystems: Theoretical Foundations for a New Ecological Synthesis*. Princeton University Press, Princeton
30. Loreau M, de Mazancourt C (2013). Biodiversity and ecosystem stability: a synthesis of underlying mechanisms. *Ecology Letters* 16(Suppl. 1):106–115. DOI: 10.1111/ele.12073
31. Louca S, Parfrey LW, Doebeli M (2016) Decoupling function and taxonomy in the global ocean microbiome. *Science* 353:1272–1277. <https://doi.org/10.1126/science.aaf4507>
32. Louca S, Polz MF, Mazel F, Albright MB, Huber JA, O'Connor MI, Ackermann M, Hahn AS, Srivastava DS, Crowe SA, Doebeli M, Parfrey LW (2018) Function and functional redundancy in microbial systems. *Nat Ecol Evol* 2:936–943. <https://doi.org/10.1038/s41559-018-0519-1>
33. Maestre FT, Callaway RM, Valladares F, Lortie CJ (2009) Refining the stress-gradient hypothesis for competition and facilitation in plant communities. *J Ecol* 97:199–205. <https://doi.org/10.1111/j.1365-2745.2008.01476.x>
34. Mehlferber EC, Arnault G, Joshi B, Partida-Martinez LP, Patras KA, Simonin M, Koskella B (2024) A cross-systems primer for synthetic microbial communities. *Nat Microbiol* 9:2765–2773. <https://doi.org/10.1038/s41564-024-01827-2>
35. Morris JJ, Johnson ZI, Szul MJ, Keller M, Zinser ER (2011) Dependence of the cyanobacterium *Prochlorococcus* on hydrogen peroxide scavenging microbes for growth at the ocean's surface. *PLoS ONE* 6:e16805. <https://doi.org/10.1371/journal.pone.0016805>
36. Nadell CD, Drescher K, Foster KR (2016) Spatial structure, cooperation and competition in biofilms. *Nat Rev Microbiol* 14:589–600. <https://doi.org/10.1038/nrmicro.2016.84>
37. Oprei A, Schreckinger J, Kholiavko T, Frossard A, Mutz M, Risse-Buhl U (2023) Long-term functional recovery and associated microbial community structure after sediment drying and bedform migration. *Front Ecol Evol* 11:1112519. <https://doi.org/10.3389/fevo.2023.1112519>
38. Pande S, Kost C (2017) Bacterial unculturability and the formation of intercellular metabolic networks. *Trends Microbiol* 25:349–361. <https://doi.org/10.1016/j.tim.2017.02.015>
39. Philippot L, Griffiths BS, Langenheder S (2021) Microbial community resilience across ecosystems and multiple disturbances. *Microbiol Mol Biol Rev* 85:e00026-20. <https://doi.org/10.1128/MMBR.00026-20>
40. Piccardi P, Vessman B, Mitri S (2019) Toxicity drives facilitation between 4 bacterial species. *Proc Natl Acad Sci USA* 116:15979–15984. <https://doi.org/10.1073/pnas.1906172116>
41. Popper KR (1959) *The Logic of Scientific Discovery*. Hutchinson, London
42. Power ME, Tilman D, Estes JA, Menge BA, Bond WJ, Mills LS, Daily G, Castilla JC, Lubchenco J, Paine RT (1996) Challenges in the quest for keystones. *BioScience* 46:609–620. <https://doi.org/10.2307/1312990>
43. Renes SE, Sjöstedt J, Fetzer I, Langenheder S (2020) Disturbance history can increase functional stability in the face of both repeated disturbances of the same type and novel disturbances. *Sci Rep* 10:11333. <https://doi.org/10.1038/s41598-020-68104-0>
44. Sanchez A (2019). Defining higher-order interactions in synthetic ecology: lessons from physics and quantitative genetics. *Cell Systems* 9:519–520. <https://doi.org/10.1016/j.cels.2019.11.009>
45. Sanchez-Gorostiaga A, Bajić D, Osborne ML, Poyatos JF, Sanchez A (2019). High-order interactions distort the functional landscape of microbial consortia. *PLOS Biology* 17:e3000550. <https://doi.org/10.1371/journal.pbio.3000550>
46. Shade A, Peter H, Allison SD, Baho DL, Berga M, Bürgmann H, Huber DH, Langenheder S, Lennon JT, Martiny JBH, Matulich KL, Schmidt TM, Handelsman J (2012) Fundamentals of microbial community resistance and

- resilience. *Front Microbiol* 3:417. <https://doi.org/10.3389/fmicb.2012.00417>
47. Spragge F, Bakkeren E, Jahn MT, Araujo EBN, Pearson CF, Wang X, Pankhurst L, Cunrath O, Foster KR (2023) Microbiome diversity protects against pathogens by nutrient blocking. *Science* 382:eadj3502. <https://doi.org/10.1126/science.adj3502>
 48. Stouffer DB, Bascompte J (2011) Compartmentalization increases food-web persistence. *Proc Natl Acad Sci USA* 108:3648–3652. <https://doi.org/10.1073/pnas.1014353108>
 49. van Elsas JD, Chiurazzi M, Mallon CA, Elhottová D, Křišťůfek V, Salles JF (2012) Microbial diversity determines the invasion of soil by a bacterial pathogen. *Proc Natl Acad Sci USA* 109:1159–1164. <https://doi.org/10.1073/pnas.1109326109>
 50. Weiss S, Van Treuren W, Lozupone C, Faust K, Friedman J, Deng Y, Xia LC, Xu ZZ, Ursell L, Alm EJ, Birmingham A, Cram JA, Fuhrman JA, Raes J, Sun F, Zhou J, Knight R (2016). Correlation detection strategies in microbial data sets vary widely in sensitivity and precision. *The ISME Journal* 10:1669–1681. <https://doi.org/10.1038/ismej.2015.235>
 51. West SA, Griffin AS, Gardner A (2007) Social semantics: altruism, cooperation, mutualism, strong reciprocity and group selection. *J Evol Biol* 20:415–432. <https://doi.org/10.1111/j.1420-9101.2006.01258.x>
 52. Yachi S, Loreau M (1999) Biodiversity and ecosystem productivity in a fluctuating environment: the insurance hypothesis. *Proc Natl Acad Sci USA* 96:1463–1468. <https://doi.org/10.1073/pnas.96.4.1463>
 53. Zachar I, Boza G (2022) The evolution of microbial facilitation: sociogenesis, symbiogenesis, and transition in individuality. *Front Ecol Evol* 10:798045. <https://doi.org/10.3389/fevo.2022.798045>