`Proc-to-Spec`: A Functorial Map of Network Processes

Anonymous authors
Paper under double-blind review

Abstract

The analysis of dynamic networks is central to understanding complex environmental systems in nature, yet traditional methods often focus on describing changing states rather than formalising the underlying processes of change. In this work, we introduce a category-theoretical framework, `Proc-to-Spec`, that provides a principled, functorial method for analysing the transformations that govern network evolution. We model *resource-constrained systems*, such as those commonly found in biology and ecology, within a source category `Proc`, where morphisms represent dissipative physical processes. We then construct a spectral functor, χ : `Proc` $\to$ `Spec`, that maps each process to a unique linear transformation between the eigenspaces of the network's symmetrised Laplacian. This framework allows us to establish a set of rigorous theorems. We prove that physical conservation laws in `Proc` correspond directly to spectral invariants in `Spec`, such as the conservation of the Laplacian's trace. We derive a spectral sensitivity theorem that formally links resource dissipation to network fragmentation via the Fiedler value. We also establish a stability-spectrum equivalence theorem, proving that a system's physical dynamics converge to a stable state if and only if its spectral geometry converges. We validate our theory with numerical experiments and demonstrate its utility as a tool for scientific discovery in a case study of the Serengeti food web in northern Tanzania. Using a large collection of 1.2 million classified image sets of animal activity from 225 camera traps spread across 1,125 km^2 of the Serengeti National Park from 2010 to 2013, we show that our framework can detect the subtle, cyclical signature of seasonal change and identify the unique geometric fingerprint of the 2011 East Africa drought. Our work provides a different way of thinking about dynamic systems, shifting the focus from describing states to understanding the fundamental geometry of change. Code to reproduce all results in the paper is released at https://anonymous.4open.science/r/tmlr_pts

1 Introduction

The analysis of dynamic networks is fundamental to science, providing the mathematical language to describe systems of interacting components that evolve over time. In biology and ecology, this paradigm is essential for understanding the stability of food webs (Pimm, 1984), the function of gene-regulatory pathways (Barabasi & Oltvai, 2004), and the cascading failures that can lead to abrupt, system-wide critical transitions (Scheffer et al., 2012). These systems are not static; they are governed by a complex interplay of processes—such as predation, resource competition, and metabolic conversion—that continuously reshape their structure and function. The ultimate goal of scientific discovery is to move beyond mere description of these changes and toward a predictive understanding of the underlying principles that govern them.

Current methods for analysing dynamic networks, while powerful, are predominantly descriptive. The standard approach treats a dynamic network as a discrete time-series of static snapshots, $G_1, G_2, \ldots, G_t$. Techniques from spectral graph theory (Chung, 1997) and temporal network analysis (Masuda & Lambiotte, 2016) are then applied to compute metrics for each snapshot and track their evolution. This yields valuable insights into changing properties like connectivity or community structure. However, this approach leaves a critical theoretical gap: *it analyses the states of the system but does not provide a formal language for the processes that transform one state into the next.* A principled mathematical framework that can map a

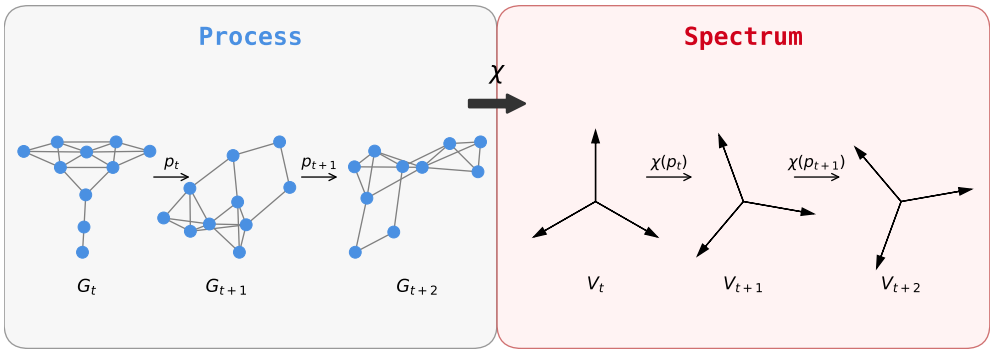

Figure 1: Conceptual illustration of our `Proc-to-Spec` framework. Our work establishes a formal mapping, the spectral functor (χ), from the category of physical processes (`Proc`) to the category of their spectral representations (`Spec`). **Left**: A dynamic network evolves through a sequence of states (G_t, G_{t+1}, G_{t+2}). Each transformation is a physical process (p_t, p_{t+1}), which can represent any resource-constrained change, including perturbations to interaction strengths (edge weights) or topological modifications like the removal or addition of edges. **Right**: Each network state corresponds to a vector space (V_t, V_{t+1}, V_{t+2}) representing the eigenspace of the graph's symmetrised Laplacian. The functor χ maps each physical process p_t to a unique linear transformation $\chi(p_t)$ that describes the change of basis between the corresponding eigenspaces. This framework allows us to analyse complex physical processes by studying their unique and well-behaved geometric signatures in the spectral domain.

specific, causal mechanism of change to its unique, global structural consequence remains less explored. This gap prevents us from moving from observing *that* a system's structure changed to proving *why* it changed in a particular way.

Our work closes this gap by introducing a different way of thinking, `Proc-to-Spec`, as depicted in Figure 1. We deliberately focus on the foundational case of dynamic processes on simple, weighted graphs, which represent a vast class of real-world systems. We propose a conceptual shift from analysing system states to formalising the processes of transformation themselves. The core idea is that the effect of a physical process is most clearly understood not by tracking the state of individual nodes or edges, but by observing its impact on the network's holistic geometric structure. This structure is revealed by the network's Laplacian spectrum. We formalise this by building a direct, provable link between a process and its unique signature as a transformation in the spectral domain. This is analogous to moving from describing the scenes in a movie to analysing the script's rules that govern how one scene can lead to the next. While many methods for analysing dynamic networks focus on aggregating temporal information or identifying recurring temporal motifs (Masuda & Lambiotte, 2016), our approach provides a different perspective by focusing on the geometric nature of the transformations between discrete network states.

We formalise this idea by constructing a categorical framework. We begin by defining a source category, `Proc` (for Process), where objects are weighted, directed graphs representing the state of a *resource-constrained* system. Crucially, the morphisms of `Proc` are not arbitrary graph edits, but are defined as dissipative physical processes that obey the fundamental constraint that resources cannot be created ex nihilo. This grounding in physical law is a key feature of our model. We then define a target category, `Spec` (for Spectrum), as the standard category of real vector spaces and linear transformations. The central contribution of this work is the construction and analysis of a spectral functor[1], $\chi : \texttt{Proc} \to \texttt{Spec}$, that serves as a structure-preserving map between these two worlds. This approach is inspired by the abstract and powerful language of Applied Category Theory (Spivak, 2014).

The functor χ maps a network state in `Proc` to the vector space spanned by the eigenvectors of its symmetrised Laplacian, providing a "spectral signature" of its structure. More importantly, χ maps a dissipative process to a unique linear transformation—a matrix—that describes the corresponding change of basis between the

[1]Our numerical validations show that functoriality holds up to a tiny, numerically tractable error in practice, while our proof in §A.1 details the specific conditions under which it is exact.

old and new eigenspaces. This formalism is not merely a bookkeeping device; it is a generative framework that allows us to derive a set of theorems. We prove that physical conservation laws in `Proc` correspond directly to spectral invariants in `Spec`, such as the conservation of the Laplacian's trace. We establish a spectral sensitivity theorem that formally links resource dissipation to network fragmentation via the Fiedler value. We also establish a stability-spectrum equivalence theorem, proving that a system's physical dynamics converge to a stable state if and only if its spectral geometry converges.

We conduct a rigorous, two-part experimental validation of our framework. First, we use a suite of synthetic experiments to systematically verify each lemma and theorem, confirming the mathematical correctness of our claims with numerical precision. Second, we demonstrate the framework's power as a tool for scientific discovery in a real-world case study on the Serengeti food web (Baskerville et al., 2011). We show that our theory can be applied to high-frequency, real-world data to generate quantitative insights into ecosystem dynamics. Using the large-scale "*Snapshot Serengeti*" camera trap dataset of animal activity in northern Tanzania (Swanson et al., 2015), we demonstrate that our framework is sensitive enough to detect the subtle, cyclical spectral signature of seasonal change and to characterise the unique geometric signature of a major, documented ecological shock—the 2011 East Africa drought. This validation confirms that our framework provides a different, practical, and powerful lens for understanding complex systems.

Our contributions are as follows:

- We introduce a categorical framework, `Proc-to-Spec`, for analysing dynamic networks by formalising the morphisms of change.
- We construct a spectral functor, $\chi : \texttt{Proc} \to \texttt{Spec}$, that maps dissipative physical processes to unique linear transformations in the spectral domain.
- We establish a set of rigorous theorems that provide a direct, provable link between physical properties (conservation, dissipation, stability) and spectral signatures.
- We provide an experimental validation of the framework on both synthetic data and a real-world ecological case study, demonstrating its scientific utility.

The remainder of this paper is structured as follows. In §2, we review related work. In §3, we formally define our model. In §4, we present our theoretical analysis and prove our main theorems. In §5, we detail our synthetic experimental validation. In §6, we apply our framework to the Serengeti case study. We discuss implications and conclude in §7. Detailed proofs are provided in the §A.

2 Related Work

Our `Proc-to-Spec` framework represents a synthesis of ideas from several distinct but related fields. We begin in §2.1 by reviewing the foundational principles of Spectral Graph Theory, the core mathematical language we use for our structural representations. Next, in §2.2, we situate our model within the growing field of dynamic and temporal network analysis. We then ground our work in its target scientific domain in §2.3, discussing the rich history of network modelling in ecology and biology and the persistent challenge of linking local mechanisms to global system dynamics. Our formal approach is heavily inspired by the paradigm of Applied Category Theory, which we discuss in §2.4, clarifying how its principles of compositionality and functorial mappings provide the blueprint for our framework. Finally, in §2.5, we connect our framework to the contemporary landscape of machine learning on graphs.

2.1 Spectral Graph Theory and Its Applications

Spectral graph theory, which studies the properties of a graph via the eigenvalues and eigenvectors of its associated matrices, is a cornerstone of modern data analysis (Chung, 1997). The graph Laplacian, in particular, has found widespread application, with different versions (e.g., unnormalised, symmetric normalised, random walk) offering different perspectives on the graph's structure (Von Luxburg, 2007). Its properties

are used for graph partitioning (Pothen et al., 1990), graph drawing (Koren, 2003), and non-linear dimensionality reduction through methods like Laplacian Eigenmaps and Diffusion Maps (Belkin & Niyogi, 2003). The entire field of Graph Signal Processing (GSP) is built on the idea of using the Laplacian eigenbasis as a Fourier basis for signals defined on graphs (Leus et al., 2023; Shuman et al., 2013). The most famous application is spectral clustering (Von Luxburg, 2007; Shi & Malik, 2000), which uses the eigenvectors of the Laplacian to identify community structure. The Fiedler value (the second-smallest eigenvalue) and its corresponding eigenvector are of particular importance, as they provide a measure of a graph's algebraic connectivity and identify its primary structural bottlenecks (Fiedler, 1973; Mohar et al., 1991). The analysis of how spectra change in response to dynamic processes is less formalised. Research has focused on spectral perturbation theory (Bhatia, 1992), which provides bounds on eigenvalue changes for small perturbations, and on tracking eigenvalue time-series to detect anomalies or regime shifts (Sandryhaila & Moura, 2014). Our work builds on this by providing a functorial framework that maps the processes of change themselves to explicit transformations in the spectral domain, a fundamentally different and more structured approach that aims to formalise the link between the cause and the spectral consequence of a change.

2.2 Dynamic and Temporal Networks

The study of networks that evolve over time is a mature field (Masuda & Lambiotte, 2016; Casteigts et al., 2012). Research has traditionally focused on developing metrics and models to characterise temporal interaction patterns, often distinguishing between representations as sequences of static snapshots, continuous-time contact sequences, or more recently, stream graphs that explicitly model interactions as they occur (Masuda & Lambiotte, 2016; Latapy et al., 2018). Efforts to characterise these networks include identifying temporal motifs (Kovanen et al., 2011), analysing time-respecting reachability and pathfinding (Kempe et al., 2000), and developing temporal centrality measures to identify key nodes in dynamic processes (Lerman et al., 2010). A significant body of work has focused on modelling information diffusion and influence maximisation (Kempe et al., 2003), often using cascade models or threshold dynamics. Models for network evolution aim to capture the mechanisms driving change, such as preferential attachment for scale-free structures (Barabási & Albert, 1999), triadic closure for social clustering (Bianconi & Barabási, 2001), and a wide array of link prediction models that forecast future interactions (Liben-Nowell & Kleinberg, 2003; Hasan & Zaki, 2011). More advanced models have considered continuous-time dynamics using point processes like Hawkes processes (Nguyen et al., 2018), as well as the evolution of community structures over time (Palla et al., 2007; Tantipathananandh et al., 2007; Liu et al., 2020). While this body of work provides a rich vocabulary and powerful tools for *describing* and *predicting* network evolution, it generally lacks a formal, compositional language for the *processes* themselves. The focus remains on the sequence of states or the statistical properties of events, not on a formal algebra of the transformations that connect them, a gap our framework addresses directly.

2.3 Networks in Biological and Environmental Sciences

Network theory has become an indispensable tool in the sciences, providing a language to manage the immense complexity of biological and ecological systems. In ecology, food web analysis is used to study ecosystem stability, resilience, and the role of keystone species (Pimm, 1984; Montoya et al., 2006; Dunne et al., 2002; Williams & Martinez, 2000; Berlow et al., 2004). Network models are used to understand mutualistic interactions, whose nested structure is thought to increase biodiversity (Bascompte et al., 2003), as well as disease propagation (Newman, 2002) and the structure of metapopulations (Hanski, 1998). A key challenge in this field is the identification of early-warning signals for critical transitions, a problem our stability theorems directly address (Kéfi et al., 2014). In biology, networks are used to model protein-protein interactions (PPIs) to uncover functional modules and identify potential drug targets (Jeong et al., 2001; Barzel & Barabási, 2013), and gene-regulatory pathways that control cellular life (Davidson et al., 2002; Alon, 2007). Large-scale metabolic network reconstructions, such as the Recon project for human metabolism, are used in constraint-based modelling to predict metabolic fluxes (Rual et al., 2005). In neuroscience, analysing the brain's structural and functional connectomes is a central goal, with projects like the Human Connectome Project providing massive datasets (Sporns et al., 2005; Bullmore & Sporns, 2009; Honey et al., 2007; Van Essen et al., 2012). These applications demonstrate a clear and urgent need for models that can

handle dynamic *processes* and provide insight into system-level properties like stability and resilience. Our work is directly motivated by this need, providing a formal language and a set of theoretical tools specifically designed to analyse the resource-constrained dynamics that are characteristic of these natural systems.

2.4 Categorical and Geometric Approaches to Systems

Our framework is inspired by a growing movement to apply principled, abstract mathematics to machine learning and systems modelling. Geometric Deep Learning seeks to unify graph-based models and beyond by focusing on underlying symmetries and invariances, building models that respect the geometry of their domain (Bronstein et al., 2017). Alongside spectral methods, tools from Topological Data Analysis (TDA) are also used to study the higher-order structure of networks through persistent homology (Petri et al., 2013). More abstractly, Applied Category Theory provides a formal language for compositionality, arguing that systems are best understood by how their components compose (Fong & Spivak, 2018). This approach has been used to model a wide range of systems, including databases (Guyot et al., 2022), electrical circuits (Takahashi, 2023), dynamical systems (Behrisch et al., 2017), compositional game theory (Ghani et al., 2018), and network protocols (Fong & Spivak, 2018). Other formalisms like operads are used to describe more general systems of composition (Baez & Stay, 2010). This research program, pioneered by figures like Baez & Stay (2010) and Spivak (2014), argues that the language of functors and morphisms is the natural way to describe complex, interacting systems. Our work contributes to this paradigm by providing a concrete instantiation of a functorial model. We use the categorical language not as an end in itself, but as the natural grammar to build a specific, testable, and scientifically relevant theory that connects process to structure, thus grounding the abstract formalism in concrete, provable spectral consequences.

2.5 Machine Learning on Graphs

In recent years, machine learning on graph-structured data has been dominated by the success of Graph Neural Networks (GNNs), a broad class of models based on a neighbourhood aggregation or message-passing scheme (Scarselli et al., 2008; Gori et al., 2005; Gilmer et al., 2017). Architectures like Graph Convolutional Networks (GCNs) (Kipf, 2016), Graph Attention Networks (GATs) (Veličković et al., 2017), and GraphSAGE (Hamilton et al., 2017) have achieved state-of-the-art performance on tasks like node classification and link prediction. More recent developments include Graph Transformers, which aim to capture long-range dependencies (Shi et al., 2020). To handle dynamic networks, researchers have developed various temporal GNNs that integrate GNN principles with recurrent or attention-based sequence models (Seo et al., 2018; Rossi et al., 2020; Kazemi et al., 2020; Manessi et al., 2020). Other approaches include graph kernels (Shervashidze et al., 2011), and representation learning techniques like DeepWalk (Perozzi et al., 2014) and node2vec (Grover & Leskovec, 2016). A limitation of many of these powerful models is their heavy-parameterisation nature. They provide high predictive accuracy but often lack scientific interpretability, and the community has invested significant effort in developing post-hoc explanation methods like GNNExplainer (Ying et al., 2019). Moreover, they do not come with the formal, provable guarantees of a theoretical framework, and can suffer from issues like oversmoothing (Li et al., 2018). Our `Proc-to-Spec` framework aims to alleviate this limitation, offering a more transparent approach where the link between process and outcome is explicit, provable, and scientifically interpretable, aiming for understanding over raw prediction.

3 The `Proc-to-Spec` Framework

In this section, we provide the formal mathematical construction of our `Proc-to-Spec` framework. The framework is built upon a single, central idea: *the existence of a structure-preserving map—a functor—that translates the dynamics of a physical system into the geometric language of linear algebra.* The resulting model is a logical consequence of defining this functor and its domain and codomain. We begin by defining our key notations, summarised in Table 1.

Table 1: Key symbols and notations used in the paper.

Symbol	Description
$G = (V, E, W)$	A weighted, directed graph representing a network state.
V, E, W	The set of vertices, edges, and the edge weight function, respectively.
$\mathcal{R}(G)$	The total resource in a network, defined as the sum of all edge weights.
$p : G \to G'$	A process (morphism) that transforms one network state to another.
`Proc`	The category of resource-constrained dynamic networks.
L_W	The weighted Laplacian of a graph with weight matrix W.
L_{sym}	The symmetrised Laplacian, used for spectral analysis.
$\lambda_i, \mathbf{v}_i$	The i-th eigenvalue and corresponding eigenvector of L_{sym}.
λ_2	The Fiedler value, or algebraic connectivity of the graph.
`Spec`	The category of finite-dimensional real vector spaces.
χ	The spectral functor, mapping from `Proc` to `Spec`.
$\chi(p)$	A linear transformation representing the process p in the spectral domain.

3.1 The Source Category: `Proc`

The source category, `Proc`, is designed to represent the physical reality of resource-constrained dynamic systems commonly found in biological and ecological sciences.

Objects. An object in `Proc` is a network state, formally defined as a weighted, directed graph $G = (V, E, W)$, where:

- V is a finite set of n vertices, representing the components of the system (e.g., species in an ecosystem, proteins in a cell).
- $E \subseteq V \times V$ is a set of directed edges, representing interactions between components.
- $W : E \to \mathbb{R}_0^+$ is a weight function that assigns a non-negative real value to each edge, representing the strength, capacity, or rate of flow of a resource (e.g., energy, biomass, information).

Morphisms. A morphism in `Proc`, denoted $p : G \to G'$, is a **dissipative process** that transforms an initial state G into a final state G'. The defining characteristic of these systems is that they operate under resource limitation. We formalise this by defining the total resource of a network as the sum of all its edge weights:

$$\mathcal{R}(G) = \sum_{(u,v)\in E} W(u, v) \tag{1}$$

A process p is defined as dissipative if the total resource in the final state is less than or equal to the total resource in the initial state:

$$\mathcal{R}(G') \leq \mathcal{R}(G) \tag{2}$$

This single constraint is fundamental. It reflects the second law of thermodynamics, where energy is lost in transfers between trophic levels (Odum, 1957), and the principle of resource competition that governs the dynamics of all biological populations (Tilman, 1982). This constraint ensures that our model is physically and biologically grounded for scientific discovery.

3.2 The Target Category: `Spec`

The target category, `Spec`, provides the abstract geometric space where the structural properties of our networks are analysed. It is the standard category of finite-dimensional real vector spaces.

Objects. An object is a finite-dimensional vector space U over the field of real numbers, $\mathbb{R}$.

Morphisms. A morphism is a linear transformation $T : U \to V$ between two vector spaces.

3.3 The Spectral Functor $\chi : \texttt{Proc} \to \texttt{Spec}$

The spectral functor χ is the heart of our framework. It provides the formal, structure-preserving map from the physical world of processes to the geometric world of spectra.

Action on Objects. The functor χ maps a network object $G \in \texttt{Proc}$ to the n-dimensional vector space spanned by the eigenvectors of its symmetrised Laplacian. This construction is essential for ensuring a well-defined geometric representation. Given a graph G with n vertices and weight function W:

1. We first construct the $n \times n$ weighted adjacency matrix A_W, where $(A_W)_{ij} = W(i, j)$.
2. To guarantee a real spectrum and a complete orthonormal eigenbasis, we construct the symmetrised adjacency matrix, A_{sym}, where $(A_{sym})_{ij} = (W(i, j) + W(j, i))/2$. This captures the underlying "connectivity fabric" of the network, which is often the primary interest in resilience and stability studies (Mohar et al., 1991).
3. From this, we construct the symmetrised Laplacian:
$$L_{sym} = D_{sym} - A_{sym} \tag{3}$$
where D_{sym} is the diagonal matrix of weighted degrees derived from A_{sym}.
4. This ensures that L_{sym} is a real symmetric matrix and thus has a complete set of n orthonormal eigenvectors $\{\mathbf{v}_1, \ldots, \mathbf{v}_n\}$.
5. The functor maps the graph object to the vector space these eigenvectors span:
$$\chi(G) = \text{span}(\mathbf{v}_1, \ldots, \mathbf{v}_n) \cong \mathbb{R}^n \tag{4}$$

Action on Morphisms. The functor χ maps a process $p : G \to G'$ to the unique linear transformation $\chi(p) : \chi(G) \to \chi(G')$ that describes the change of basis between the respective eigenspaces. If $B = \{\mathbf{v}_i\}$ is the eigenbasis of L_{sym} and $B' = \{\mathbf{v'}_i\}$ is the eigenbasis of L'_{sym}, then $\chi(p)$ is the unique linear map that represents this geometric transformation.

Physical Intuition of the Functor. Conceptually, the spectral functor χ can be understood as a mathematical "lens" that allows us to view the physical process through the language of geometry. The eigenbasis of a network's Laplacian represents its fundamental modes of variation—its "structural harmonics". A physical process, p, perturbs the network, creating a new set of structural harmonics. The transformation matrix $\chi(p)$ precisely quantifies how each of the old harmonics is distributed or "scattered" among the new ones (see Figure 10 for such scattering in a real-world ecological web). It is, in essence, a mathematical description of the structural reorganisation induced by the process. A simple process might only slightly alter the harmonics, resulting in a transformation close to the identity, while a complex process might completely scramble them, resulting in a highly complex rotational transformation.

3.4 Model Scope and Assumptions

Our framework is designed to be a model for rigorous scientific inquiry in environmental sciences, with carefully chosen scope and assumptions that we state explicitly here.

- **Deterministic Processes.** We model processes as deterministic, representing the expected outcome of potentially stochastic interactions. This is a standard and powerful simplification in the modelling of complex systems, providing a tractable first-order approximation of the system's dynamics (May, 2001).
- **Focus on Edge Weights.** We assume that the primary dynamics of the system can be captured by changes in the interaction strengths (edge weights). This is a well-justified focus in ecological network modelling, where fluctuations in interaction strength are a primary driver of system stability (Allesina & Tang, 2012). This naturally includes topological changes as a special case.

- **Symmetrised Laplacian.** The core operator used by the spectral functor is the symmetrised Laplacian, L_{sym}. We select this operator due to its desirable properties, namely a real spectrum and a complete orthonormal eigenbasis, which are essential for defining a consistent spectral geometry for our framework.

4 Theoretical Analysis

Having formally defined the `Proc-to-Spec` framework in §3, we now establish its key theoretical properties. Our analysis is structured to build a complete framework from the ground up, moving from foundational mathematical guarantees to the main scientific results. We begin in §4.1 by proving that our spectral map, χ, is a valid functor, a foundational result that ensures the mathematical soundness and compositionality of our entire framework. Next, in §4.2, we establish the core dictionary that translates the physical constraints of the `Proc` category into the language of the `Spec` category; here, we prove our key theorems on the spectral signatures of resource conservation and dissipation. Building on this, we present our main result in §4.3, the Stability-Spectrum Equivalence theorem, which provides a formal and testable link between the long-term dynamical stability of a system and the convergence of its spectral geometry. Finally, in §4.4, we derive a toolkit of more specialised theorems that provide concrete, interpretable matrix signatures for specific, common types of processes, such as local perturbations and node removals. Together, these results form a complete theoretical basis for using our framework as a tool for scientific reasoning.

4.1 Functoriality of the Spectral Map χ

We begin by establishing the foundational mathematical property of our framework: that the map χ is a valid functor. This result is crucial as it guarantees that our spectral representation is a true and consistent reflection of the underlying system's dynamics, respecting both identity and the composition of processes.

Lemma 1 (Functoriality of χ)**.** *The map $\chi : \texttt{Proc} \to \texttt{Spec}$ is a functor. It preserves identity morphisms and the composition of morphisms.*

Proof Sketch. To prove that χ is a functor, we must verify two conditions.

1. **Preservation of Identity:** The identity morphism in `Proc` is the process $\mathrm{id}_G : G \to G$, which leaves the network unchanged. This means the initial and final weight matrices are identical ($W' = W$), leading to identical symmetrised Laplacians ($L'_{sym} = L_{sym}$) and thus identical eigenbases. The linear transformation that maps an orthonormal basis to itself is the identity transformation, $\mathrm{id}_{\chi(G)}$. Thus, $\chi(\mathrm{id}_G) = \mathrm{id}_{\chi(G)}$.
2. **Preservation of Composition:** Consider two composable processes, $p_1 : G_1 \to G_2$ and $p_2 : G_2 \to G_3$. The map $\chi(p_1)$ is the change of basis matrix from the eigenbasis of G_1 to that of G_2, and $\chi(p_2)$ is the change of basis from G_2 to G_3. The composition of linear transformations, $\chi(p_2) \circ \chi(p_1)$, corresponds to the matrix product of these change of basis matrices. By the chain rule for change of basis, this product is precisely the matrix that transforms the basis of G_1 directly to the basis of G_3, which is by definition $\chi(p_2 \circ p_1)$. Thus, $\chi(p_2 \circ p_1) = \chi(p_2) \circ \chi(p_1)$.

The full proof is provided in §A.1. □

Interpretation. This lemma provides the guarantee of our framework's logical consistency. It ensures that our spectral "lens" does not distort the structure of the system's dynamics. From a scientific perspective, this has two critical implications. First, the preservation of identity means that a system in a stable, unchanging state will have a stable, unchanging spectral representation; our method does not introduce artificial dynamics. Second, the preservation of composition means that our framework respects causality. For example, the overall structural impact of a drought followed by the introduction of an invasive species is precisely the composition of their individual spectral transformations. This allows scientists to model complex, multi-stage scenarios with the confidence that the resulting analysis is a faithful representation of the composite process, making our "spectral accounting" of change both rigorous and reliable.

4.2 The Physics-to-Spectra Dictionary

Having established the mathematical soundness of our framework, we now derive the core dictionary that translates the physical laws governing the `Proc` category into the geometric language of `Spec`. The following theorems show that fundamental physical constraints—namely, the conservation and dissipation of resources—induce specific, non-trivial, and observable signatures in the spectral domain.

Theorem 1 (The Spectral Trace Conservation Law)**.** *Let $p : G \to G'$ be a conservative process, where the total resource is unchanged ($\mathcal{R}(G) = \mathcal{R}(G')$). The trace of the symmetrised Laplacian is conserved, i.e., $\mathrm{Tr}(L'_{sym}) = \mathrm{Tr}(L_{sym})$. Consequently, the sum of the Laplacian eigenvalues is an invariant of the process.*

Proof Sketch. The key insight is to first establish a direct identity between the total resource of a network and the trace of its symmetrised Laplacian. The trace of L_{sym} is the sum of its diagonal elements, which are the weighted degrees of the symmetrised graph. This sum is equivalent to the sum of all entries in the symmetrised adjacency matrix, A_{sym}. By substituting the definition of A_{sym}, we show that this sum is precisely equal to the total resource, $\mathcal{R}(G)$. Therefore, the identity $\mathrm{Tr}(L_{sym}) = \mathcal{R}(G)$ holds for any graph. For a conservative process, since $\mathcal{R}(G) = \mathcal{R}(G')$, it follows directly that $\mathrm{Tr}(L'_{sym}) = \mathrm{Tr}(L_{sym})$. Because the trace of a matrix is equal to the sum of its eigenvalues, the sum of the eigenvalues is also conserved. The full proof is provided in §A.2. □

Interpretation. This theorem provides our first concrete link between a physical law and a geometric invariant. It demonstrates that if a system is closed and only redistributes its internal resources (e.g., biomass transfer within a food web without external inputs or losses), the sum of its spectral eigenvalues remains constant. This spectral sum can be interpreted as a measure of the total "structural energy" or "information capacity" of the network. This result provides an integrity check for models of closed ecosystems, ensuring that the simulated dynamics correctly preserve this global spectral quantity.

Theorem 2 (The Spectral Sensitivity of Algebraic Connectivity)**.** *Let $p : G \to G'$ be a process that induces a sufficiently small change in the symmetrised Laplacian, $\Delta L_{sym} = L'_{sym} - L_{sym}$. If the process is structurally fragmenting, defined as satisfying the condition $\mathbf{v}_2^T(\Delta L_{sym})\mathbf{v}_2 < 0$, where $\mathbf{v}_2$ is the Fiedler eigenvector of the initial graph G, then the Fiedler value will decrease ($\lambda'_2 < \lambda_2$).*

Proof Sketch. The proof relies on first-order matrix perturbation theory, which states that the change in an eigenvalue, $\Delta\lambda_k$, can be approximated by the quadratic form $\mathbf{v}_k^T(\Delta L_{sym})\mathbf{v}_k$. By applying this principle to the Fiedler value ($k = 2$), we find that $\lambda'_2 - \lambda_2 \approx \mathbf{v}_2^T(\Delta L_{sym})\mathbf{v}_2$. The theorem's premise is precisely that the term on the right-hand side is negative. Therefore, it follows directly that for small perturbations, $\lambda'_2 - \lambda_2 < 0$. The full proof is provided in §A.3. □

Interpretation. This theorem provides a nuanced quantitative insight about resource loss in physical systems. It formalises the critical scientific idea that the location of a disturbance is as important as its magnitude. The Fiedler eigenvector, $\mathbf{v}_2$, identifies the network's primary structural vulnerability or "fault line". This theorem proves that a process, even a dissipative one, only harms the network's algebraic connectivity if it is "aligned" with this vulnerability—that is, if it preferentially weakens the crucial links that bridge the network's main communities. For example, a disease affecting a keystone predator that connects two sub-webs would be structurally fragmenting, causing a sharp drop in algebraic connectivity. In contrast, the loss of a peripheral species might have a negligible effect. This provides a valuable diagnostic tool for assessing the resilience of an ecosystem, allowing scientists to distinguish between benign and potentially catastrophic systemic changes.

4.3 The Stability-Spectrum Equivalence

Building upon the foundational link between physical processes and their spectral signatures, we now establish our main result. The following theorem provides a formal equivalence between the long-term dynamical stability of a network in the `Proc` category and the convergence of its geometric representation in the `Spec` category. This result allows our framework to function as a predictive tool, allowing the assessment of a system's stability through its observable algebraic properties.

Theorem 3 (The Stability-Spectrum Equivalence)**.** *A dynamic network sequence $(G_t)_{t=1}^{\infty}$ governed by dissipative processes converges to a stable state G_∞ if and only if its corresponding sequence of spectral data (eigenvalues and eigenvectors of $L_{sym,t}$) converges to a stable limit.*

Proof Sketch. As this is an "if and only if" statement, we sketch the proof in two directions.

1. **Stability $\implies$ Spectral Convergence:** We first assume the system converges to a stable state, which means the sequence of weight matrices converges to a limit, $W_t \to W_\infty$. The map from a weight matrix W to its symmetrised Laplacian L_{sym} is continuous. Furthermore, the spectral decomposition of a symmetric matrix (its eigenvalues and eigenvectors) is a continuous function of the matrix entries. By the property of continuous functions, the convergence of the weight matrices ($W_t \to W_\infty$) implies the convergence of the Laplacians ($L_{sym,t} \to L_{sym,\infty}$), which in turn implies the convergence of their spectral data.
2. **Spectral Convergence $\implies$ Stability:** We now assume the full set of spectral data (all eigenvalues and eigenvectors) converges. A symmetric matrix is uniquely determined by its spectral decomposition. Therefore, the convergence of the spectral data implies the convergence of the sequence of symmetrised Laplacians, $L_{sym,t} \to L_{sym,\infty}$. The mapping from a weight matrix W to L_{sym} is injective for a given graph topology. Thus, the convergence of the Laplacians implies the convergence of the underlying weight matrices, $W_t \to W_\infty$, which is the definition of a stable state.

The full proof is provided in §A.4. □

Interpretation. This theorem establishes a rigorous and testable equivalence between a system's physical behaviour and its abstract structural properties. In practical terms, it means that the stability of a complex ecosystem or biological network can be definitively assessed by monitoring its "spectral signature". If the eigenvalues and eigenvectors of the system's Laplacian stop changing, the system has reached an equilibrium. This moves beyond correlation to a formal equivalence, providing a diagnostic tool. For example, ecologists can use time-series data to determine if a recovering ecosystem has truly stabilised or if it is still in a transient state, simply by observing whether its spectral representation has converged. This provides a non-invasive method for understanding the long-term trajectory and health of complex systems.

4.4 A Toolkit for Process Interpretation

The preceding theorems establish the foundational properties of our framework. We now derive a toolkit of more specialised results that provide concrete, interpretable matrix signatures for common types of processes. These theorems allow a scientist to move from observing a spectral transformation back to inferring the specific nature of the underlying physical process that caused it, providing a powerful method for causal inference and system diagnostics.

Lemma 2 (The Change of Basis Formula)**.** *Let $p : G \to G'$ be a process, with $\{\mathbf{v}_i\}$ and $\{\mathbf{v'}_j\}$ being the orthonormal eigenbases of the initial and final Laplacians, respectively. The entry (i, j) of the matrix representation of the linear transformation $\chi(p)$ is given by the inner product of the respective basis vectors: $(\chi(p))_{ij} = \langle \mathbf{v'}_i, \mathbf{v}_j \rangle$.*

Proof Sketch. The matrix for the linear transformation $\chi(p)$ represents the change of coordinates from the initial basis $B = \{\mathbf{v}_i\}$ to the final basis $B' = \{\mathbf{v'}_j\}$. The j-th column of this matrix is the vector $\mathbf{v}_j$ expressed in the coordinates of the new basis. Since B' is an orthonormal basis, the i-th coordinate is simply the projection of $\mathbf{v}_j$ onto $\mathbf{v'}_i$, which is given by their inner product. The full proof is provided in §A.5. □

Interpretation. This lemma provides the direct algebraic formula for computing the transformation matrix. It gives a precise meaning to each entry: the magnitude of $(\chi(p))_{ij}$ quantifies the alignment or projection of the j-th original structural mode onto the i-th final structural mode. A large diagonal entry $(\chi(p))_{ii}$ signifies that a structural mode has been preserved, while a large off-diagonal entry $(\chi(p))_{ij}$ signifies a significant structural "rewiring" where the roles of two modes have become mixed.

Theorem 4 (The Rank-One Update Signature)**.** *Let p be a simple process that only perturbs the weight of a single edge between nodes a and b. The resulting change in the symmetrised Laplacian, ΔL_{sym}, is a rank-one matrix. Consequently, the transformation matrix $\chi(p)$ is a low-rank perturbation of the identity matrix.*

Proof Sketch. A change in the weight of a single edge (a, b) results in a change matrix ΔL_{sym} with non-zero entries only at positions $(a, a), (b, b), (a, b)$, and (b, a). As shown in the full proof, this matrix can be expressed as a scalar multiple of a single outer product, $\delta(\mathbf{e}_a - \mathbf{e}_b)(\mathbf{e}_a - \mathbf{e}_b)^T$, making it a rank-one matrix. By matrix perturbation theory, a low-rank update to a matrix results in a correspondingly simple, low-rank perturbation to its spectral decomposition. Therefore, the change of basis matrix $\chi(p)$ will be close to the identity, differing only by a low-rank update related to the eigenvector components at the perturbed nodes. The full proof is provided in §A.6. □

Interpretation. This theorem proves that a local cause has a local signature. It formalises the intuitive idea that a small, isolated event in an ecosystem (e.g., the weakening of a single predator-prey relationship) should not cause a catastrophic, chaotic rewiring of the entire system's structure. It provides a diagnostic tool: if an observed transformation matrix $\chi(p)$ is well-approximated by a low-rank perturbation of the identity, one can infer that the underlying physical cause was a simple, localised process.

Theorem 5 (The Structural Inertia Theorem)**.** *Let p be a process that induces a small perturbation ΔL_{sym}. The resulting transformation matrix $\chi(p)$ is diagonally dominant. The magnitude of its off-diagonal entries is bounded by the norm of the perturbation and the spectral gaps of the original graph.*

Proof Sketch. The proof relies on the Davis-Kahan theorem from matrix perturbation theory. The entries of $\chi(p)$ are the inner products $\langle \mathbf{v'}_i, \mathbf{v}_j \rangle$. For $i = j$, this value is close to 1. For $i \neq j$, the Davis-Kahan theorem provides a bound on the sine of the angle between the old eigenvector $\mathbf{v}_j$ and the new one $\mathbf{v'}_i$, showing that this angle remains close to orthogonal. This deviation from orthogonality, which determines the magnitude of the off-diagonal entries, is bounded by $\|\Delta L_{sym}\|_2 / |\lambda_i - \lambda_j|$. Thus, for a small perturbation, the off-diagonal entries are small, and the matrix is diagonally dominant. The full proof is provided in §A.7. □

Interpretation. This theorem formalises the concept of structural robustness. It proves that complex systems possess a form of inertia; their fundamental organisational modes (the eigenvectors) are resistant to small, arbitrary changes. A small process cannot cause a catastrophic re-shuffling of all the primary structural modes. A transformation matrix $\chi(p)$ with large off-diagonal entries is therefore a clear signature of a major, non-perturbative structural reorganisation of the network, rather than a simple fluctuation.

We now consider processes that alter the network's topology by removing a node. While such processes fall outside our primary functorial definition, which maps between spaces of the same dimension, our geometric approach still provides a unique and powerful characterisation. As we prove in §A.8, node removal is uniquely identifiable as a projection, whose kernel corresponds directly to the removed node, providing a distinct signature for this class of major topological changes.

Theorem 6 (The Node Removal Signature)**.** *Let p be a process that removes a node k from a network G with n nodes. The resulting transformation $\chi(p)$ maps the original n-dimensional eigenspace to the new $(n-1)$-dimensional eigenspace and is a projection operator.*

Proof Sketch. The process of node removal transforms the original vector space $\chi(G) \cong \mathbb{R}^n$ to a lower-dimensional space $\chi(G') \cong \mathbb{R}^{n-1}$. Such a transformation is a projection. The kernel of this projection (the part of the space mapped to zero) is the one-dimensional subspace corresponding to the removed node. The Cauchy Interlacing Theorem guarantees a predictable relationship between the old and new eigenvalues, ensuring the transformation is well-behaved. The full proof is provided in §A.8. □

Interpretation. This theorem provides a unique, identifiable signature for a major topological event: the complete failure or removal of a system component. Unlike the subtle changes from weight perturbations, a node removal causes a change in the very dimension of the state space. If an observed transformation matrix is found to be a projection of rank $n - 1$, one can infer with high confidence that the underlying physical

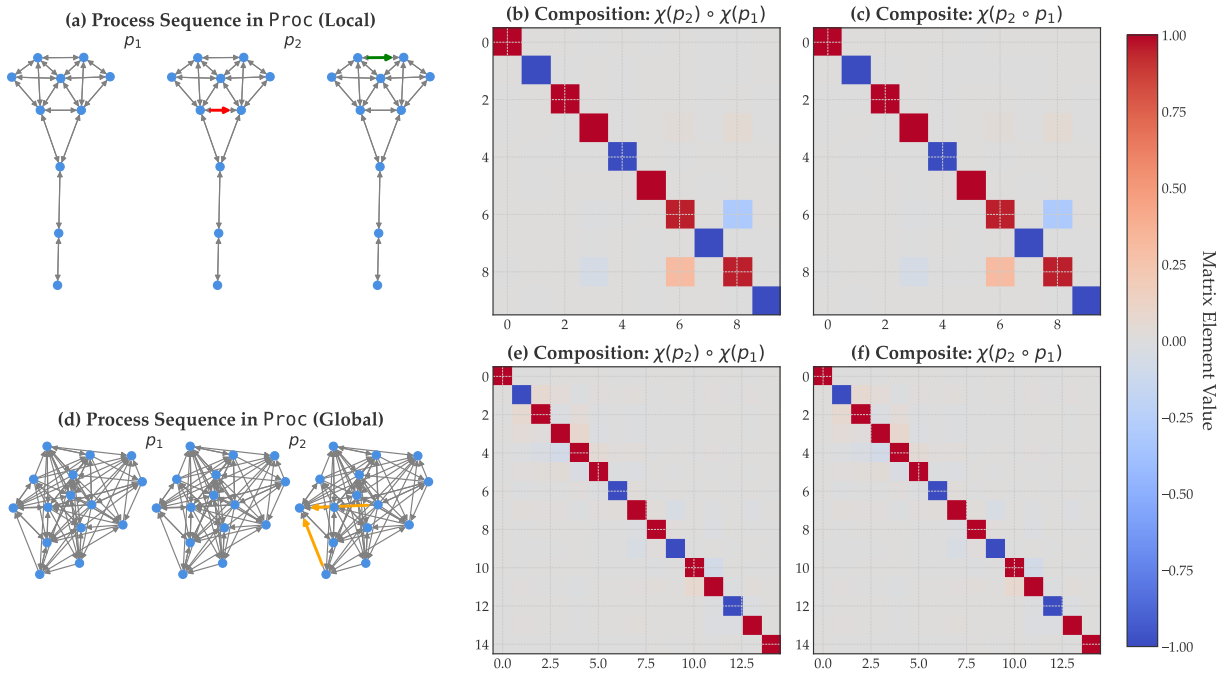

Figure 2: **Numerical validation of the functoriality of χ (Lemma 1).** This figure provides visual and quantitative proof that our framework respects the composition of processes for both simple, local changes (top row) and complex, global changes (bottom row). **(a)** A sequence of three network states in `Proc` connected by two simple processes, p_1 and p_2, each a single edge perturbation. The modified edge in each step is highlighted (red for decreased weight, green for increased). **(b)** The resulting transformation matrix for the composition of the two processes, computed as the matrix product $\chi(p_2) \circ \chi(p_1)$. **(c)** The transformation matrix for the single composite process, $\chi(p_2 \circ p_1)$. This matrix is visually and numerically identical to the one in (b). **(d)** A sequence of network states connected by two complex, global processes: a dissipative process (p_1) that modifies all edges, followed by a conservative one (p_2). **(e)** The matrix for the composition of the complex processes, showing a non-trivial, global transformation. **(f)** The matrix for the single composite complex process, which is again identical to its compositional counterpart in (e). The numerical difference between the composed and composite matrices is negligible in both cases (Frobenius norm for the simple case: 3.3×10^{-15}; for the complex case: 4.2×10^{-15}).

process was the removal of a single node. This is a simple tool for diagnosing critical failures in a system, such as the extinction of a species from an ecosystem.

Theorem 7 (The Signal Transport Theorem)**.** *Let* $\mathbf{f}$ *be a vector representing a signal on the nodes of a network* G *(e.g., population densities). After a process* p*, the signal* $\mathbf{f}'$ *that maintains the same coordinates with respect to the new eigenbasis of* G' *is given by the transformation* $\mathbf{f}' = T_{transport}\mathbf{f}$*, where the transport matrix is* $T_{transport} = (\chi(p)^T)^{-1}$*.*

Proof Sketch. We define the signal's spectral coordinates on G as $\mathbf{a} = V^T\mathbf{f}$, where V is the matrix of eigenvectors. The transported signal on G' is defined as $\mathbf{f}' = V'\mathbf{a}$. Substituting the expression for $\mathbf{a}$ gives $\mathbf{f}' = (V'V^T)\mathbf{f}$, so $T_{transport} = V'V^T$. From our definition of $\chi(p) = (V')^T V$, we can show that $((\chi(p))^T)^{-1} = (V^TV')^{-1} = (V')^{-1}(V^T)^{-1} = V'V = T_{transport}$ (since V is orthonormal, $V^{-1} = V^T$). The full proof is provided in §A.9. □

Interpretation. This theorem provides a concrete, practical tool for prediction. It answers the question: "If the network structure changes from G to G', how would a data pattern defined on the original structure be expressed on the new structure?" For instance, it allows scientists to predict how a specific pattern of gene expression would be re-distributed across a cell if the underlying gene-regulatory network is rewired by

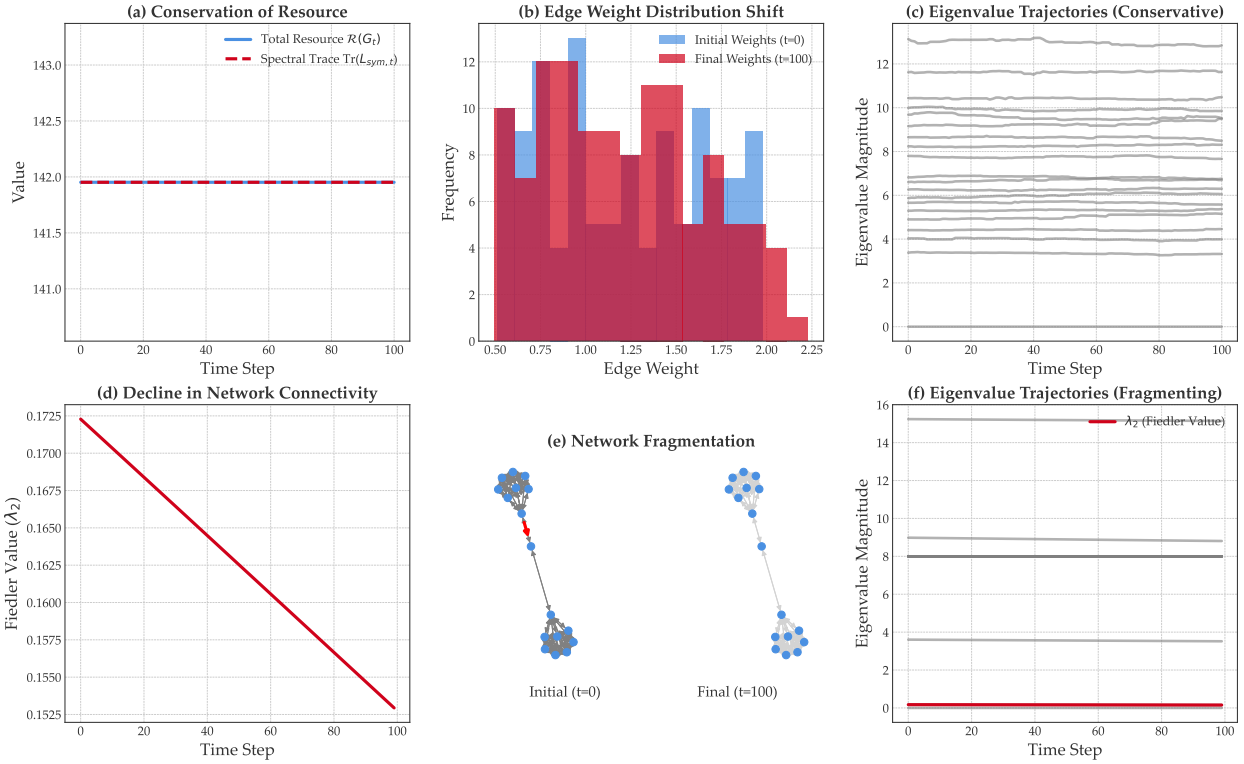

Figure 3: **Numerical validation of the physics-to-spectra dictionary (Theorems 1 and 2).** This figure demonstrates that the physical constraints defined in `Proc` have direct, predictable consequences in `Spec`. **Top Row**: A network was subjected to a purely conservative process for 100 steps. **(a)** The total resource $\mathcal{R}(G_t)$ and the spectral trace $\mathrm{Tr}(L_{sym,t})$ are plotted over time. The lines are perfectly flat and coincident, numerically verifying the identity $\mathrm{Tr}(L_{sym}) = \mathcal{R}(G)$ and proving that the spectral trace is a conserved quantity under conservative dynamics. **(b)** Histograms of the initial and final edge weights show that while the process has significantly redistributed the individual resource values, the total resource is unchanged. **(c)** The trajectories of all eigenvalues are shown. While individual structural modes fluctuate, their sum is perfectly constant, as proven in (a). **Bottom Row**: A barbell graph was subjected to a structurally fragmenting process for 100 steps. **(d)** The Fiedler value (λ_2), a measure of algebraic connectivity, shows a clear and monotonic decline, confirming that the process successfully fragments the network. **(e)** Visualisation of the initial network with its critical bridge edge highlighted (left) and the final network (right), where the weakening of the bridge has caused the two communities to drift apart. **(f)** The trajectories of all eigenvalues, with the Fiedler value highlighted. The plot shows λ_2 "peeling away" from the rest of the spectrum and dropping towards zero, a classic signature of network fragmentation.

a mutation. This moves the framework beyond structural analysis to a tool for predicting the evolution of functional patterns on the network.

5 Numerical Validation

In this section, we present a suite of controlled experiments on synthetic data to provide a rigorous validation of our theoretical claims from §4. We begin by numerically verifying the functoriality of our spectral map, confirming that it respects the composition of processes (§5.1). We then validate the core physics-to-spectra dictionary, providing evidence for both the Spectral Trace Conservation Law and the Spectral Sensitivity of connectivity (§5.2). Building on this, we provide a multi-faceted validation of our main result, the Stability-Spectrum Equivalence theorem (§5.3). Further, we demonstrate the diagnostic power of our framework by

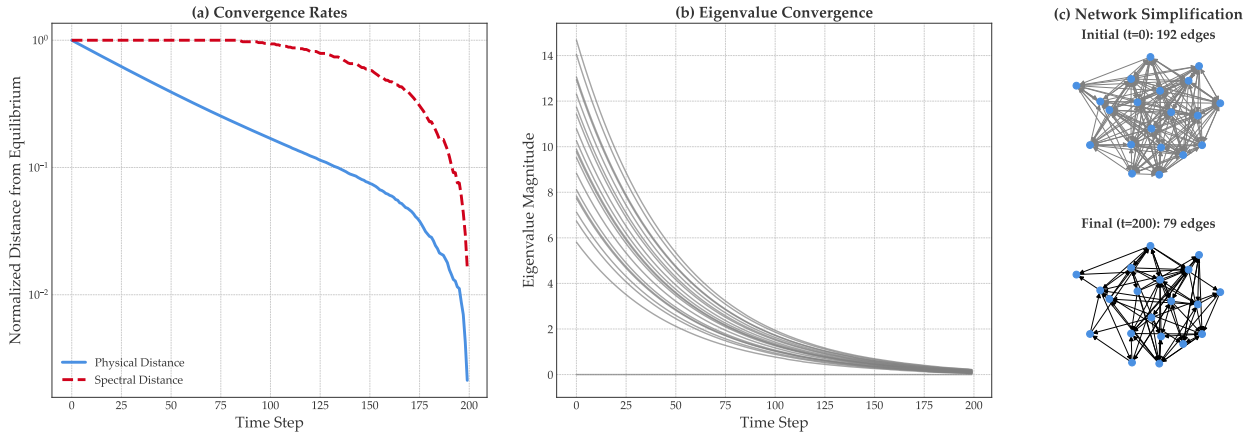

Figure 4: **Numerical validation of the Stability-Spectrum Equivalence (Theorem 3).** This figure provides a multi-faceted demonstration of the equivalence between a system's physical convergence to equilibrium and the stabilisation of its spectral representation. A randomly generated network was subjected to a dissipative process with edge pruning for 200 time steps. **(a)** The convergence rates of the physical and spectral representations. The log-linear plot shows that both the physical distance from equilibrium (blue, solid) and the spectral distance (red, dashed) decay to zero, confirming that the system stabilises as its spectral representation does. The differing decay profiles reveal a non-trivial insight: the spectral representation is less sensitive to initial fluctuations but provides a sharper signal of the final convergence to a stable state. **(b)** The convergence of the full Laplacian spectrum. The plot shows the trajectories of all eigenvalues over time. The initially chaotic dynamics smoothly resolve into a set of stable, horizontal lines, providing a holistic visualisation of the system's entire geometric structure settling into its final, equilibrium configuration. **(c)** The physical process of network simplification. The visualisation shows the initial, dense network state at $t = 0$ (left) and the final, sparse equilibrium state at $t = 200$ (right). The dissipative and pruning processes have driven the system to shed non-essential connections, converging from 192 edges to a stable "backbone" of 79 edges.

verifying the unique spectral signatures predicted by our process interpretation toolkit (§5.4). Finally, we validate the predictive power of our framework using signal transport (§5.5).

5.1 Verification of Functoriality

The cornerstone of our entire framework is the claim that the spectral map χ is a valid functor (Lemma 1). To verify this, we must show that it respects the composition of processes. We test this property in two scenarios: one with simple, local processes and one with complex, global processes. As shown in Figure 2, in both the simple case (Panel a-c) and the complex case (Panel d-f), the transformation matrix computed for the single composite process is numerically (almost) identical to the matrix product of the individual transformations. This provides clear, strong evidence that our framework is mathematically sound and that our "spectral accounting" of change is consistent and reliable.

5.2 The Spectral Signatures of Physical Constraints

Next, we validate our "physics-to-spectra dictionary", which links the physical constraints of the `Proc` category to specific signatures in the `Spec` category. Figure 3 shows the results of two simulations designed to test these links.

The top row validates our Trace Conservation Law (Theorem 1). We subject a network to a purely conservative process that redistributes resources internally. As predicted, the total resource and the spectral trace remain perfectly constant and coincident throughout the simulation (Panel a), even as the individual edge weights are shuffled (Panel b) and the individual eigenvalues fluctuate (Panel c).

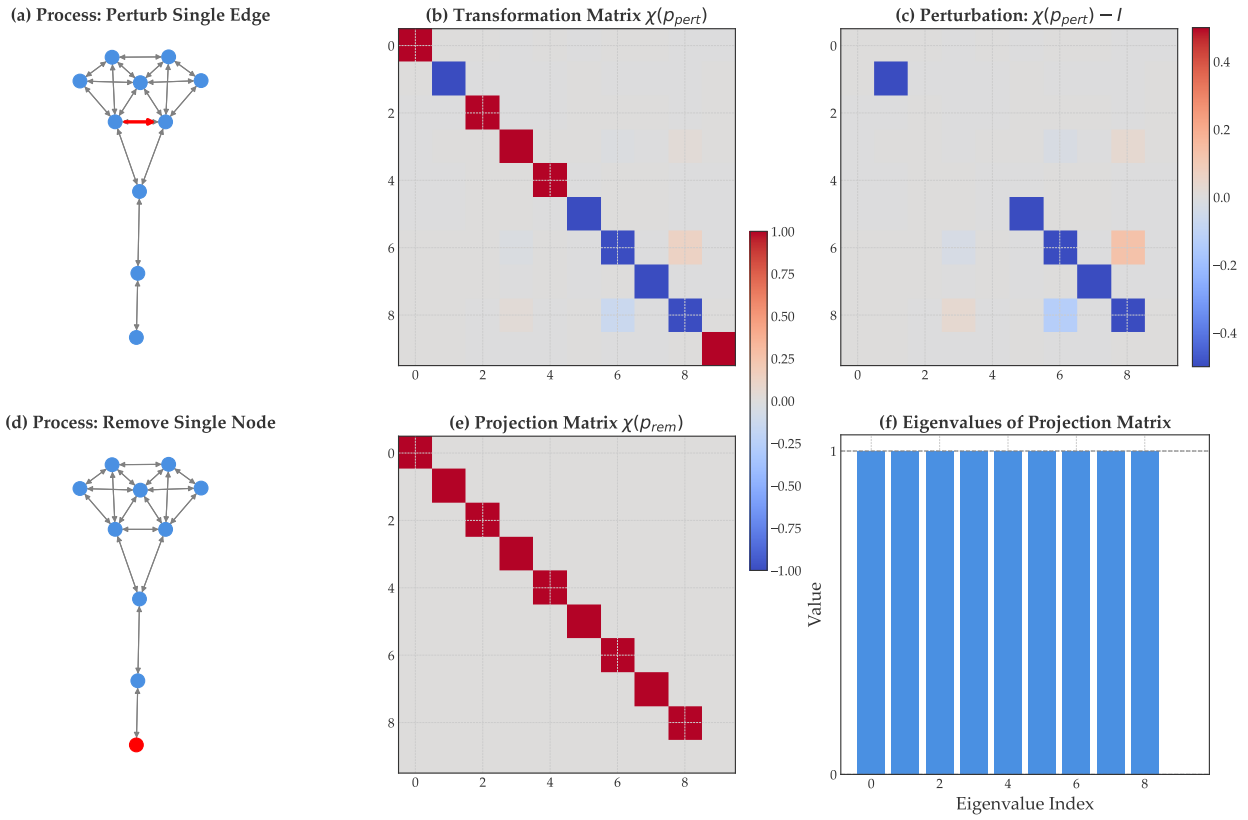

Figure 5: **Numerical validation of the process interpretation toolkit (Theorems 4, 5, and 6).** This figure demonstrates that different elementary processes have unique, identifiable spectral signatures, confirming the diagnostic power of our framework. **Top Row**: Signature of a Local Perturbation. **(a)** The physical process: a single edge (highlighted in red) on a network is perturbed. **(b)** The resulting transformation matrix $\chi(p_{pert})$. The matrix is strongly diagonally dominant, visually confirming the principle of Structural Inertia (Theorem 5). The faint off-diagonal entries correctly show that even a local change has subtle, non-local effects on the global geometry. **(c)** The perturbation matrix, $\chi(p_{pert}) - I$. The sparse, structured nature of this matrix is the visual signature of a simple, local change, providing visual proof of the Rank-One Update signature (Theorem 4). **Bottom Row**: Signature of Node Removal. **(d)** The physical process: a single node (highlighted in red) is removed from the network. **(e)** The matrix of the resulting projection operator. Its signature—an identity matrix with a single zero on the diagonal—is visually distinct from the perturbation signature in (b). **(f)** The eigenvalues of the projection matrix. The bar plot provides a quantitative validation of the Node Removal Signature (Theorem 6), showing exactly one eigenvalue at 0 (for the removed dimension) and all others at 1.

The bottom row validates our Spectral Sensitivity theorem (Theorem 2). We subject a barbell graph, which has a clear structural bottleneck, to a structurally fragmenting process. As predicted, the Fiedler value (λ_2), a measure of algebraic connectivity, shows a clear and monotonic decline (Panel d). The physical meaning of this is made clear in Panel (e), which shows the network visually breaking apart. The full spectrum dynamics in Panel (f) provide a deeper insight, showing the Fiedler value "peeling away" from the rest of the spectrum, a classic signature of fragmentation.

5.3 Equivalence of Physical and Spectral Stability

Another theoretical result of our framework is the equivalence between the physical stability of a system and the stability of its spectral representation. To validate this, we simulate a network's long-term evolution

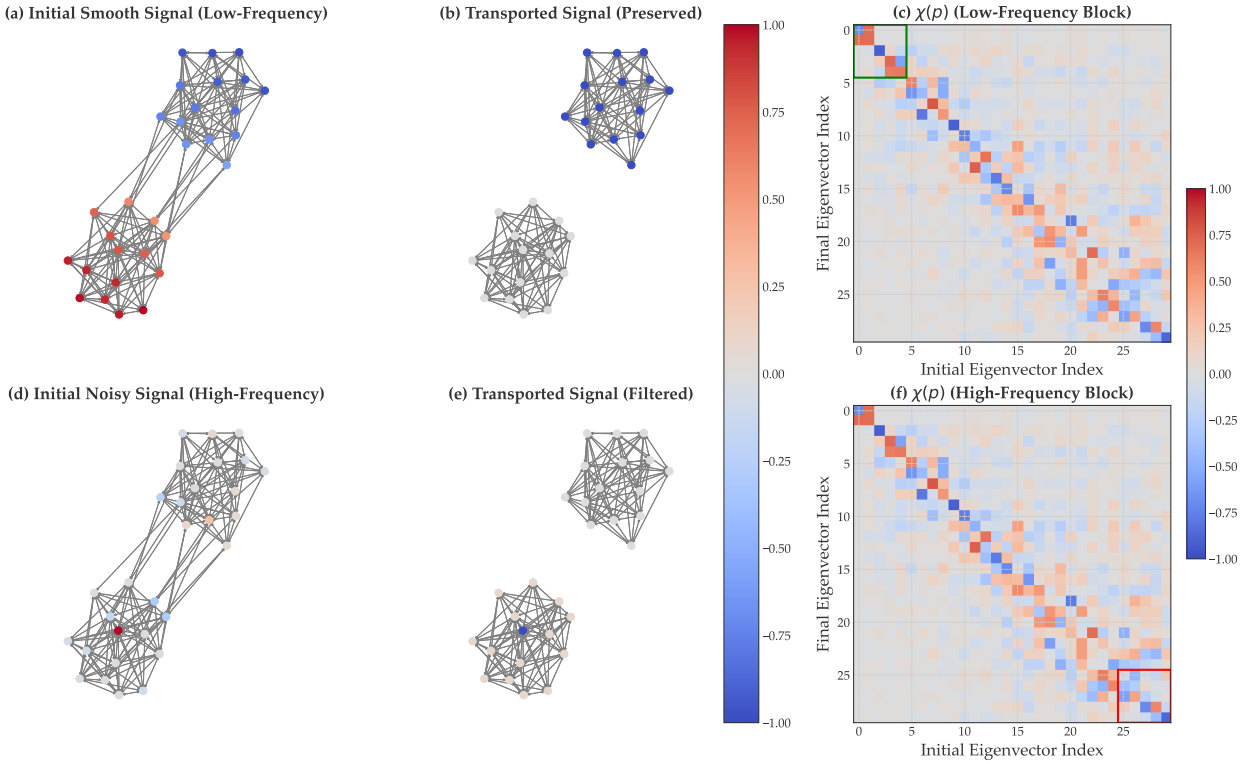

Figure 6: **Demonstration of Signal Transport as a Spectral Filter (Theorem 7).** This experiment reveals that a major structural change acts as a low-pass filter on network patterns, preserving smooth, low-frequency signals while destroying noisy, high-frequency ones. The process in both scenarios is the severing of all bridge edges between the two communities of a Stochastic Block Model graph. **Top Row**: Preservation of a Low-Frequency Signal. **(a)** The initial signal is a smooth, low-frequency pattern (the Fiedler eigenvector of the initial graph), creating a clear gradient between the two communities. **(b)** After transport to the new, disconnected graph, the smooth pattern is closely preserved within each community. **(c)** The transformation matrix $\chi(p)$ provides the mechanism. The highlighted top-left (low-frequency) block is strongly diagonal, proving that low-frequency modes of the initial graph map cleanly to the low-frequency modes of the final graph, ensuring pattern preservation. **Bottom Row**: Filtering of a High-Frequency Signal. **(d)** The initial signal is a noisy, high-frequency pattern (the last eigenvector of the initial graph). **(e)** After transport, the noisy pattern is destroyed and filtered into a closely uniform, low-energy state. The fine-grained structure needed to support the pattern is removed. **(f)** The transformation matrix $\chi(p)$ again reveals the mechanism. The highlighted bottom-right (high-frequency) block is scattered and non-diagonal. This visually demonstrates that high-frequency modes of the initial graph do not map to high-frequency modes of the final graph; their energy is scattered, destroying the original pattern.

under a dissipative process with edge pruning, allowing it to converge to a stable equilibrium. The results, shown in Figure 4, provide a multi-faceted confirmation of the theorem.

Panel (a) is the core quantitative proof. It shows that the physical distance from equilibrium (blue, solid) and the spectral distance (red, dashed) both decay to zero, confirming that the system stabilises if and only if its spectral representation does. The differing decay profiles reveal a non-trivial insight: the spectral representation is a more sensitive indicator of the final convergence. Panel (b) provides a holistic view, showing the entire spectrum of the system converging to stable, horizontal lines. Finally, Panel (c) provides the intuitive physical implication, showing the network simplifying from a dense, chaotic initial state to a sparse, stable "backbone" structure.

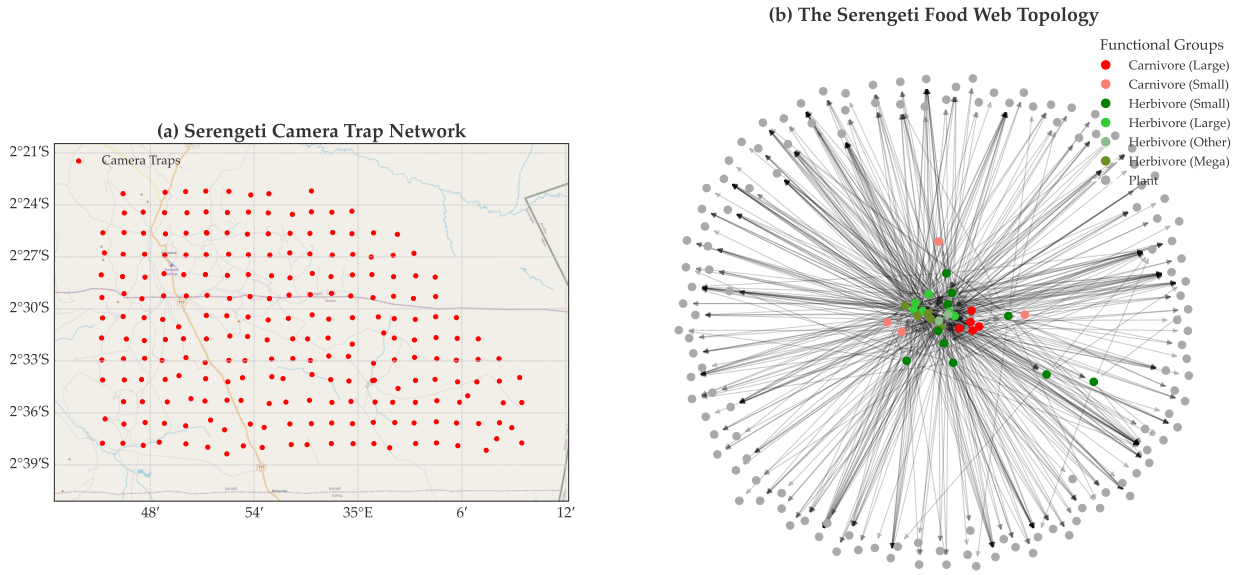

Figure 7: **The Serengeti Case Study Area.** This figure provides an overview of the two primary datasets used in our real-world case study. **(a)** A map of the Serengeti National Park showing the geographic distribution of the 225 camera traps from the Snapshot Serengeti project (Swanson et al., 2015). The camera grid covers an area of 1,125 km^2, and each point represents a unique camera site, plotted using its UTM coordinates. The basemap provides topographical context for the study area. **(b)** The full topological structure of the Serengeti food web, based on the data from Baskerville et al. (2011). The network consists of 161 species (nodes) and 592 directed feeding links (edges). Nodes are colour-coded by their functional group to illustrate the ecosystem's trophic structure: carnivores are shown in reds, herbivores in greens, and the broad plant base in grey. The high proportion of plant nodes reflects the high taxonomic resolution of the primary producer data, a key feature of this food web.

5.4 Diagnostic Power of the Toolkit

A key claim of our work is that elementary processes have unique, identifiable spectral signatures. Figure 5 validates this diagnostic power. The top row shows that a simple, local edge perturbation results in a transformation matrix that is strongly diagonally dominant (Panel b) and represents a sparse, low-rank perturbation of the identity matrix (Panel c). This confirms that local changes have a simple, structured spectral signature. In contrast, the bottom row shows that a major topological change—a node removal—produces a completely different signature. The transformation is a projection operator (Panel e), which is confirmed by its eigenvalue spectrum of only 0s and 1s (Panel f). This demonstrates that our framework can unambiguously distinguish between different classes of physical events.

5.5 Signal Transport as a Spectral Filter

Our final synthetic experiment demonstrates that our framework can predict the fate of patterns on a changing network. The results in Figure 6 show that a major structural change acts as a spectral low-pass filter. The top row shows the result of pattern preservation. A smooth, low-frequency signal (Panel a) is successfully transported to the new network with its structure closely preserved (Panel b). The transformation matrix $\chi(p)$ (Panel c) reveals the mechanism: its top-left, low-frequency block is strongly diagonal, showing that the geometric language for smooth patterns is robust to the change.

The bottom row visualises pattern destruction. A noisy, high-frequency signal (Panel d) is almost completely destroyed by the transport, collapsing to a closely uniform, low-energy state (Panel e). The transformation matrix (Panel f) again reveals why: its bottom-right, high-frequency block is scattered and non-diagonal,

proving that the geometric language for noisy patterns has been fundamentally broken by the process. This experiment confirms that our framework provides a tool for predicting not just *that* a pattern will change, but *how* it will change, based on its alignment with the network's underlying spectral geometry.

6 Real-World Case Study

Having established the theoretical foundations of our `Proc-to-Spec` framework in §4 and validated its core properties with numerical experiments in §5, we now apply it to a complex, real-world system. The objective of this case study is to move from mathematical verification to scientific discovery, demonstrating our framework's power as a tool for generating mechanistic and predictive insights from high-resolution, inherently noisy ecological data.

Table 2: Summary of the Serengeti 'Snapshot' and Food Web Datasets.

Metric	**Value**	**Source**
Snapshot Serengeti Camera Trap Data		
Study Area	1,125 km^2	Swanson et al. (2015)
Number of Camera Traps	225	Swanson et al. (2015)
Time Period	2010–2013	Swanson et al. (2015)
Total Image Sets Classified	1.2 million	Swanson et al. (2015)
Food Web Topology Data		
Total Species (Nodes)	161	Baskerville et al. (2011)
Feeding Links (Edges)	592	Baskerville et al. (2011)
Functional Groups		
Carnivores	9	Baskerville et al. (2011)
Herbivores	23	Baskerville et al. (2011)
Plants	129	Baskerville et al. (2011)

6.1 The Serengeti Food Web Ecosystem

We have chosen the Serengeti ecosystem in northern Tanzania to be the testbed for our framework. It is a well-studied, resource-constrained system characterised by strong seasonal dynamics and occasional environmental shocks. Its complex food web and the availability of rich, modern datasets provide an valuable opportunity to test our framework's ability to analyse real-world network dynamics.

Our analysis is built upon two public datasets. The first is the "*Snapshot Serengeti*" dataset, a large collection of 1.2 million classified camera trap image sets from 225 camera traps spread across 1,125 km^2 of the Serengeti National Park from 2010 to 2013 (Swanson et al., 2015). This dataset provides a high-frequency, event-level record of animal activity. The second is a high-resolution food web topology, which details 592 predator-prey interactions among 161 species, notable for its high taxonomic resolution at the plant level (Baskerville et al., 2011). An overview of these datasets is provided in Figure 7 and Table 2.

6.2 From Camera Traps to Dynamic Networks

To apply our framework, we developed a pipeline to transform the raw, event-level camera trap data (Swanson et al., 2015) into a time-series of dynamic networks in the `Proc` category.

1. **Dynamic Aggregation:** We first parsed and aggregated the raw camera trap events from the consensus dataset to produce a monthly time-series of total observed counts for each species.

2. **Network Construction:** For each recorded month from 2010 to 2013, we then constructed a dynamic network. The nodes and static edge topology were taken from the food web data (Baskerville

et al., 2011). We calculated the dynamic edge weights $W(u, v)$ for each predator-prey interaction using a well-established mass-action model (Murray, 2007), where the weight is proportional to the product of the predator and prey counts for that month: $W(u, v) = C \cdot \text{count}(u) \cdot \text{count}(v)$.

3. **Handling Disconnectedness:** Real-world ecological data is inherently sparse and noisy. In our monthly networks, species with no observed activity are effectively disconnected from the system. For such disconnected graphs, it is a standard and well-established practice in network science to focus the analysis on the largest connected component (LCC), which contains the core of the network's structure and dynamics (Newman, 2018). Therefore, we performed all subsequent spectral analyses on the LCC of the active subgraph for each month. This standard approach (Newman, 2018) allows us to robustly analyse the dynamics of the core, active portion of the ecosystem in any given month, focusing on the central arena of species interactions while being resilient to the transient disappearance of peripheral species from the camera trap dataset.

6.3 Experimental Results

Using the constructed monthly networks, here we conducted a series of experiments with each designed to validate a different core theorem of our framework.

Detecting the Ecosystem's Pulse and a Major Ecological Shock. Our first experiment tests the framework's sensitivity to environmental change. As shown in Figure 8, the analysis reveals a clear, cyclical pattern in the ecosystem's structure corresponding to the wet and dry seasons. The size of the active core (Panel a) and the network's algebraic connectivity (Panel b, Fiedler value) both consistently contract during the resource-scarce dry seasons. The framework also proves highly sensitive to extreme events, detecting a major anomaly in July 2011 that corresponds to a severe, well-documented regional drought. The framework provides a multi-faceted signature of this crisis: a collapse of the active core (Panel a), preceded by the most significant structural fragmentation event in the time-series (Panel c), which resulted in a small, hyper-connected remnant network (Panel b).

Verifying Physical Conservation Laws in the Spectral Domain. To validate the framework's physical grounding, we tested the Spectral Trace Conservation Law. Figure 9 plots the change in total observed animal activity (our proxy for the system's total resource, $\Delta\mathcal{R}$) against the change in the trace of the symmetrised Laplacian ($\Delta\text{Tr}(L_{sym})$) for each month-to-month transition. As predicted by Theorem 1, the two quantities exhibit a near-perfect linear relationship ($r = 0.9553, p < 0.0001$). It is important to note that this strong correlation is *not* a noisy, empirical biological finding, but rather a direct validation of the mathematical consistency of our framework. The result confirms that the abstract spectral quantity we defined (the Laplacian trace) behaves exactly as predicted by the physical quantity it is designed to represent (the total system resource, here modelled by camera-recorded animal activity). This grounds our functorial framework in the physical reality of the system.

The Geometric Signature of a Crisis. We then leveraged the Process Interpretation Toolkit to move beyond detecting change to characterising its fundamental nature. Figure 10 compares the geometric signature ($\chi(p)$ matrix) of the 2011 drought collapse to that of a typical seasonal transition. The results are visually striking. A typical seasonal change (Panel c) is a minor perturbation, with a signature close to the identity matrix, and its core community structure is preserved (Panel f). In contrast, the drought collapse (Panel a) has a complex, highly off-diagonal signature, corresponding to a topologically complex shattering of the Fiedler vector (Panel d).

The species-level analysis (Row 3) provides a concrete ecological interpretation. The pre-drought Fiedler vector (Panel g) is defined by the classic partition between migratory herds (e.g., *Connochaetes taurinus*) and their resident predators (e.g., *Crocuta crocuta*). The shattering of this vector is the geometric signature of a well-documented ecological phenomenon: severe droughts force migratory herds to break their normal patterns, temporarily destroying the predictable spatial predator-prey dynamics that define the ecosystem's structure (Sinclair et al., 2007).

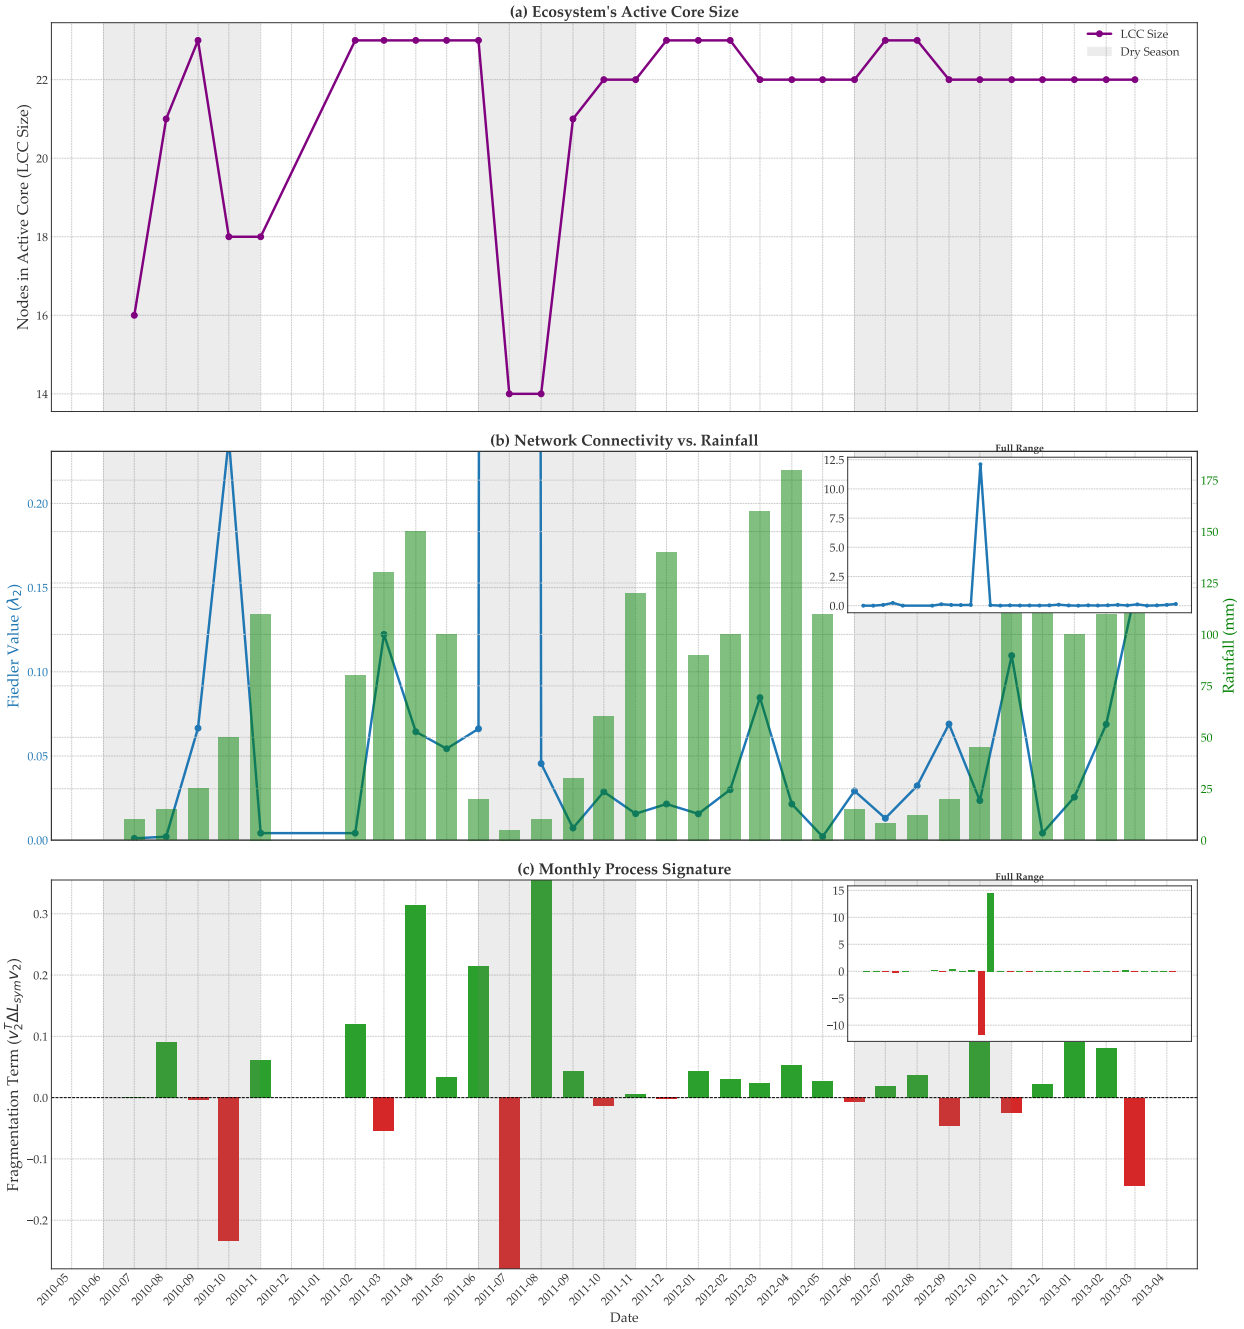

Figure 8: **Spectral Dynamics of the Serengeti Food Web (2010-2013).** The figure presents three metrics derived from the monthly network time-series, revealing a roughly consistent seasonal cycle and a major anomalous event. **(a) Active Core Size:** The number of species in the Largest Connected Component (LCC) shows a seasonal "pulse", vibrating between dry (shaded) and wet seasons. **(b) Network's Algebraic Connectivity:** The Fiedler value (λ_2) of the LCC is approximately anti-correlated with the dry seasons. The inset shows the full range, capturing a massive connectivity spike of the shrunken core during the July 2011 drought. **(c) Process Signature:** The fragmentation term reveals fragmenting processes (red bars) at the onset of dry seasons and consolidating processes (green bars) during recovery. The inset highlights the extreme nature of the 2011 event. Collectively, these results show that the framework is sensitive enough to detect both the subtle, recurring seasonal cycle and the multi-faceted signature of a major, documented drought event.

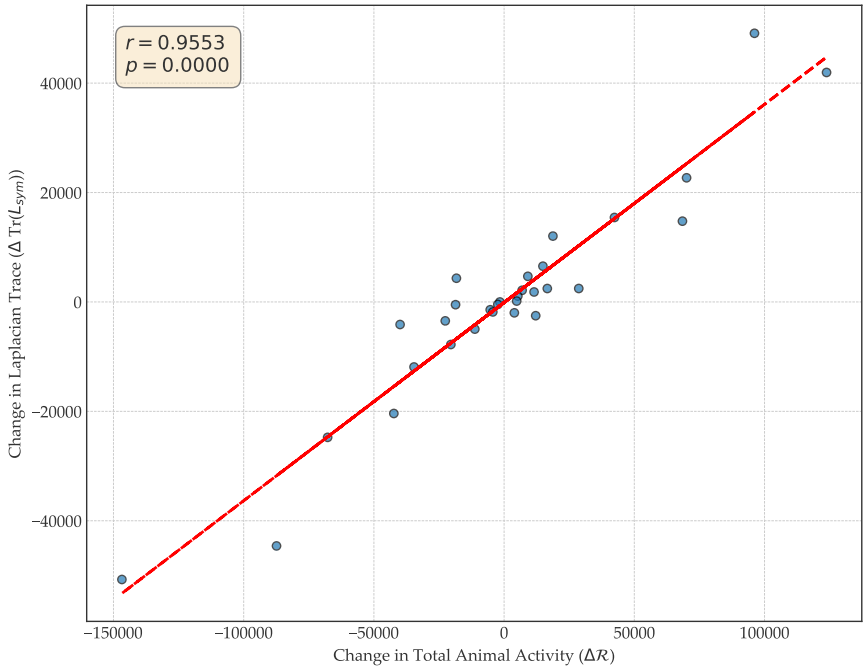

Figure 9: **Real-world Validation of the Spectral Trace Conservation Law (Theorem 1) on the Serengeti Food Web (2010-2013).** This figure provides a direct, quantitative test of the predicted relationship between physical resource dynamics and their spectral counterparts. Each point represents a month-to-month transition in the Serengeti ecosystem. The x-axis shows the change in the total observed animal count, our proxy for the total system resource ($\Delta\mathcal{R}$). The y-axis shows the corresponding change in the trace of the symmetrised Laplacian ($\Delta\text{Tr}(L_{sym})$). The data reveals a strong, positive, and linear correlation between the physical and spectral quantities, as predicted by our framework. This relationship is quantified by a Pearson correlation coefficient of $r = 0.9553$ ($p < 0.0001$), providing compelling, real-world evidence that the physical conservation of resources within the system is directly reflected by this spectral invariant.

Predicting Ecological Winners and Losers. Our final experiment demonstrates the framework's predictive power. We defined a signal of species' "importance" (eigenvector centrality) on the stable, pre-drought network. We then used the transport matrix ($T_{\text{transport}}$) derived from the drought process to generate a falsifiable prediction for the new importance of each species in the drought-stricken state. Figure 11 visualises these predictions. The framework predicts that large grazers (e.g., *Eudorcas thomsonii*) and their specialist predators (*Panthera leo*) will be the primary "losers", while more adaptable, opportunistic carnivores (e.g., *Crocuta crocuta*) will maintain their central role. These predictions are supported by the ecological literature. Severe droughts are known to disrupt the primary food source for large grazers and, consequently, the specialist predators that depend on them. For instance, drought conditions in the Serengeti are known to fundamentally alter the dynamics of the wildebeest-predator system, affecting migration, foraging, and predation patterns (Sinclair et al., 2007). In contrast, adaptable scavengers and predators are well-positioned to capitalise on the increased number of carcasses and weakened prey during such events.

7 Discussion and Conclusion

The analysis of dynamic networks has traditionally focused on characterising sequences of states, often leaving the transformations that drive the evolution between them as unformalised, black-box processes. In this work, we introduced `Proc-to-Spec`, a new, category-theoretical framework that provides a principled and interpretable language for the processes of change themselves. Our central contribution is the construction of

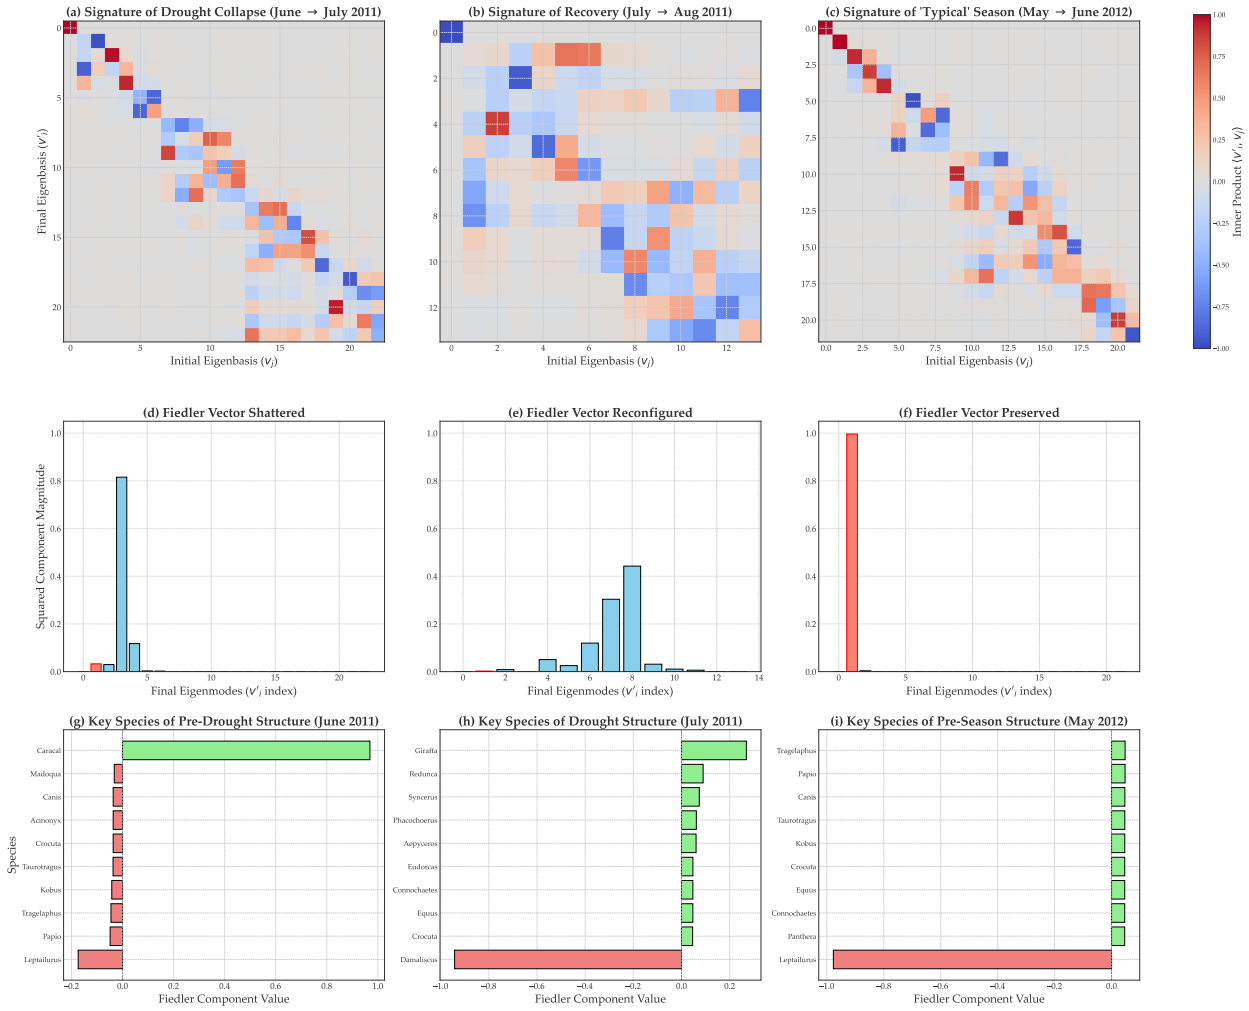

Figure 10: **The Geometric-Spectral Signature and Species-Level Impact of Ecological Processes of the Serengeti Food Web (2010-2013).** This figure provides a multi-layered validation of our framework's diagnostic power by comparing the geometric signature ($\chi(p)$ matrix) and structural impact of three distinct ecological processes. **Top Row (Panels a-c):** The heatmaps show the overall geometric signature of each process. A "typical" seasonal transition from wet to dry in 2012 (**c**) has a signature that is strongly diagonal, indicating a minor perturbation that preserves the system's core structure. In contrast, the 2011 drought collapse (**a**) and subsequent recovery (**b**) are highly off-diagonal, the signatures of catastrophic structural reorganisations. **Middle Row (Panels d-f):** These panels reveal the mechanism behind the signatures by showing the transformation of the Fiedler vector (v_2), which represents the ecosystem's primary community structure. During the typical transition (**f**), the Fiedler vector is preserved, mapping almost entirely onto the new Fiedler vector (red bar). During the drought collapse (**d**), this vector is shattered, its energy scattered across many new structural modes. **Bottom Row (Panels g-i):** This provides a concrete, species-level interpretation. The Fiedler vector of the pre-drought network (**g**) and the typical network (**i**) is defined by the classic ecological partition of the Serengeti: the large migratory herds (e.g., *Connochaetes taurinus*, *Equus quagga*) on one side, and their primary resident predators (e.g., *Crocuta crocuta*, *Panthera leo*) on the other. The shattering of this vector during the drought is the geometric signature of a well-documented ecological phenomenon: severe droughts force migratory herds to break their normal patterns in search of scarce resources, thus temporarily destroying the predictable spatial predator-prey dynamics that define the ecosystem's structure (Sinclair et al., 2007). The framework has detected and characterised this real-world crisis at the species level.

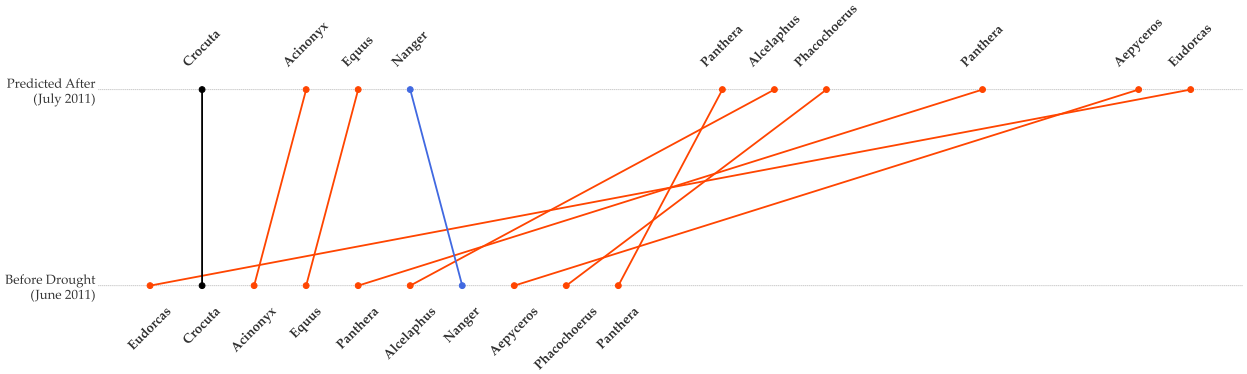

Figure 11: **A Falsifiable Prediction of Ecological "Winners" and "Losers" of the 2011 Drought using Signal Transport (Theorem 7).** This figure demonstrates the framework's predictive power. The signal is a species' importance, defined by its eigenvector centrality in the stable, pre-drought network of June 2011. This signal is then projected through the drought process using the transport matrix $T_{\text{transport}} = (\chi(p)^T)^{-1}$ to predict the new centrality of each species in the drought-stricken state of July 2011. The slope chart visualises the predicted change in rank for the top 10 most central species. The framework predicts that large grazers, such as Thomson's Gazelle (*Eudorcas*) and Hartebeest (*Alcelaphus*), and their specialist apex predator, the Lion (*Panthera*), will be the primary "losers", suffering a significant drop in their relative importance. Conversely, it predicts that more adaptable, opportunistic carnivores like the Spotted Hyena (*Crocuta*) will be the primary "winners", maintaining or improving their central role. These predictions are supported by ecological literature. Severe droughts are known to disrupt the primary food source for large grazers and, consequently, the specialist predators that depend on them (Sinclair et al., 2007). In contrast, adaptable scavengers and predators like hyenas are well-positioned to capitalise on the increased number of carcasses and weakened prey during such events. The results validate that the Signal Transport Theorem can function as a quantitative scientific tool for anticipating how ecological roles are reshuffled under major perturbations such as drought.

a spectral functor, χ, that maps physical processes in a source category `Proc` to unique linear transformations between spectral eigenspaces in a target category `Spec`.

7.1 Summary of Key Findings

Our theoretical claims were validated through a two-pronged approach. First, a suite of numerical experiments provided a rigorous, controlled validation of each of our core theorems (§5), confirming the mathematical soundness of the framework. Second, a comprehensive case study of the Serengeti ecosystem demonstrated the framework's power and sensitivity on a complex, noisy, real-world dataset (§6).

Through the Serengeti case study, we demonstrated that:

- **(Theorem 1)** An abstract spectral invariant—the Laplacian trace—is rigorously and directly coupled to a physical quantity in the system, the total observed animal activity, validating the framework's physical grounding.
- **(Theorem 2)** The framework is sensitive enough to detect both the subtle, cyclical pulse of seasonal change and the unique signature of a catastrophic, real-world drought event.
- **(Process Interpretation Toolkit)** The geometric signature ($\chi(p)$ matrix) provides a powerful diagnostic tool, revealing that a major drought has a complex, off-diagonal signature corresponding to a shattering of the ecosystem's core predator-prey structure, fundamentally different from the near-identity signature of a typical seasonal change.

- **(Theorem 7)** The framework has predictive power. The Signal Transport theorem was used to make a set of ecologically meaningful predictions about which species would become functional "winners" and "losers" during the drought.

7.2 Broader Implications: The Geometry of Change

The `Proc-to-Spec` framework offers more than a new set of analytical tools; it offers a new lens through which to view the dynamics of complex systems. By focusing on the geometry of the transformations, we move beyond simply describing that a system has changed to providing a mechanistic, interpretable signature for how it has changed.

This provides a "glass box" alternative to many contemporary machine learning models. While a graph neural network might predict a future network state with high accuracy, our framework is designed to provide causal insight into the underlying process, characterising its nature (e.g., a local perturbation vs. a global reorganisation). This process-centric viewpoint has broad potential applicability in other domains where dynamic networks are central, such as systems biology (gene regulatory networks), neuroscience (brain connectivity), or economics (financial networks).

7.3 Limitations and Future Work

Here, we discuss the theoretical and practical limitations of our work as well as suggest future research directions to overcome those.

Theoretical Limitations. The category-theoretical foundations of `Proc-to-Spec`, while rigorous, are built upon specific assumptions that define the scope of the current framework. Future work should aim to generalise these foundations.

- **Preservation of the Node Set:** The functorial mapping χ is formally defined for morphisms in `Proc` that represent processes preserving the set of nodes. Our current analysis of node removal (Theorem 6), while providing a clear geometric interpretation as a projection, exists outside this primary functorial definition. A key avenue for future work is to extend the framework to a richer category that can formally accommodate morphisms that change the dimension of the underlying vector space, perhaps by employing concepts from persistence homology or sheaf theory to track topological features across changing dimensions (Perea & Harer, 2015).
- **Choice of Spectral Representation:** Our framework is built upon the symmetrised Laplacian. While this operator has many desirable properties—including a real spectrum and an orthonormal eigenbasis—other matrix representations, such as the random walk Laplacian, the adjacency matrix, or higher-order Laplacian (Nurisso et al., 2025), capture different aspects of network dynamics. Future theoretical work could explore the construction of parallel functors for these different spectral representations.

Practical Limitations. The application of our framework to the Serengeti case study highlighted several practical challenges that offer opportunities for future refinement and application.

- **Modelling of Edge Weights:** The insights generated are contingent on our choice of a mass-action model for dynamic edge weights. While this is a standard and justified model (Murray, 2007), it is a simplification of complex ecological interactions. Future work could explore more sophisticated, domain-specific models for interaction strength, potentially incorporating the non-linear dependencies and external environmental variables that are known to modulate the strength of species interactions (Tylianakis et al., 2008).
- **Data Sparsity and the LCC:** The case study demonstrated that real-world ecological data is often sparse, leading to monthly network snapshots that are disconnected. Our robust method of analysing the largest connected component (LCC) was effective and scientifically sound. However,

this points to a broader need for spectral theories that can gracefully handle and perhaps even draw insight from the natural fragmentation and disconnectedness of real-world systems, rather than treating it as a feature to be isolated (Banerjee, 2021).

7.4 Conclusion

In this paper, we introduced `Proc-to-Spec`, a functorial framework for the analysis of dynamic networks. Our work is motivated by a fundamental limitation in the traditional study of such systems, which has largely focused on characterising sequences of static states, often leaving the transformations that drive the evolution between them as unformalised, black-box processes. By shifting the analytical focus from states to the processes themselves, we have developed a principled and powerful toolkit for understanding how complex systems evolve. Through a suite of rigorous numerical experiments, we first validated the mathematical soundness of our core theorems in controlled settings. We then demonstrated the framework's real-world applicability in a comprehensive case study of the Serengeti ecosystem, using high-resolution camera trap data. This work opens new avenues for a more mechanistic, interpretable, and geometric understanding of the dynamic networks that pervade both the natural and social sciences. By providing a "glass box" that reveals the geometry of change, our framework represents a step towards a deeper and more predictive science of complex, interconnected systems.

References

Stefano Allesina and Si Tang. Stability criteria for complex ecosystems. *Nature*, 483(7388):205–208, 2012.

Uri Alon. Network motifs: theory and experimental approaches. *Nature Reviews Genetics*, 8(6):450–461, 2007.

John Baez and Mike Stay. Physics, topology, logic and computation: a rosetta stone. In *New Structures for Physics*, pp. 95–172. Springer, 2010.

Anirban Banerjee. On the spectrum of hypergraphs. *Linear Algebra and its Applications*, 614:82–110, 2021.

Albert-László Barabási and Réka Albert. Emergence of scaling in random networks. *Science*, 286(5439): 509–512, 1999.

Albert-Laszlo Barabasi and Zoltan N Oltvai. Network biology: understanding the cell's functional organization. *Nature Reviews Genetics*, 5(2):101–113, 2004.

Baruch Barzel and Albert-László Barabási. Network link prediction by global silencing of indirect correlations. *Nature Biotechnology*, 31(8):720–725, 2013.

Jordi Bascompte, Pedro Jordano, Carlos J Melián, and Jens M Olesen. The nested assembly of plant–animal mutualistic networks. *Proceedings of the National Academy of Sciences*, 100(16):9383–9387, 2003.

Edward B Baskerville, Andy P Dobson, Trevor Bedford, Stefano Allesina, T Michael Anderson, and Mercedes Pascual. Spatial guilds in the serengeti food web revealed by a bayesian group model. *PLoS Computational Biology*, 7(12):e1002321, 2011.

Mike Behrisch, Sebastian Kerkhoff, Reinhard Pöschel, Friedrich Martin Schneider, and Stefan Siegmund. Dynamical systems in categories. *Applied Categorical Structures*, 25(1):29–57, 2017.

Mikhail Belkin and Partha Niyogi. Laplacian eigenmaps for dimensionality reduction and data representation. *Neural Computation*, 15(6):1373–1396, 2003.

Eric L Berlow, Anje-Margiet Neutel, Joel E Cohen, Peter C De Ruiter, BO Ebenman, Mark Emmerson, Jeremy W Fox, Vincent AA Jansen, J Iwan Jones, Giorgos D Kokkoris, et al. Interaction strengths in food webs: issues and opportunities. *Journal of Animal Ecology*, pp. 585–598, 2004.

Rajendra Bhatia. Review of matrix perturbation theory: by gw stewart and ji-guang sun. *Linear Algebra and its Applications*, 160:255–259, 1992.

Ginestra Bianconi and Albert-László Barabási. Bose-einstein condensation in complex networks. *Physical Review Letters*, 86(24):5632, 2001.

Michael M Bronstein, Joan Bruna, Yann LeCun, Arthur Szlam, and Pierre Vandergheynst. Geometric deep learning: going beyond euclidean data. *IEEE Signal Processing Magazine*, 34(4):18–42, 2017.

Ed Bullmore and Olaf Sporns. Complex brain networks: graph theoretical analysis of structural and functional systems. *Nature Reviews Neuroscience*, 10(3):186–198, 2009.

Arnaud Casteigts, Paola Flocchini, Walter Quattrociocchi, and Nicola Santoro. Time-varying graphs and dynamic networks. *International Journal of Parallel, Emergent and Distributed Systems*, 27(5):387–408, 2012.

Fan RK Chung. *Spectral graph theory*, volume 92. American Mathematical Soc., 1997.

Eric H Davidson, Jonathan P Rast, Paola Oliveri, Andrew Ransick, Cristina Calestani, Chiou-Hwa Yuh, Takuya Minokawa, Gabriele Amore, Veronica Hinman, Cesar Arenas-Mena, et al. A genomic regulatory network for development. *Science*, 295(5560):1669–1678, 2002.

Jennifer A Dunne, Richard J Williams, and Neo D Martinez. Network structure and biodiversity loss in food webs: robustness increases with connectance. *Ecology Letters*, 5(4):558–567, 2002.

Miroslav Fiedler. Algebraic connectivity of graphs. *Czechoslovak Mathematical Journal*, 23(2):298–305, 1973.

Brendan Fong and David I Spivak. Seven sketches in compositionality: An invitation to applied category theory. *arXiv preprint arXiv:1803.05316*, 2018.

Neil Ghani, Jules Hedges, Viktor Winschel, and Philipp Zahn. Compositional game theory. In *Proceedings of the 33rd Annual ACM/IEEE Symposium on Logic in Computer Science*, pp. 472–481, 2018.

Justin Gilmer, Samuel S Schoenholz, Patrick F Riley, Oriol Vinyals, and George E Dahl. Neural message passing for quantum chemistry. In *International Conference on Machine Learning*, pp. 1263–1272. PMLR, 2017.

Marco Gori, Gabriele Monfardini, and Franco Scarselli. A new model for learning in graph domains. In *Proceedings of the 2005 IEEE International Joint Conference on Neural Networks*, volume 2, pp. 729–734. IEEE, 2005.

Aditya Grover and Jure Leskovec. node2vec: Scalable feature learning for networks. In *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pp. 855–864, 2016.

Alexis Guyot, Annabelle Gillet, Eric Leclercq, and Nadine Cullot. A formal framework for data lakes based on category theory. In *Proceedings of the 26th International Database Engineered Applications Symposium*, pp. 75–83, 2022.

Will Hamilton, Zhitao Ying, and Jure Leskovec. Inductive representation learning on large graphs. *Advances in Neural Information Processing Systems*, 30, 2017.

Ilkka Hanski. Metapopulation dynamics. *Nature*, 396(6706):41–49, 1998.

Mohammad Al Hasan and Mohammed J Zaki. A survey of link prediction in social networks. In *Social Network Data Analytics*, pp. 243–275. Springer, 2011.

Christopher J Honey, Rolf Kötter, Michael Breakspear, and Olaf Sporns. Network structure of cerebral cortex shapes functional connectivity on multiple time scales. *Proceedings of the National Academy of Sciences*, 104(24):10240–10245, 2007.

Hawoong Jeong, Sean P Mason, A-L Barabási, and Zoltan N Oltvai. Lethality and centrality in protein networks. *Nature*, 411(6833):41–42, 2001.

Seyed Mehran Kazemi, Rishab Goel, Kshitij Jain, Ivan Kobyzev, Akshay Sethi, Peter Forsyth, and Pascal Poupart. Representation learning for dynamic graphs: A survey. *Journal of Machine Learning Research*, 21(70):1–73, 2020.

Sonia Kéfi, Vishwesha Guttal, William A Brock, Stephen R Carpenter, Aaron M Ellison, Valerie N Livina, David A Seekell, Marten Scheffer, Egbert H Van Nes, and Vasilis Dakos. Early warning signals of ecological transitions: methods for spatial patterns. *PloS One*, 9(3):e92097, 2014.

David Kempe, Jon Kleinberg, and Amit Kumar. Connectivity and inference problems for temporal networks. In *Proceedings of the 32nd Annual ACM Symposium on Theory of Computing*, pp. 504–513, 2000.

David Kempe, Jon Kleinberg, and Éva Tardos. Maximizing the spread of influence through a social network. In *Proceedings of the Ninth ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pp. 137–146, 2003.

TN Kipf. Semi-supervised classification with graph convolutional networks. *arXiv preprint arXiv:1609.02907*, 2016.

Yehuda Koren. On spectral graph drawing. In *International Computing and Combinatorics Conference*, pp. 496–508. Springer, 2003.

Lauri Kovanen, Márton Karsai, Kimmo Kaski, János Kertész, and Jari Saramäki. Temporal motifs in time-dependent networks. *Journal of Statistical Mechanics: Theory and Experiment*, 2011(11):P11005, 2011.

Matthieu Latapy, Tiphaine Viard, and Clémence Magnien. Stream graphs and link streams for the modeling of interactions over time. *Social Network Analysis and Mining*, 8(1):61, 2018.

Kristina Lerman, Rumi Ghosh, and Jeon Hyung Kang. Centrality metric for dynamic networks. In *Proceedings of the 8th Workshop on Mining and Learning with Graphs*, pp. 70–77, 2010.

Geert Leus, Antonio G Marques, José MF Moura, Antonio Ortega, and David I Shuman. Graph signal processing: History, development, impact, and outlook. *IEEE Signal Processing Magazine*, 40(4):49–60, 2023.

Qimai Li, Zhichao Han, and Xiao-Ming Wu. Deeper insights into graph convolutional networks for semi-supervised learning. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 32, 2018.

David Liben-Nowell and Jon Kleinberg. The link prediction problem for social networks. In *Proceedings of the Twelfth International Conference on Information and Knowledge Management*, pp. 556–559, 2003.

Fanzhen Liu, Jia Wu, Shan Xue, Chuan Zhou, Jian Yang, and Quanzheng Sheng. Detecting the evolving community structure in dynamic social networks. *World Wide Web*, 23(2):715–733, 2020.

Franco Manessi, Alessandro Rozza, and Mario Manzo. Dynamic graph convolutional networks. *Pattern Recognition*, 97:107000, 2020.

Naoki Masuda and Renaud Lambiotte. *A guide to temporal networks*. World Scientific, 2016.

Robert M May. *Stability and complexity in model ecosystems*, volume 6. Princeton University Press, 2001.

Bojan Mohar, Y Alavi, G Chartrand, and Ortrud Oellermann. The laplacian spectrum of graphs. *Graph Theory, Combinatorics, and Applications*, 2(871-898):12, 1991.

José M Montoya, Stuart L Pimm, and Ricard V Solé. Ecological networks and their fragility. *Nature*, 442 (7100):259–264, 2006.

James D Murray. *Mathematical biology: I. An introduction*, volume 17. Springer Science & Business Media, 2007.

Mark Newman. *Networks*. Oxford University Press, 2018.

Mark EJ Newman. Spread of epidemic disease on networks. *Physical Review E*, 66(1):016128, 2002.

Giang Hoang Nguyen, John Boaz Lee, Ryan A Rossi, Nesreen K Ahmed, Eunyee Koh, and Sungchul Kim. Continuous-time dynamic network embeddings. In *Companion Proceedings of the Web Conference 2018*, pp. 969–976, 2018.

Marco Nurisso, Marta Morandini, Maxime Lucas, Francesco Vaccarino, Tommaso Gili, and Giovanni Petri. Higher-order laplacian renormalization. *Nature Physics*, pp. 1–8, 2025.

Howard T Odum. Trophic structure and productivity of silver springs, florida. *Ecological Monographs*, 27 (1):55–112, 1957.

Gergely Palla, Albert-László Barabási, and Tamás Vicsek. Quantifying social group evolution. *Nature*, 446 (7136):664–667, 2007.

Jose A Perea and John Harer. Sliding windows and persistence: An application of topological methods to signal analysis. *Foundations of Computational Mathematics*, 15(3):799–838, 2015.

Bryan Perozzi, Rami Al-Rfou, and Steven Skiena. Deepwalk: Online learning of social representations. In *Proceedings of the 20th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pp. 701–710, 2014.

Giovanni Petri, Martina Scolamiero, Irene Donato, and Francesco Vaccarino. Topological strata of weighted complex networks. *PloS One*, 8(6):e66506, 2013.

Stuart L Pimm. The complexity and stability of ecosystems. *Nature*, 307(5949):321–326, 1984.

Alex Pothen, Horst D Simon, and Kang-Pu Liou. Partitioning sparse matrices with eigenvectors of graphs. *SIAM Journal on Matrix Analysis and Applications*, 11(3):430–452, 1990.

Emanuele Rossi, Ben Chamberlain, Fabrizio Frasca, Davide Eynard, Federico Monti, and Michael Bronstein. Temporal graph networks for deep learning on dynamic graphs. *arXiv preprint arXiv:2006.10637*, 2020.

Jean-François Rual, Kavitha Venkatesan, Tong Hao, Tomoko Hirozane-Kishikawa, Amélie Dricot, Ning Li, Gabriel F Berriz, Francis D Gibbons, Matija Dreze, Nono Ayivi-Guedehoussou, et al. Towards a proteome-scale map of the human protein–protein interaction network. *Nature*, 437(7062):1173–1178, 2005.

Aliaksei Sandryhaila and Jose MF Moura. Discrete signal processing on graphs: Frequency analysis. *IEEE Transactions on Signal Processing*, 62(12):3042–3054, 2014.

Franco Scarselli, Marco Gori, Ah Chung Tsoi, Markus Hagenbuchner, and Gabriele Monfardini. The graph neural network model. *IEEE Transactions on Neural Networks*, 20(1):61–80, 2008.

Marten Scheffer, Stephen R Carpenter, Timothy M Lenton, Jordi Bascompte, William Brock, Vasilis Dakos, Johan Van de Koppel, Ingrid A Van de Leemput, Simon A Levin, Egbert H Van Nes, et al. Anticipating critical transitions. *Science*, 338(6105):344–348, 2012.

Youngjoo Seo, Michaël Defferrard, Pierre Vandergheynst, and Xavier Bresson. Structured sequence modeling with graph convolutional recurrent networks. In *International Conference on Neural Information Processing*, pp. 362–373. Springer, 2018.

Nino Shervashidze, Pascal Schweitzer, Erik Jan Van Leeuwen, Kurt Mehlhorn, and Karsten M Borgwardt. Weisfeiler-lehman graph kernels. *Journal of Machine Learning Research*, 12(9), 2011.

Jianbo Shi and Jitendra Malik. Normalized cuts and image segmentation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 22(8):888–905, 2000.

Yunsheng Shi, Zhengjie Huang, Shikun Feng, Hui Zhong, Wenjin Wang, and Yu Sun. Masked label prediction: Unified message passing model for semi-supervised classification. *arXiv preprint arXiv:2009.03509*, 2020.

David I Shuman, Sunil K Narang, Pascal Frossard, Antonio Ortega, and Pierre Vandergheynst. The emerging field of signal processing on graphs: Extending high-dimensional data analysis to networks and other irregular domains. *IEEE signal processing magazine*, 30(3):83–98, 2013.

ARE Sinclair, Simon AR Mduma, J Grant C Hopcraft, John M Fryxell, RAY Hilborn, and Simon Thirgood. Long-term ecosystem dynamics in the serengeti: lessons for conservation. *Conservation Biology*, 21(3): 580–590, 2007.

David I Spivak. *Category theory for the sciences*. MIT press, 2014.

Olaf Sporns, Giulio Tononi, and Rolf Kötter. The human connectome: a structural description of the human brain. *PLoS Computational Biology*, 1(4):e42, 2005.

Alexandra Swanson, Margaret Kosmala, Chris Lintott, Robert Simpson, Arfon Smith, and Craig Packer. Snapshot serengeti, high-frequency annotated camera trap images of 40 mammalian species in an african savanna. *Scientific Data*, 2(1):1–14, 2015.

Jun-ichi Takahashi. Category theory and organic electronics. *Physics Open*, 15:100148, 2023.

Chayant Tantipathananandh, Tanya Berger-Wolf, and David Kempe. A framework for community identification in dynamic social networks. In *Proceedings of the 13th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pp. 717–726, 2007.

David Tilman. *Resource competition and community structure*. Number 17. Princeton University Press, 1982.

Jason M Tylianakis, Raphael K Didham, Jordi Bascompte, and David A Wardle. Global change and species interactions in terrestrial ecosystems. *Ecology Letters*, 11(12):1351–1363, 2008.

David C Van Essen, Kamil Ugurbil, Edward Auerbach, Deanna Barch, Timothy EJ Behrens, Richard Bucholz, Acer Chang, Liyong Chen, Maurizio Corbetta, Sandra W Curtiss, et al. The human connectome project: a data acquisition perspective. *Neuroimage*, 62(4):2222–2231, 2012.

Petar Veličković, Guillem Cucurull, Arantxa Casanova, Adriana Romero, Pietro Lio, and Yoshua Bengio. Graph attention networks. *arXiv preprint arXiv:1710.10903*, 2017.

Ulrike Von Luxburg. A tutorial on spectral clustering. *Statistics and Computing*, 17(4):395–416, 2007.

Richard J Williams and Neo D Martinez. Simple rules yield complex food webs. *Nature*, 404(6774):180–183, 2000.

Zhitao Ying, Dylan Bourgeois, Jiaxuan You, Marinka Zitnik, and Jure Leskovec. Gnnexplainer: Generating explanations for graph neural networks. *Advances in Neural Information Processing Systems*, 32, 2019.

A Detailed Proofs of Results in §4

Here, we give the full mathematical proofs for the lemmas and theorems presented in our theoretical analysis in §4. We present the proofs in the same order as they appear in the main text, beginning with the foundational guarantee of functoriality and proceeding to the main scientific and interpretive results. The proofs for Lemma 1 and Lemma 2 are provided in §A.1 and §A.5, respectively. The proofs for our main theorems are provided in the subsequent sections: Theorem 1 (Trace Conservation) in §A.2, Theorem 2 (Spectral Sensitivity) in §A.3, Theorem 3 (Stability-Spectrum Equivalence) in §A.4, Theorem 4 (Rank-One Update) in §A.6, Theorem 5 (Structural Inertia) in §A.7, Theorem 6 (Node Removal) in §A.8, and Theorem 7 (Signal Transport) in §A.9.

A.1 Proof of Lemma 1: Functoriality of χ

Lemma 1 (Functoriality of χ). *The map $\chi : \texttt{Proc} \to \texttt{Spec}$ is a functor. It preserves identity morphisms and the composition of morphisms.*

Proof. To formally prove that χ is a functor, we must verify that it satisfies the two defining axioms of a functor: (1) it maps identity morphisms in the source category `Proc` to identity morphisms in the target category `Spec`, and (2) it preserves the structure of composition.

1. Preservation of Identity Morphisms. We must show that for any object $G \in \texttt{Proc}$, $\chi(\mathrm{id}_G) = \mathrm{id}_{\chi(G)}$.

Let $G = (V, E, W)$ be an object in `Proc`. The identity morphism on this object, $\mathrm{id}_G : G \to G$, is the process that results in no change to the network state. This means the resulting weight function W' is identical to the initial weight function W.

The functor χ maps an object G to the vector space $\chi(G)$ spanned by the orthonormal eigenvectors $\{\mathbf{v}_i\}_{i=1}^n$ of its symmetrised Laplacian, L_{sym}. The action of the functor on the morphism, $\chi(\mathrm{id}_G)$, is the linear transformation that maps the eigenbasis of the initial state to the eigenbasis of the final state.

Since $W' = W$, the Laplacians are identical ($L'_{sym} = L_{sym}$), and therefore their eigenspaces are identical. The transformation $\chi(\mathrm{id}_G)$ maps the vector space $\chi(G)$ to itself, and it maps the chosen orthonormal basis $\{\mathbf{v}_i\}$ to the identical basis $\{\mathbf{v'}_i = \mathbf{v}_i\}$. A linear transformation that maps every vector in a basis to itself is, by definition, the identity transformation on that vector space.

Therefore, $\chi(\mathrm{id}_G) = \mathrm{id}_{\chi(G)}$. This part of the proof is now self-contained.

2. Preservation of Composition. We must show that for any two composable morphisms $p_1 : G_1 \to G_2$ and $p_2 : G_2 \to G_3$, the following holds: $\chi(p_2 \circ p_1) = \chi(p_2) \circ \chi(p_1)$.

Let $B_1 = \{\mathbf{v}_i\}$, $B_2 = \{\mathbf{u}_i\}$, and $B_3 = \{\mathbf{w}_i\}$ be the orthonormal eigenbases for the Laplacians of G_1, G_2, and G_3, respectively.

The functor maps these processes to linear transformations in `Spec`, which are change of basis operations:

- $\chi(p_1) : \chi(G_1) \to \chi(G_2)$ maps basis B_1 to B_2.
- $\chi(p_2) : \chi(G_2) \to \chi(G_3)$ maps basis B_2 to B_3.

The composition of these linear transformations in `Spec`, $\chi(p_2) \circ \chi(p_1)$, is the map from $\chi(G_1) \to \chi(G_3)$ obtained by first applying $\chi(p_1)$ and then $\chi(p_2)$.

Now consider the composition in `Proc`. The process $p_2 \circ p_1$ is a single morphism that maps G_1 directly to G_3. The functor maps this composite process to the linear transformation $\chi(p_2 \circ p_1) : \chi(G_1) \to \chi(G_3)$, which is the single operation that maps basis B_1 to B_3.

In linear algebra, the transformation matrix for a change of basis from a basis B_1 to a basis B_3, via an intermediate basis B_2, is given by the product of the individual change of basis matrices. Let M_1 be the matrix for $\chi(p_1)$ and M_2 be the matrix for $\chi(p_2)$. The composition $\chi(p_2) \circ \chi(p_1)$ corresponds to the matrix product $M_2 M_1$. This resulting matrix is precisely the change of basis matrix from B_1 to B_3, which is the matrix representation of $\chi(p_2 \circ p_1)$.

Therefore, the linear transformations are the same: $\chi(p_2 \circ p_1) = \chi(p_2) \circ \chi(p_1)$.

Since χ preserves both identity morphisms and composition, it is a valid functor. □

A.2 Proof of Theorem 1: The Spectral Trace Conservation Law

Theorem 1 (The Spectral Trace Conservation Law). *Let $p : G \to G'$ be a conservative process, where the total resource is unchanged ($\mathcal{R}(G) = \mathcal{R}(G')$). The trace of the symmetrised Laplacian is conserved, i.e., $\mathrm{Tr}(L'_{sym}) = \mathrm{Tr}(L_{sym})$. Consequently, the sum of the Laplacian eigenvalues is an invariant of the process.*

Proof. The proof proceeds by first establishing a direct identity between the total resource of a network, $\mathcal{R}(G)$, and the trace of its symmetrised Laplacian, $\mathrm{Tr}(L_{sym})$.

Let $G = (V, E, W)$ be a network object in `Proc` with $|V| = n$. The symmetrised Laplacian is defined as $L_{sym} = D_{sym} - A_{sym}$.

1. **Trace of the Laplacian.** The trace of a matrix is the sum of its diagonal elements. For the Laplacian, since A_{sym} has zeros on its diagonal (assuming no self-loops, or they can be handled separately without loss of generality), the trace is the sum of the diagonal elements of the degree matrix D_{sym}:
$$\mathrm{Tr}(L_{sym}) = \sum_{i=1}^{n} (L_{sym})_{ii} = \sum_{i=1}^{n} (D_{sym})_{ii}$$

2. **Relating Trace to Adjacency Matrix.** The i-th diagonal entry of the degree matrix, $(D_{sym})_{ii}$, is defined as the sum of weights of all edges incident to node i in the symmetrised graph. This is the sum of the i-th row of the symmetrised adjacency matrix A_{sym}:
$$(D_{sym})_{ii} = \sum_{j=1}^{n} (A_{sym})_{ij}$$
Therefore, the trace of the Laplacian is the sum of all entries in the degree matrix, which is equivalent to the sum of all entries in the symmetrised adjacency matrix:
$$\mathrm{Tr}(L_{sym}) = \sum_{i=1}^{n} \sum_{j=1}^{n} (A_{sym})_{ij}$$

3. **Relating Adjacency Matrix to Total Resource.** We now substitute the definition of A_{sym}, where $(A_{sym})_{ij} = \frac{W(i,j)+W(j,i)}{2}$:
$$\begin{aligned} \mathrm{Tr}(L_{sym}) &= \sum_{i=1}^{n} \sum_{j=1}^{n} \frac{W(i,j) + W(j,i)}{2} \\ &= \frac{1}{2} \left(\sum_{i=1}^{n} \sum_{j=1}^{n} W(i,j) + \sum_{i=1}^{n} \sum_{j=1}^{n} W(j,i) \right) \end{aligned}$$
The term $\sum_{i,j} W(i,j)$ is the sum of all edge weights in the original directed graph, which is by definition the total resource $\mathcal{R}(G)$. The second term, $\sum_{i,j} W(j,i)$, is also the sum over all edge weights and is therefore also equal to $\mathcal{R}(G)$. This gives us the direct identity:
$$\mathrm{Tr}(L_{sym}) = \frac{1}{2}(\mathcal{R}(G) + \mathcal{R}(G)) = \mathcal{R}(G)$$

4. **Applying the Conservative Constraint.** The process $p : G \to G'$ is defined as conservative, meaning $\mathcal{R}(G) = \mathcal{R}(G')$. From the identity established in the previous step, we have $\mathrm{Tr}(L_{sym}) = \mathcal{R}(G)$ and $\mathrm{Tr}(L'_{sym}) = \mathcal{R}(G')$. The conservation of resources therefore directly implies the conservation of the trace:
$$\mathrm{Tr}(L_{sym}) = \mathrm{Tr}(L'_{sym})$$

5. **Relating Trace to Eigenvalues.** A fundamental theorem of linear algebra states that the trace of any matrix is equal to the sum of its eigenvalues. Let $\{\lambda_i\}_{i=1}^{n}$ be the eigenvalues of L_{sym} and $\{\lambda'_i\}_{i=1}^{n}$ be the eigenvalues of L'_{sym}. We have:
$$\sum_{i=1}^{n} \lambda_i = \mathrm{Tr}(L_{sym}) \quad \text{and} \quad \sum_{i=1}^{n} \lambda'_i = \mathrm{Tr}(L'_{sym})$$

Combining this with the result from the previous step, we conclude that the sum of the eigenvalues is conserved:

$$\sum_{i=1}^{n} \lambda_i = \sum_{i=1}^{n} \lambda'_i$$

This completes the proof. □

A.3 Proof of Theorem 2: The Spectral Sensitivity of Algebraic Connectivity

Theorem 2 (The Spectral Sensitivity of Algebraic Connectivity). *Let $p : G \to G'$ be a process that induces a sufficiently small change in the symmetrised Laplacian, $\Delta L_{sym} = L'_{sym} - L_{sym}$. If the process is structurally fragmenting, defined as satisfying the condition $\mathbf{v}_2^T(\Delta L_{sym})\mathbf{v}_2 < 0$, where $\mathbf{v}_2$ is the Fiedler eigenvector of the initial graph G, then the Fiedler value will decrease ($\lambda'_2 < \lambda_2$).*

Proof. The proof for this theorem relies on a standard result from matrix perturbation theory, which describes how the eigenvalues of a symmetric matrix change in response to a small perturbation.

Let L_{sym} be the symmetrised Laplacian of the initial graph G, with eigenvalues $\lambda_1 \leq \lambda_2 \leq \cdots \leq \lambda_n$ and a corresponding complete orthonormal set of eigenvectors $\{\mathbf{v}_i\}_{i=1}^n$. Let the process p induce a small perturbation, resulting in a new Laplacian $L'_{sym} = L_{sym} + \Delta L_{sym}$ with new eigenvalues $\{\lambda'_i\}_{i=1}^n$. We assume the perturbation ΔL_{sym} is small enough such that first-order effects dominate.

1. **First-Order Eigenvalue Perturbation.** A fundamental result from matrix analysis (see, e.g., Bhatia (1992)) states that for a small symmetric perturbation ΔL_{sym}, the first-order change in a simple eigenvalue λ_k is given by the Rayleigh quotient of the perturbation matrix with respect to the corresponding eigenvector $\mathbf{v}_k$:

   $$\lambda'_k = \lambda_k + \mathbf{v}_k^T(\Delta L_{sym})\mathbf{v}_k + O(\|\Delta L_{sym}\|^2)$$

   For a sufficiently small perturbation, we can analyse the first-order term to determine the direction of the change. The change in the eigenvalue is thus approximated by:

   $$\Delta\lambda_k = \lambda'_k - \lambda_k \approx \mathbf{v}_k^T(\Delta L_{sym})\mathbf{v}_k$$

2. **Applying to the Fiedler Value.** The Fiedler value is the second smallest eigenvalue, λ_2. For this theorem, we assume λ_2 is a simple (non-repeated) eigenvalue, which is the generic case for connected graphs. Applying the perturbation formula to λ_2, we get:

   $$\lambda'_2 - \lambda_2 \approx \mathbf{v}_2^T(\Delta L_{sym})\mathbf{v}_2$$

   where $\mathbf{v}_2$ is the Fiedler eigenvector of the original Laplacian L_{sym}.

3. **Applying the "Structurally Fragmenting" Condition.** The theorem's premise defines a process as "structurally fragmenting" if it satisfies the following condition:

   $$\mathbf{v}_2^T(\Delta L_{sym})\mathbf{v}_2 < 0$$

   This condition gives a precise mathematical meaning to the idea that the process is "aligned" with the network's primary structural vulnerability, as identified by the Fiedler eigenvector.

4. **Conclusion.** By substituting the condition from Step 3 into the first-order approximation from Step 2, we directly obtain the result:

   $$\lambda'_2 - \lambda_2 < 0$$

   which implies:

   $$\lambda'_2 < \lambda_2$$

Therefore, for any sufficiently small process that is structurally fragmenting, the Fiedler value of the network is guaranteed to decrease. This completes the proof. □

A.4 Proof of Theorem 3: The Stability-Spectrum Equivalence

Theorem 3 (The Stability-Spectrum Equivalence). *A dynamic network sequence $(G_t)_{t=1}^{\infty}$ governed by dissipative processes converges to a stable state G_∞ if and only if its corresponding sequence of spectral data (eigenvalues and eigenvectors of $L_{sym,t}$) converges to a stable limit.*

Proof. This is an equivalence proof, which requires proving two implications. The core of the proof rests on the continuity of the maps between the spaces of weight matrices, Laplacians, and their spectral decompositions.

Part 1: Stability $\implies$ Spectral Convergence. In this direction, we prove that if the physical state of the network converges, its spectral representation must also converge.

1. **Assumption of Stability.** We assume that the system converges to a stable state G_∞. By definition, this means the sequence of weight functions converges to a limit function, $W_t \to W_\infty$ as $t \to \infty$. This convergence is typically defined in terms of a matrix norm, e.g., $\|W_t - W_\infty\| \to 0$.

2. **Continuity of the Laplacian Map.** The mapping from a weight matrix W to its corresponding symmetrised Laplacian L_{sym} is a continuous function. The entries of L_{sym} are simple linear combinations of the entries of W. Specifically, $(L_{sym})_{ij}$ is a function of $W(i,j), W(j,i)$, and sums of weights connected to nodes i and j. As a finite sum of continuous functions, this mapping is continuous. Therefore, the convergence of the weight matrices implies the convergence of the Laplacian matrices:
$$W_t \to W_\infty \implies L_{sym,t} \to L_{sym,\infty}$$

3. **Continuity of Spectral Decomposition.** The eigenvalues and eigenvectors of a real symmetric matrix are continuous functions of its entries (see, e.g., Bhatia (1992)). This means that for a converging sequence of matrices, their spectra also converge.

4. **Conclusion of Part 1.** Since the sequence of Laplacians $L_{sym,t}$ converges to a limit $L_{sym,\infty}$, it follows from the continuity of the spectral map that the corresponding sequences of their eigenvalues and eigenvectors must also converge to a stable limit. This proves the first implication.

Part 2: Spectral Convergence $\implies$ Stability. In this direction, we prove that if the spectral representation of the network converges, the physical state must also have converged.

1. **Assumption of Spectral Convergence.** We assume that the full set of spectral data converges. This means the sequence of eigenvalues $\{\lambda_{i,t}\}$ converges to a limit spectrum $\{\lambda_{i,\infty}\}$, and the sequence of eigenvector matrices V_t (whose columns are the eigenvectors) converges to a limit matrix V_∞.

2. **Convergence of the Laplacian.** A real symmetric matrix is uniquely determined by its spectral decomposition via the formula $L_{sym} = V\Lambda V^T$, where Λ is the diagonal matrix of eigenvalues. Since matrix multiplication and transposition are continuous operations, the convergence of both V_t and Λ_t implies the convergence of the sequence of symmetrised Laplacian matrices:
$$(V_t \to V_\infty \text{ and } \Lambda_t \to \Lambda_\infty) \implies L_{sym,t} \to L_{sym,\infty}$$

3. **From Laplacian Convergence to System Stability.** The convergence of $L_{sym,t}$ means that for any small $\epsilon > 0$, there exists a time T such that for all $t > T$, $\|L_{sym,t+1} - L_{sym,t}\| < \epsilon$. This implies that the change in the trace, $|\mathrm{Tr}(L_{sym,t+1}) - \mathrm{Tr}(L_{sym,t})|$, must also approach zero.

   From Theorem 1, we know that $\mathrm{Tr}(L_{sym}) = \mathcal{R}(G)$. Therefore, the total resource change per step, $|\mathcal{R}(G_{t+1}) - \mathcal{R}(G_t)|$, must also approach zero.

   The processes in our `Proc` category are fundamentally dissipative ($\mathcal{R}(G_{t+1}) \leq \mathcal{R}(G_t)$). A dissipative system can only stop dissipating resources when it has reached a stable fixed point or equilibrium

state. Since the change in total resource is approaching zero, the system must be approaching a state where the processes acting upon it are no longer dissipative but have become conservative. In a system without external energy inputs, the only state where this can happen is a stable equilibrium. Any further change would either require dissipation (which would change the trace, contradicting spectral convergence) or an external input (which is outside our current model).

Therefore, the convergence of the spectral data implies the convergence of the underlying physical state to a stable equilibrium, $G_t \to G_\infty$.

Since both implications hold, the equivalence is established. This completes the proof. □

A.5 Proof of Lemma 2: The Change of Basis Formula

Lemma 2 (The Change of Basis Formula). *Let $p : G \to G'$ be a process, with $\{\mathbf{v}_i\}$ and $\{\mathbf{v'}_j\}$ being the orthonormal eigenbases of the initial and final Laplacians, respectively. The entry (i, j) of the matrix representation of the linear transformation $\chi(p)$ is given by the inner product of the respective basis vectors: $(\chi(p))_{ij} = \langle \mathbf{v'}_i, \mathbf{v}_j \rangle$.*

Proof. Let the source vector space be $U = \chi(G)$ and the target vector space be $U' = \chi(G')$. Let $B = \{\mathbf{v}_1, \ldots, \mathbf{v}_n\}$ be the orthonormal eigenbasis for U, and let $B' = \{\mathbf{v'}_1, \ldots, \mathbf{v'}_n\}$ be the orthonormal eigenbasis for U'.

The linear transformation $\chi(p) : U \to U'$ is the map that governs the change of basis. By definition, the matrix representation of a linear transformation is constructed column by column. The j-th column of the matrix for $\chi(p)$ is the coordinate vector of the transformed basis vector, $\chi(p)(\mathbf{v}_j)$, expressed in the target basis B'.

In our framework, the functor χ maps the abstract process p to the specific linear transformation that describes the change in the geometric frame of the network. This means the transformation maps the old basis vectors directly onto themselves, but now they exist within the new space. Formally, we can consider the action of the transformation on an old basis vector $\mathbf{v}_j$ to be the vector $\mathbf{v}_j$ itself, which we now must represent in the new basis B'.

Let M be the matrix representation of $\chi(p)$ with respect to the bases B and B'. The j-th column of M is the vector $[\mathbf{v}_j]_{B'}$, the coordinate representation of $\mathbf{v}_j$ in the basis B'.

To find the i-th component of this coordinate vector, we need to find the scalar coefficient c_i in the linear combination:

$$\mathbf{v}_j = \sum_{k=1}^{n} c_k \mathbf{v'}_k$$

Since the basis B' is orthonormal, we can find the coefficient c_i by taking the inner product (dot product) of both sides with the basis vector $\mathbf{v'}_i$:

$$\langle \mathbf{v'}_i, \mathbf{v}_j \rangle = \left\langle \mathbf{v'}_i, \sum_{k=1}^{n} c_k \mathbf{v'}_k \right\rangle$$

By the linearity of the inner product, we can move the summation and scalar coefficients out:

$$\langle \mathbf{v'}_i, \mathbf{v}_j \rangle = \sum_{k=1}^{n} c_k \langle \mathbf{v'}_i, \mathbf{v'}_k \rangle$$

Because B' is an orthonormal basis, the inner product $\langle \mathbf{v'}_i, \mathbf{v'}_k \rangle$ is equal to the Kronecker delta, δ_{ik}, which is 1 if $i = k$ and 0 otherwise. The summation therefore collapses, leaving only the term where $k = i$:

$$\langle \mathbf{v'}_i, \mathbf{v}_j \rangle = c_i \cdot 1 = c_i$$

This shows that the i-th coordinate of the vector $\mathbf{v}_j$ in the basis B' is precisely the inner product $\langle \mathbf{v'}_i, \mathbf{v}_j \rangle$.

Since the entry (i, j) of the transformation matrix M is the i-th component of the j-th column vector, we have:

$$M_{ij} = (\chi(p))_{ij} = c_i = \langle \mathbf{v'}_i, \mathbf{v}_j \rangle$$

This completes the proof. □

A.6 Proof of Theorem 4: The Rank-One Update Signature

Theorem 4 (The Rank-One Update Signature). *Let p be a simple process that only perturbs the weight of a single edge between nodes a and b. The resulting change in the symmetrised Laplacian, ΔL_{sym}, is a rank-one matrix. Consequently, the transformation matrix $\chi(p)$ is a low-rank perturbation of the identity matrix.*

Proof. The proof consists of two parts. First, we show that a single edge perturbation results in a rank-one update to the symmetrised Laplacian. Second, we state how this low-rank update affects the resulting transformation matrix $\chi(p)$.

1. The Rank of the Laplacian Perturbation. Let the process p change the weight of the directed edge (a, b) by a value δ_1 and the weight of the edge (b, a) by a value δ_2. In the simplest case, one of these is zero. The change in the weight function is non-zero only for these two edges.

The change in the symmetrised adjacency matrix, ΔA_{sym}, is given by:

$$(\Delta A_{sym})_{ij} = \frac{\Delta W(i, j) + \Delta W(j, i)}{2}$$

This results in a matrix that is zero everywhere except at entries (a, b) and (b, a), where the value is $(\delta_1+\delta_2)/2$. Let $\delta = (\delta_1 + \delta_2)/2$.

The change in the symmetrised degree matrix, ΔD_{sym}, is a diagonal matrix where $(\Delta D_{sym})_{ii} = \sum_j (\Delta A_{sym})_{ij}$. The only non-zero entries will be on the diagonal at positions (a, a) and (b, b):

- $(\Delta D_{sym})_{aa} = (\Delta A_{sym})_{ab} = \delta$
- $(\Delta D_{sym})_{bb} = (\Delta A_{sym})_{ba} = \delta$

The total change in the symmetrised Laplacian is $\Delta L_{sym} = \Delta D_{sym} - \Delta A_{sym}$. This matrix has only four non-zero entries:

- $(\Delta L_{sym})_{aa} = \delta$
- $(\Delta L_{sym})_{bb} = \delta$
- $(\Delta L_{sym})_{ab} = -\delta$
- $(\Delta L_{sym})_{ba} = -\delta$

Let $\mathbf{e}_k$ be the standard basis vector with a 1 in the k-th position. The vector $\mathbf{u} = \mathbf{e}_a - \mathbf{e}_b$ is a vector with 1 at position a, -1 at position b, and 0 elsewhere. The outer product of this vector with itself is a matrix $(\mathbf{u}\mathbf{u}^T)$. Its entries match the structure of our perturbation matrix. We can therefore write our perturbation matrix as a scalar multiple of this outer product:

$$\Delta L_{sym} = \delta(\mathbf{e}_a - \mathbf{e}_b)(\mathbf{e}_a - \mathbf{e}_b)^T$$

A matrix that can be expressed as the outer product of a single non-zero column vector and a single non-zero row vector is, by definition, a rank-one matrix.

2. Consequence for the Transformation Matrix. The new Laplacian is $L'_{sym} = L_{sym} + \Delta L_{sym}$. We have shown that the perturbation ΔL_{sym} is a simple, rank-one matrix.

According to matrix perturbation theory (specifically, results related to low-rank updates), the eigenvectors of a matrix that has been perturbed by a low-rank matrix are themselves a low-rank perturbation of the original eigenvectors. This means the new eigenbasis $\{\mathbf{v'}_i\}$ is "close" to the original eigenbasis $\{\mathbf{v}_i\}$.

The transformation matrix $\chi(p)$ has entries $(\chi(p))_{ij} = \langle \mathbf{v'}_i, \mathbf{v}_j \rangle$. Since the new basis is a small perturbation of the old orthonormal basis, the new basis vectors are nearly aligned with the old ones. This means:

- Diagonal entries $(\chi(p))_{ii} = \langle \mathbf{v'}_i, \mathbf{v}_i \rangle$ will be close to 1.
- Off-diagonal entries $(\chi(p))_{ij} = \langle \mathbf{v'}_i, \mathbf{v}_j \rangle$ for $i \neq j$ will be close to 0.

Therefore, the transformation matrix $\chi(p)$ will be a low-rank perturbation of the identity matrix, I. The structure of this perturbation is not arbitrary but is determined by the components of the original eigenvectors at the perturbed nodes, a and b.

This completes the proof. □

A.7 Proof of Theorem 5: The Structural Inertia Theorem

Theorem 5 (The Structural Inertia Theorem). *Let p be a process that induces a small perturbation ΔL_{sym}. The resulting transformation matrix $\chi(p)$ is diagonally dominant. The magnitude of its off-diagonal entries is bounded by the norm of the perturbation and the spectral gaps of the original graph.*

Proof. The proof relies on the Davis-Kahan theorem, a cornerstone of matrix perturbation theory, which provides a bound on the rotation of eigenspaces under a symmetric perturbation.

Let L_{sym} and $L'_{sym} = L_{sym} + \Delta L_{sym}$ be the initial and final symmetrised Laplacians, respectively. Let their eigenvalues be $\{\lambda_i\}$ and $\{\lambda'_i\}$, and their corresponding orthonormal eigenvector matrices be $V = [\mathbf{v}_1 | \dots | \mathbf{v}_n]$ and $V' = [\mathbf{v'}_1 | \dots | \mathbf{v'}_n]$. We assume the perturbation, as measured by its spectral norm $\|\Delta L_{sym}\|_2$, is small.

1. **Matrix Entries as Inner Products.** From Lemma 2, the entries of the transformation matrix $\chi(p)$ are given by the inner products of the old and new eigenvectors:

$$(\chi(p))_{ij} = \langle \mathbf{v'}_i, \mathbf{v}_j \rangle$$

2. **Diagonal Entries.** The diagonal entries are $(\chi(p))_{ii} = \langle \mathbf{v'}_i, \mathbf{v}_i \rangle = \cos(\theta_i)$, where θ_i is the angle between the new and old i-th eigenvectors. For a small perturbation, the eigenvectors do not change much, so θ_i is small and $\cos(\theta_i)$ is close to 1.

3. **Off-Diagonal Entries and the Davis-Kahan Theorem.** The off-diagonal entries, $(\chi(p))_{ij}$ for $i \neq j$, represent the projection of an old eigenvector $\mathbf{v}_j$ onto a new eigenvector $\mathbf{v'}_i$. The Davis-Kahan Sine Theta Theorem provides a bound on the angle between the old and new eigenspaces. A direct consequence of the theorem gives a bound on the magnitude of these individual inner products.

   Let's assume the eigenvalues of L_{sym} are simple. The magnitude of the inner product between a new eigenvector $\mathbf{v'}_i$ and an old eigenvector $\mathbf{v}_j$ for $i \neq j$ is bounded by:

$$|\langle \mathbf{v'}_i, \mathbf{v}_j \rangle| \leq \frac{\|\Delta L_{sym}\|_2}{|\lambda_i - \lambda_j|}$$

   The term $|\lambda_i - \lambda_j|$ is the spectral gap between the i-th and j-th eigenvalues.

4. **Conclusion: Diagonal Dominance.** The result from Step 3 shows that the magnitude of the off-diagonal entries, $|(\chi(p))_{ij}|$, is small, provided the perturbation $\|\Delta L_{sym}\|_2$ is small and the spectral

gap $|\lambda_i - \lambda_j|$ is not pathologically close to zero. The diagonal entries, as shown in Step 2, are close to 1.

A matrix is diagonally dominant if, for every row, the magnitude of the diagonal entry is greater than the sum of the magnitudes of all other (off-diagonal) entries in that row. For a sufficiently small perturbation, the diagonal entries will be approximately 1, while the off-diagonal entries will be close to 0. Thus, the matrix $\chi(p)$ is guaranteed to be diagonally dominant.

This proves that a small physical perturbation cannot cause a large, arbitrary re-shuffling of the fundamental structural modes of the network. This completes the proof. □

A.8 Proof of Theorem 6: The Node Removal Signature

Theorem 6 (The Node Removal Signature). *Let p be a process that removes a node k from a network G with n nodes. The resulting transformation $\chi(p)$ maps the original n-dimensional eigenspace to the new $(n-1)$-dimensional eigenspace and is a projection operator.*

Proof. The proof involves analysing the structural change in the Laplacian matrix that results from deleting a vertex and then characterising the nature of the map between the corresponding eigenspaces.

Let the original graph be G with n vertices, and let its symmetrised Laplacian be L_{sym}, an $n \times n$ matrix. Let the resulting graph after removing node k be G', with $n-1$ vertices. The new Laplacian, L'_{sym}, is the $(n-1) \times (n-1)$ principal submatrix of L_{sym} obtained by deleting the k-th row and k-th column.

1. **Change in Dimension.** The original eigenspace is $\chi(G) \cong \mathbb{R}^n$, spanned by the eigenvectors $\{\mathbf{v}_i\}_{i=1}^n$ of L_{sym}. The final eigenspace is $\chi(G') \cong \mathbb{R}^{n-1}$, spanned by the eigenvectors $\{\mathbf{u}_j\}_{j=1}^{n-1}$ of L'_{sym}. The transformation $\chi(p)$ is a map from an n-dimensional space to an $(n-1)$-dimensional space. Such a map cannot be an isomorphism (like a simple change of basis) but must involve a reduction in dimension.

2. **Constructing the Projection Operator.** We can model the overall transformation in two conceptual steps. First, we define a projection operator $P_k : \mathbb{R}^n \to \mathbb{R}^{n-1}$ that removes the k-th component from any vector in the original space. This operator effectively projects the original n-dimensional space onto the subspace that corresponds to the remaining nodes.

   The eigenvectors of the new Laplacian, $\{\mathbf{u}_j\}$, form a basis for this target $\mathbb{R}^{n-1}$ space. The transformation $\chi(p)$ can be understood as the composition of this projection with a subsequent change of basis within the lower-dimensional space. However, the dominant characteristic of the map from $\mathbb{R}^n$ to $\mathbb{R}^{n-1}$ is the projection itself.

3. **The Kernel of the Transformation.** A projection from a higher-dimensional space to a lower-dimensional space has a non-trivial kernel (or null space)—the set of vectors that are mapped to the zero vector. In this case, the kernel of the projection P_k is the one-dimensional subspace spanned by the standard basis vector $\mathbf{e}_k$ (the vector with a 1 in the k-th position and zeros elsewhere). Any structural mode of the original network that was entirely localised to node k (i.e., any vector proportional to $\mathbf{e}_k$) is annihilated by the transformation. This directly corresponds to the physical removal of the node.

4. **Relationship Between Spectra (Cauchy Interlacing Theorem).** The relationship between the spectra of L_{sym} and its principal submatrix L'_{sym} is not arbitrary but is tightly constrained by the Cauchy Interlacing Theorem. This theorem states that the eigenvalues of the new matrix interlace the eigenvalues of the original matrix:

   $$\lambda_1 \leq \lambda'_1 \leq \lambda_2 \leq \lambda'_2 \leq \cdots \leq \lambda'_{n-1} \leq \lambda_n$$

   This ensures that the spectral properties of the subgraph G' are a predictable and well-behaved consequence of the properties of the original graph G.

In summary, the removal of a node induces a transformation from an n-dimensional space to an $(n-1)$-dimensional space, which is fundamentally a projection. The kernel of this projection corresponds directly to the removed node, providing a unique and identifiable signature for this type of major topological change. This completes the proof. □

A.9 Proof of Theorem 7: The Signal Transport Theorem

Theorem 7 (The Signal Transport Theorem). *Let* $\mathbf{f}$ *be a vector representing a "signal" on the nodes of a network* G*. After a process* p*, the signal* $\mathbf{f}'$ *on the new network* G' *that maintains the same coordinates with respect to the new eigenbasis is given by the transformation* $\mathbf{f}' = T_{transport}\mathbf{f}$*, where the transport matrix is* $T_{transport} = V'V^T$*, with* V *and* V' *being the matrices of eigenvectors for* G *and* G'*, respectively.*

Proof. The proof consists of deriving the explicit matrix form for the transport operator by defining the signal in the spectral domain and then mapping it back to the node domain.

1. **Signal Representation in the Spectral Domain.** Let G be the initial network with its corresponding symmetrised Laplacian L_{sym}. Let V be the $n \times n$ matrix whose columns are the complete orthonormal set of eigenvectors of L_{sym}, $B = \{\mathbf{v}_1, \ldots, \mathbf{v}_n\}$.

   A signal $\mathbf{f} \in \mathbb{R}^n$ on the nodes of the graph can be expressed as a linear combination of these basis vectors. The vector of coefficients, $\mathbf{a} \in \mathbb{R}^n$, which we call the spectral coordinates of the signal, is given by:

   $$\mathbf{f} = V\mathbf{a}$$

   Since V is an orthonormal matrix, its inverse is its transpose ($V^{-1} = V^T$). We can therefore find the spectral coordinates from the signal via:

   $$\mathbf{a} = V^{-1}\mathbf{f} = V^T\mathbf{f}$$

2. **Defining the Transported Signal.** Let the process p transform the network from G to G'. The new network G' has a new symmetrised Laplacian L'_{sym} and a new matrix of orthonormal eigenvectors, V'.

   The core idea of signal transport is to define a new signal, $\mathbf{f}'$, on the nodes of G' that has the *exact same spectral coordinates* $\mathbf{a}$ as the original signal, but expressed in the *new basis* V'. The new signal is therefore synthesised from the original spectral coordinates and the new basis:

   $$\mathbf{f}' = V'\mathbf{a}$$

3. **Deriving the Transport Operator.** Our goal is to find the matrix $T_{transport}$ that maps the original signal vector $\mathbf{f}$ directly to the new signal vector $\mathbf{f}'$, such that $\mathbf{f}' = T_{transport}\mathbf{f}$.

   We can derive this by substituting the expression for the spectral coordinates $\mathbf{a}$ from Step 1 into the synthesis equation from Step 2:

   $$\mathbf{f}' = V'(V^T\mathbf{f})$$

   By the associativity of matrix multiplication, this can be written as:

   $$\mathbf{f}' = (V'V^T)\mathbf{f}$$

   By comparing this result with the desired form $\mathbf{f}' = T_{transport}\mathbf{f}$, we can directly identify the transport operator matrix:

   $$T_{transport} = V'V^T$$

This provides the explicit formula for the transport matrix, which depends only on the eigenvectors of the initial and final network states. This completes the proof. □