Parameter-Based Eigenvalue Elasticity Analysis for Stock-and-Flow Models in a Categorical Framework

Nastaran Jamali, Dr. Nathaniel Osgood

January 2026

Abstract

We report here the extension of our categorical computing framework to perform parameter-based eigenvalue elasticity analysis for nonlinear stock-and-flow models. The work is explicitly grounded in the eigenvalue sensitivity theory formulated in the MSc thesis of Qian Zhang, which treats eigenvalues as implicit functions of model parameters and emphasizes the retention of complex eigenmodes and near-zero regimes.

At each reporting time, the nonlinear system is locally linearized to obtain the Jacobian matrix, and the full eigen-spectrum is computed. For every eigenmode and parameter, exact eigenvalue sensitivities and elasticities are evaluated. Rather than restricting attention to a single dominant eigenvalue, parameter influence is aggregated across all eigenmodes using maximum-magnitude criteria. When elasticities become ill-conditioned due to vanishing eigenvalues, dominance is assessed using raw eigenvalue sensitivities, as prescribed in Zhang's framework.

The methodology is implemented for categorically encoded stock-and-flow models and validated on a classical SIR epidemiological system. Results demonstrate that parameter dominance varies significantly across dynamical regimes, and that the most informative sensitivity signals arise precisely in oscillatory or near-zero regimes where dominant- eigenvalue-based analysis fails. The categorically grounded implementation provides a faithful computational realization of thesis-level eigenvalue sensitivity theory and a robust reporting strategy suitable for nonlinear, time-varying systems.

Executive Summary

This project implements a parameter-based eigenvalue sensitivity and elasticity analysis for nonlinear stock-and-flow models, following the theoretical framework developed in Qian Zhang's MSc thesis. The objective is to produce a computational pipeline that identifies, at each point in time, which model parameters exert the strongest structural influence on local system behavior.

The implementation operates directly on categorically specified models. Each model is simulated forward in time, locally linearized at selected time points, and analyzed using a full eigen-decomposition of the Jacobian matrix. Eigenvalue sensitivities and elasticities with respect to all parameters are computed for every eigenmode, including complex- conjugate pairs.

To avoid the limitations of dominant-eigenvalue-only analysis, parameter influence is aggregated across eigenmodes using maximum-magnitude criteria. When eigenvalues approach zero and elasticities become numerically unstable, the analysis automatically transitions to sensitivity-based dominance measures, ensuring continuity of interpretation.

The pipeline produces time-resolved tables and plots showing:

- the full eigenvalue spectrum,
- parameter elasticities and sensitivities across all modes,
- dominant parameters by elasticity and by sensitivity,
- diagnostic indicators for oscillatory and near-zero regimes.

The methodology is demonstrated in an epidemiological model of SIRS. The results show clear shifts in parameter dominance across epidemic phases and confirm that the strongest sensitivity signals arise near regime transitions. All outputs are generated by a reproducible Julia notebook accompanying this report.

1 Background

Local behavior of nonlinear dynamical systems is governed by the eigenvalues of the Jacobian matrix. Eigenvalues determine whether perturbations grow, decay, or generate oscillatory responses. While eigenvalue-based reasoning has long been used in System Dynamics, traditional applications often focus on structural elements such as feedback loops or individual Jacobian entries [5].

In applied modeling contexts, however, the primary objects of interest are frequently model parameters. Parameters represent intervention levers, policy controls, or external assumptions, and analysts seek principled ways to assess how changes in these quantities affect system behavior over time. Classical local sensitivity analysis addresses this question only at the level of state trajectories and often fails to capture changes in qualitative behavior.

Qian Zhang's thesis [2] provided a rigorous extension of eigenvalue sensitivity theory to parameters. By treating eigenvalues as implicit functions of model parameters through the Jacobian, Zhang derived exact expressions for eigenvalue sensitivities and elasticities that remain valid for nonlinear, time-varying systems. A central insight of this framework is that sensitivity information should not be restricted to a single dominant eigenvalue, particularly near regime transitions, oscillatory behavior, or zero-crossings of growth rates.

This project builds directly on Zhang's formulation [2]", saying "This project builds directly on Zhang's formulation [2], and a subsequent paper on the same work[3] and translates it into a practical, computational pipeline for stock-and-flow models whose diagrammatic structure is explicitly encoded using mechanisms from category theory. The goal is not to reinterpret feedback structure, but to report parameter dominance in a way that remains valid precisely when dominant-eigenvalue assumptions break down. By retaining all eigenmodes and explicitly handling complex and near-zero regimes, the analysis targets the regions where structural sensitivity is most informative.

2 Theoretical Foundations: Parameter-Based Eigenvalue Sensitivity

Consider a nonlinear dynamical system expressed in state–space form

$$\dot{x} = f(x, p), \qquad x \in \mathbb{R}^n,\ p \in \mathbb{R}^m,$$

where x denotes the vector of state variables (stocks) and p denotes a vector of model parameters. At a given operating point (x^*, p), the system is locally linearized as

$$\dot{\xi} = J(x^*, p)\, \xi,$$

where $J = \partial f / \partial x$ is the Jacobian matrix evaluated at (x^*, p).

Local system behavior around this point in state space x^* is governed by the eigenvalues of J. Positive real parts indicate local instability or growth, negative real parts indicate decay, and nonzero imaginary parts correspond to oscillatory dynamics. In time-varying nonlinear systems, both the Jacobian and its eigenvalues evolve along the trajectory, making eigenvalue analysis inherently local and time-dependent.

2.1 Eigenvalues as Implicit Functions of Parameters

Eigenvalue-based analysis has been previously explored in the system dynamics literature (e.g., [7, 6]). In this work, eigenvalues are treated as implicit functions of model parameters through the Jacobian:

$$\lambda = \lambda\big(J(p)\big).$$

Let λ_i be a (possibly complex) eigenvalue of J with corresponding right and left eigenvectors r_i and l_i, normalized such that

$$l_i^\top r_i = 1.$$

Following standard results in eigenvalue sensitivity analysis (see, e.g., [2]), the sensitivity of λ_i with respect to a parameter p_j is given by

$$\frac{\partial \lambda_i}{\partial p_j} = l_i^\top \left(\frac{\partial J}{\partial p_j} \right) r_i.$$

Here and throughout, the index i indexes eigenvalues of J, while j indexes model parameters. This expression follows directly from classical eigenvalue perturbation theory but is here interpreted explicitly as a *parameter sensitivity*, not a structural one. Importantly, $\partial J/\partial p_j$ is computed analytically or via automatic differentiation, preserving exact dependence on the parameter.

2.2 Parameter-Based Eigenvalue Elasticity

To obtain a dimensionless measure comparable across parameters, Zhang defines the eigenvalue elasticity with respect to parameter p_j as

$$E_{i,j} = \frac{\partial \lambda_i}{\partial p_j} \frac{p_j}{\lambda_i}$$

whenever $\lambda_i \neq 0$. This elasticity, which can be notionally written as $\frac{\frac{\partial \lambda_i}{\lambda_i}}{\frac{\partial p_j}{p_j}}$, quantifies the proportional change in eigenvalue λ_i induced by a proportional change in parameter p_j.

Unlike loop-based elasticities, parameter-based elasticities do not require identifying feedback loops, and remain well-defined for a broad range of model structures. Because many interventions are commonly characterized explored through changes to parameter values directly support intervention-oriented interpretation.

2.3 Treatment of Complex Eigenvalues

A central contribution of Zhang's thesis is the explicit treatment of complex eigenvalues. For a complex-conjugate pair

$$\lambda_i = a \pm ib,$$

the sensitivity $\partial\lambda_i/\partial p_j$ and elasticity $E_{i,j}$ are themselves complex-valued. Rather than discarding these modes or projecting them prematurely onto the real axis, Zhang emphasizes the value of retaining the full complex quantity.

Accordingly, the realization of this framework reports the following:

- the magnitude $|E_{i,j}|$,
- the real part $\mathrm{Re}(E_{i,j})$,
- the imaginary part $\mathrm{Im}(E_{i,j})$,

for each eigenmode and each parameter. This separation allows growth/decay effects and oscillatory effects to be interpreted independently.

2.4 Beyond Dominant Eigenvalue Assumptions

Classical analyses often focus exclusively on the dominant eigenvalue, defined as the eigenvalue with the largest real part.[5] Zhang explicitly cautions against this restriction, noting that:

- near-zero eigenvalues indicate regime transitions where sensitivity is highest,
- oscillatory modes may dominate behavior even when their real parts are small,
- parameter influence may be distributed across multiple eigenmodes.

Therefore, this project computes eigenvalue sensitivities and elasticities for *all* eigenmodes at each time step. Parameter dominance is then assessed using aggregated criteria across modes, rather than relying on a single dominant eigenvalue.

2.5 Normalization of Left and Right Eigenvectors

For each eigenvalue λ_i of the Jacobian matrix, corresponding right and left eigenvectors r_i and l_i are computed numerically from the full eigen-decomposition. The normalization condition adopted throughout this work is

$$l_i^\top r_i = 1,$$

consistent with the formulation presented in Zhang's Sections 3.2.2.1–3.2.2.3.

This normalization is applied uniformly to both real and complex eigenmodes. Eigenvalue sensitivities are computed using the bilinear form

$$\frac{\partial \lambda_i}{\partial p_j} = l_i^\top \left(\frac{\partial J}{\partial p_j} \right) r_i,$$

without introducing complex conjugation.

This choice reflects the computational realization implemented in the accompanying notebook and preserves consistency with Zhang's original sensitivity expression, which is formulated in terms of a transpose rather than a Hermitian inner product. All reported elasticities and dominance measures are therefore interpreted with respect to this normalization convention.

2.6 Aggregation Across Eigenmodes

To summarize parameter influence without discarding information, we adopt the aggregation strategy implied by Zhang's analysis. For each parameter p_j at time t_k, we compute:

$$\max_i |E_{i,j}(t_k)|, \qquad \max_i |\mathrm{Re}(E_{i,j}(t_k))|, \qquad \max_i |\mathrm{Im}(E_{i,j}(t_k))|.$$

As above, index i ranges over the eigenvalues. These quantities identify which parameter exerts the strongest influence on any local mode of the system at that time.

When all $|\lambda_i|$ fall below a specified tolerance, elasticities become ill-conditioned. In these cases, Zhang recommends reverting to raw sensitivities $\partial\lambda_i/\partial p_j$, which remain finite. This fallback criterion is implemented explicitly in the present work.

2.7 Interpretive Principle

The guiding interpretive principle of this project is therefore:

> *If a system enters an oscillatory regime or approaches a zero-crossing of eigenvalues, parameter-based sensitivity analysis becomes more informative, not less.*

This principle directly motivates retaining complex eigenmodes, aggregating across modes, and avoiding dominant-eigenvalue-only simplifications.

3 Categorical Stock-and-Flow Representation

All models in this project are encoded using the `StockAndFlowF` schema, an attributed C-set (ACSet) representation of stock-and-flow systems. This categorical formulation enables a principled separation between the *syntactic structure* of a model and its *semantic interpretation* as a system of differential equations, following recent developments in compositional system dynamics [1] and the broader framework of applied category theory [4].

At the syntactic level, a `StockAndFlowF` instance consists of several typed object classes and morphisms:

- **Stocks**, representing state variables that accumulate flows,
- **Flows**, representing rate processes that transfer quantity between stocks,
- **Auxiliary variables**, which encode intermediate algebraic expressions, including for flows
- **Parameters**, which remain constant during simulation but influence flow expressions.

These objects are related by a fixed ACSet schema that specifies incidence relations (e.g., which flows affect which stocks) and functional dependencies (e.g., which auxiliaries and parameters feed into which flows).

Crucially, this categorical representation is *purely structural*: it describes how components are connected – characterizing,, as it were, the grammar of the system that defines the syntax – but not yet how the system evolves in time. It bears emphasis that this representation *categorically characterizes the mathematics of the stock and flow diagrams themselves*, without committing to a specific semantics. Semantics are introduced via a functorial interpretation that maps a `StockAndFlowF` instance to semantic domains. While the authors have explored elsewhere alternative semantic domains[10]—such as those involving discrete transition semantics—we consider here the classic semantic domain involving ordinary differential equation (ODE) system of the form

$$\dot{x} = f(x, p)$$

where x is the vector of stock levels and p is the vector of parameters. The vector field f is generated compositionally from the flow definitions and their associated auxiliary expressions.

Because the entire construction is compositional, the same categorical structure also induces higher-order semantic objects in a systematic way. In particular:

- The Jacobian matrix $J(x, p) = \partial f/\partial x$ is obtained by differentiating the induced vector field with respect to the stocks.
- Parameter derivatives $\partial J/\partial p_j$ are obtained by differentiating the Jacobian with respect to individual parameters.

In the implementation, these derivatives are computed using automatic differentiation, but their existence and structure are guaranteed by the categorical organization of the model.

This syntax–semantics separation is essential for the eigenvalue elasticity analysis performed in this project. The categorical layer ensures that the analysis pipeline is entirely generic: no part of the eigenvalue, sensitivity, or elasticity computation depends on model-specific assumptions beyond those mathematically encoded in the `StockAndFlowF` schema. Any stock-and-flow model expressed as an ACSet (that is, as an instance of that schema) can therefore be passed through the same sequence of steps:

$$\text{ACSet} \longrightarrow f(x,p) \longrightarrow J(x,p) \longrightarrow \{\lambda_k, r_k, l_k\}.$$

From a categorical perspective, this can be viewed as a chain of structure-preserving mappings: the combinatorial graph of stocks, flows, and auxiliaries determines a differentiable dynamical system, whose linearization determines a spectrum of eigenmodes. Eigenvalue sensitivities and elasticities then provide a quantitative link back from system behavior to underlying structural and parametric features.

An important consequence of this formulation is that the analysis naturally extends beyond reduced stock-flow only graphs. Auxiliary variables and parameters appear explicitly in the ACSet and therefore contribute to the Jacobian and its parameter derivatives. This is particularly important for parameter-based eigenvalue elasticity analysis, where sensitivities are computed with respect to parameters rather than structural links.

Categorical description Formally, a `StockAndFlowF` model is an attributed C-set whose underlying category includes multiple objects (each representing a *type* of quantity in that can be present in a stock and flow diagram – stocks, flows, auxiliary variables, and parameters) together with (typed) morphisms encoding their relationships. Stocks represent state variables, flows represent rate processes, auxiliary variables mediate functional dependencies, and parameters act as exogenous constants.

The incidence structure is captured by morphisms mapping flows to their source and target stocks, as well as morphisms mapping auxiliary variables to the stocks and parameters on which they depend. Attribute maps allow for assigning algebraic expressions and operators to flows and auxiliaries. A given stock and flow diagram is then an instance of this schema – a structure preserving mapping (functor) that maps the schema category into the category of FinSet, in which objects are sets of finite size, and morphisms are functions between those sets. For a given such instance (associated with a particular stock and flow diagram), such a functor (call it G) , preserves structure, mapping objects to objects and morphisms to morphisms, mapping composites $f \mathbin{⨟} g$ of two morphisms f and g in the schema category to the composites of their images $G(f) \mathbin{⨟} G(g)$, and mapping the obligatory identity morphisms id_c on each object c in the schema to the identity morphism on its image $id_{G(c)}$. Thus, for example, the object S in the schema representing the type stocks is mapped to the a particular set $G(S)$, each of whose elements consists of a particular stock. Similarly, the object I representing the type inflow would be mapped to the set $G(I)$, each of whose elements would consist of an inflows, and a morphism in the schema $downstream : I \to S$ going from the type I to type S would be mapped by the functor for this particular instance (representing a particular stock and flow model) to a function mapping each inflow within the set $G(I)$ of inflows into the particular stock in the set $G(S)$ into which it flows.

The operational semantics of the model are obtained by interpreting such a categorically encode instance in the category of smooth vector fields, yielding an ordinary differential equation $\dot{x} = f(x, p)$.

Concrete morphisms in `StockAndFlowF`. In the concrete implementation used in this project, the categorical structure is realized through explicit incidence maps, such as the *downstream* morphism mentioned above. For other types of variables, for example, the morphism `:lvv` encodes dependencies of auxiliary variables on other auxiliary variables, while `:lpvv` captures dependencies of auxiliary variables on parameters. Similarly, morphisms such as `:lvs` encode links from auxiliary variables to stocks.

These morphisms—and the functions to which they are mapped by the instance functor—

determine the dependency graph through which perturbations in parameters propagate to flows and ultimately to stock dynamics. As a result, the structure of the Jacobian matrix $J(x, p)$ and the parameter derivatives $\partial J/\partial p_j$ are not specified manually, but are induced by the categorical incidence relations encoded in the ACSet.

3.1 Relation between ACSet and Semantics in Catlab

In terms of implementation, the `StockAndFlowF` representation used in this work follows the ACSet (Attributed C-set) formalism as implemented in the StockFlow.jl package contributed to by Xiaoyan Li and the authors, as well as the larger Catlab ecosystem. In this setting, a model is defined purely in terms of objects, morphisms, and attributes, independent of any particular semantic interpretation. Semantics are then assigned by functorial mappings into target categories such as smooth dynamical systems.

This separation allows the same categorical structure to support multiple analytical interpretations. In the present work, the primary semantic interpretation maps a `StockAndFlowF` instance to an ODE system and its associated Jacobian. However, the same structure could in principle support alternative semantics, such as discrete-time dynamics or stochastic interpretations.

Crucially, this ACSet-based formulation enables eigenvalue elasticity analysis to be implemented generically: the analysis operates on semantic objects (Jacobians and eigenvalues) that are systematically derived from the categorical structure, rather than on hand-crafted equations. This ensures that the parameter-based elasticity pipeline applies uniformly across models, regardless of their size or specific functional forms.

Overall, the categorical `StockAndFlowF` representation provides a principled foundation for implementing eigenvalue sensitivity and elasticity analysis in a way that is general, modular, extensible, and mathematically transparent. It ensures that the numerical results reported later in this report are not *ad hoc* computations, but consequences of a well-defined structural semantics.

From categorical structure to eigenvalue-based analysis. The categorical stock-and-flow representation introduced in this section provides more than a convenient modeling syntax: it establishes a precise structural foundation from which local dynamic behavior can be derived in a systematic and model-agnostic manner. In particular, once the vector field $f(x, p)$ is functorially induced from the `StockAndFlowF` ACSet, the Jacobian matrix $J(x, p)$ and its parameter derivatives $\partial J / \partial p_j$ become canonical objects associated with the model at any operating point.

This structural grounding enables the transition to eigenvalue-based sensitivity analysis without introducing *ad hoc* assumptions or model-specific algebra. Eigenvalue Elasticity Analysis (EEA) is therefore not applied to a hard-coded or an externally specified system of equations, but rather emerges naturally as a semantic layer built on top of the categorical representation. The following section leverages this structure to quantify how parameters influence local system behavior through their effects on the eigenvalues of the Jacobian.

In particular, this formulation allows sensitivity analysis to remain meaningful even in regimes where the dominant eigenvalue is zero or complex, a point that becomes central in the parameter-based elasticity analysis developed in the following section.

4 Computational Implementation of Parameter-Based Eigenvalue Elasticity Analysis

This section describes the computational realization of the parameter-based eigenvalue sensitivity and elasticity framework introduced in the preceding sections. The emphasis is on implementation choices that are required to faithfully operationalize Zhang's theoretical formulation for nonlinear, time-varying systems, rather than on reproducing analytic derivations.

All results reported in this study are generated by an accompanying dedicated Julia notebook (`ParameterBasedE.ipynb`), which serves as an executable specification of the analysis pipeline.

4.1 Software Stack

The implementation is written entirely in Julia and relies on the following components:

- `OrdinaryDiffEq.jl` for numerical integration of the nonlinear ODE system,
- `ForwardDiff.jl` for automatic differentiation of the vector field and its Jacobian,
- `LinearAlgebra` for eigen-decomposition of dense Jacobian matrices,
- `Catlab.jl` and `StockAndFlowF` for categorical model representation,
- `DataFrames.jl` and `Plots.jl` for structured reporting and visualization.

The choice of automatic differentiation is essential: all Jacobians and parameter derivatives are computed directly from the induced vector field, without manual linearization or model-specific algebra.

4.2 Categorical Data Structures

Models are specified using the `StockAndFlowF` ACSet schema, which encodes stocks, flows, auxiliary variables, and parameters as typed objects with explicit incidence relations. This representation characterizes the dependency structure through which parameter perturbations propagate to the system dynamics.

From a single categorical specification, the pipeline functorially derives:

- the nonlinear vector field $f(x, p)$,
- the Jacobian matrix $J(x, p) = \partial f / \partial x$,
- parameter derivatives $\partial J / \partial p_j$ for all parameters.

Because the ACSet structure characterizes auxiliary variables and parameter dependencies explicitly, parameter-based sensitivity analysis is performed on the full semantic model, not on a reduced stock-only representation.

4.3 Jacobian and Parameter-Derivative Strategy

At each saved simulation time t_k, the system state $x(t_k)$ is obtained by numerical integration. The Jacobian matrix $J(x(t_k), p)$ is computed using forward-mode automatic differentiation applied to the induced vector field.

For each parameter p_j, the partial derivative $\partial J/\partial p_j$ is computed numerically by differentiating the Jacobian with respect to p_j. This use of automatic differentiation serves as a computational substitute for the symbolic differentiation assumed in Zhang's theoretical development and does not alter the underlying sensitivity formulation. This step is central to Zhang's formulation: eigenvalue sensitivities are defined directly with respect to parameters, not inferred indirectly from structural proxies. No sparsity assumptions or dominance-based truncations are applied at this stage; the entire Jacobian and all parameter derivatives are retained.

4.4 Eigen-Decomposition, Normalization, and Complex Handling

At each reporting time, a full eigen-decomposition of the Jacobian matrix is performed. All eigenvalues are retained, including complex-conjugate pairs and modes with small real parts.

For each eigenvalue λ_k, corresponding right and left eigenvectors r_k and l_k are computed and normalized to satisfy

$$l_k^\top r_k = 1,$$

as required by Zhang's sensitivity formula. This normalization is enforced explicitly in the implementation to ensure numerical consistency of sensitivity and elasticity values.

Complex eigenvalues are handled without projection or filtering. Sensitivities and elasticities are computed as complex quantities, and their real parts, imaginary parts, and magnitudes are stored separately. This design choice preserves oscillatory sensitivity information and avoids misclassification of oscillatory regimes as numerically unstable.

A tolerance λ_{tol} is introduced to identify near-zero eigenvalues. When $|\lambda_k| < \lambda_{\text{tol}}$, elasticities involving λ_k are marked as undefined, but sensitivities remain well-defined and are retained for dominance analysis.

4.5 Numerical Diagnostics and Conditioning Checks

At each reporting time, numerical diagnostics are computed to assess the reliability of the eigen-decomposition and the resulting sensitivity and elasticity values. In particular, the condition number of the eigenvector matrix R is evaluated. Large condition numbers indicate near-defective Jacobians, in which case eigenvalue sensitivities and elasticities may be numerically unreliable.

In addition, pairwise distances between eigenvalues are monitored. When two or more eigenvalues satisfy

$$|\lambda_i - \lambda_k| < \delta,$$

the system is flagged as operating in a near-multiplicity regime. Such regimes are associated with heightened structural sensitivity and reduced numerical stability of individual eigenmodes.

These diagnostics do not alter the computed sensitivities or elasticities, but are reported alongside the results as interpretive indicators. They provide explicit signals identifying regimes where caution is required in interpreting modal dominance, fully consistent with Zhang's discussion of ill-conditioned and transitional regimes.

4.6 Aggregation and Reporting Outputs

Rather than selecting a single dominant eigenvalue, parameter influence is aggregated across eigenmodes. For each parameter p_j at time t_k, the notebook computes:

$$\max_k |E_{k,j}|, \qquad \max_k |\mathrm{Re}(E_{k,j})|, \qquad \max_k |\mathrm{Im}(E_{k,j})|.$$

When elasticities are undefined for all eigenmodes at a given time, dominance is assessed using the corresponding maximum absolute sensitivities:

$$\max_k |S_{k,j}|.$$

The reporting layer produces structured tables and plots that include:

- the full eigenvalue spectrum over time,
- aggregated parameter elasticities and sensitivities,
- dominant parameters identified by elasticity and by sensitivity,
- diagnostic notes indicating oscillatory or near-zero regimes.

This reporting strategy ensures that parameter dominance is identified consistently across all dynamical regimes, including those where dominant-eigenvalue-based analysis would otherwise fail.

5 Algorithmic Pipeline

At each saved time t_i:

1. Solve the ODE to obtain $x(t_i)$.
2. Compute $J(x(t_i), p)$ via automatic differentiation.

3. Compute all eigenpairs (λ_k, r_k, l_k).

4. Compute $\partial J/\partial p_j$ for each parameter.

5. Evaluate sensitivities $\partial \lambda_k/\partial p_j$.

6. Compute elasticities $E_{k,j}$.

7. Aggregate via $\max |E_{k,j}|$ and $\max |S_{k,j}|$.

8. Identify dominant parameters by elasticity and sensitivity.

Complex eigenvalues are retained throughout.

6 Aggregation, Dominance, and Reporting Strategy

To translate multi-modal sensitivity information into interpretable results, the notebook implements a structured aggregation and reporting strategy, consistent with Qian Zhang's thesis formulation.

For each parameter p_j at time t_k, the maximum absolute elasticity across eigenmodes is computed as

$$\max_k |E_{k,j}|,$$

with analogous quantities computed for the real and imaginary parts separately. This aggregation ensures that oscillatory modes and secondary modes contribute to dominance assessment when relevant.

In time steps where elasticities are undefined for all eigenmodes (typically due to $|\lambda_k|$ falling below tolerance), the analysis falls back to parameter dominance based on sensitivities:

$$\max_k |S_{k,j}|.$$

The notebook therefore reports, at each time step:

- maximum absolute elasticity per parameter across all modes,
- maximum absolute sensitivity per parameter across all modes,
- dominant parameter by elasticity,
- dominant parameter by sensitivity,
- and a diagnostic note describing the validity regime.

This dual reporting strategy avoids suppressing information near critical transitions and makes explicit when elasticity-based dominance is well-defined versus when sensitivity-based dominance is more appropriate.

All plots and tables in the case study section are generated directly from these aggregated outputs, ensuring that reported dominance reflects the full spectral structure of the model rather than a single selected eigenvalue.

6.1 Formal Definition of Parameter Dominance

Parameter dominance is defined operationally, following the aggregation logic implied by Zhang's thesis.

For each parameter p_j at time t_k, eigenvalue elasticities $E_{i,j}(t_k)$ are computed for all eigenmodes i satisfying $|\lambda_i(t_k)| > \lambda_{\mathrm{tol}}$. If at least one such eigenmode exists, elasticity-based dominance is defined as

$$p_j^*(t_k) = \arg\max_{p_j} \left(\max_i |E_{i,j}(t_k)| \right).$$

If no eigenmode satisfies the elasticity validity condition (i.e., $|\lambda_i(t_k)| \leq \lambda_{\mathrm{tol}}$ for all i), elasticities are considered ill-conditioned. In this case, dominance is determined using raw eigenvalue sensitivities:

$$p_j^*(t_k) = \arg\max_{p_j} \left(\max_i \left| \frac{\partial \lambda_i}{\partial p_j}(t_k) \right| \right).$$

This two-tier dominance rule ensures that parameter influence is always reported, even in near-zero or transitional regimes. Importantly, such regimes are not treated as failures of the method, but as structurally sensitive phases where parameter leverage is maximal, in direct accordance with Zhang's analysis.

6.2 Tolerance Selection and Interpretation of Missing Elasticities

A numerical tolerance λ_{tol} is introduced to identify eigenvalues whose magnitudes are sufficiently close to zero that elasticity computation becomes ill-conditioned. When $|\lambda_i| \ll 1$, the normalization factor $1/\lambda_i$ amplifies numerical noise and can lead to misleading elasticity values.

In such cases, elasticities are intentionally marked as undefined. This behavior is not a numerical failure, but an explicit diagnostic signal indicating that the system is operating near a critical transition. Zhang emphasizes that these regimes correspond to maximal structural sensitivity, and that reliance on elasticity alone is inappropriate.

Accordingly, missing elasticities in the reported tables are interpreted as *regime markers*, not errors, and arise directly from the explicit tolerance-based logic implemented in the computational pipeline. Dominance is still reported using raw sensitivities, ensuring continuity of analysis across all dynamical phases.

This design choice makes explicit the distinction between: (i) regimes where proportional influence is meaningful, and (ii) regimes where absolute sensitivity provides a more favourable and sensible interpretive lens.

7 Case Study: Parameter-Based Eigenvalue Elasticity Analysis of the SIRS Model

This section presents a case study applying the parameter-based eigenvalue elasticity framework to a SIRS epidemiological model with waning immunity. The purpose of this case study is not merely to report numerical outputs, but to examine how the full-spectrum, time-resolved analysis advocated in Qian Zhang's thesis behaves in a setting with reinfection feedback and persistent modal interaction.

As in the rest of this paper, all eigenmodes of the Jacobian are retained at every reporting time. No reduction to a single dominant eigenvalue is imposed, and complex eigenvalues are explicitly preserved rather than discarded. Parameter influence is evaluated across the entire spectrum, with elasticity- and sensitivity-based dominance reported separately when informative.

7.1 Model Specification and Experimental Setup

The SIRS model consists of three stocks: Susceptible (S), Infected (I), and Recovered (R). In contrast to the classical SIR formulation, the SIRS structure includes a feedback flow from R back to S, representing waning of immunity. The model is specified categorically using the `StockAndFlowF` schema, from which the system of ordinary differential equations, the Jacobian matrix, and all parameter derivatives are generated automatically.

The model equations are:

$$\dot{S} = \text{Const1} - \beta S I + \omega R,$$

$$\dot{I} = \beta S I - \gamma I,$$

$$\dot{R} = \gamma I - \omega R,$$

where β is the infection rate (implicitly incorporates effects of contact rate, probability of transmission per discordant contact, and population size), γ is the recovery rate, ω is the waning-immunity rate, and Const1 is a constant inflow term (representing, for example, the arrival of immigrants or births).

The parameter set therefore includes:

- β: infection rate,
- γ: recovery rate,
- ω: waning-immunity rate,
- Const1: constant inflow term.

The system is simulated over a fixed time horizon from the initial condition

$$S(0) = 990, \qquad I(0) = 5, \qquad R(0) = 0,$$

with parameter values

$$\beta = 0.0004, \qquad \gamma = 0.05, \qquad \omega = 0.01, \qquad \text{Const1} = 1.0.$$

Spectral analysis is performed at selected reporting times corresponding to the initial invasion phase, the transient regime, and later-time approach toward a stabilized reinfection structure. At each reporting time t_k, the Jacobian $J(x(t_k), p)$ is computed and fully diagonalized.

No assumptions are made regarding dominant eigenvalues, real-valued spectra, or modal separation.

7.2 Evolution of the Eigenvalue Spectrum

The local dynamic behavior of the SIRS model is governed by the eigenvalues of the Jacobian matrix. Figure 1 shows the complete eigenvalue spectrum in the complex plane across all reporting times.

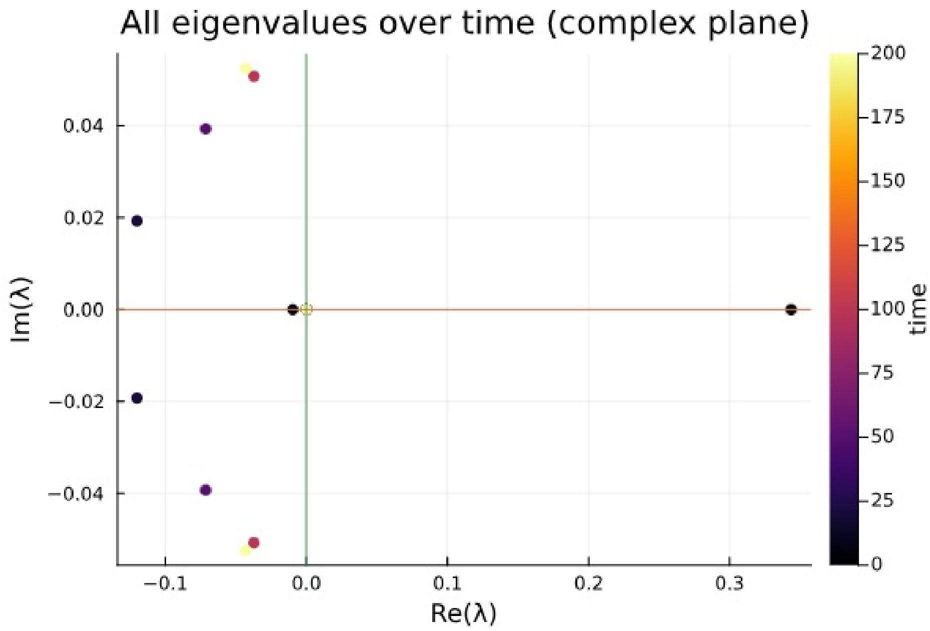

Figure 1: All eigenvalues of the SIRS Jacobian over time in the complex plane. Each point corresponds to one eigenvalue at a reporting time.

Several important features are visible. At the initial reporting time, one eigenvalue has a positive real part, indicating local growth of the infectious subsystem. As the trajectory evolves, this unstable growth regime disappears and the nonzero dynamic modes move into the left half-plane. From the transient phase onward, the spectrum is characterized by a complex-conjugate pair with negative real part, while one mode remains at or very near the origin.

This behavior differs qualitatively from the simpler SIR case. In the SIRS system, the feedback induced by waning immunity introduces a persistent reinfection loop, and this loop is reflected in the emergence and continuation of oscillatory modal structure. Rather than a purely monotonic decay regime, the model exhibits damped oscillatory local dynamics over a substantial part of the simulation horizon.

To clarify the temporal behavior of these modes, Figures 2 and 3 plot the real and imaginary components of all eigenvalues as functions of time.

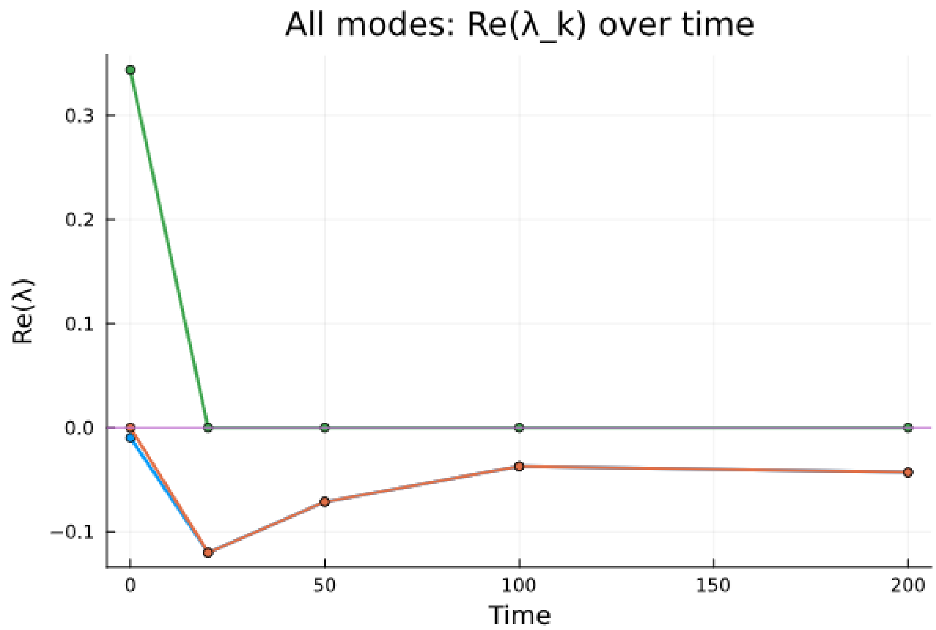

Figure 2: Real parts of all eigenvalues, $\mathrm{Re}(\lambda_k)$, over time for the SIRS model.

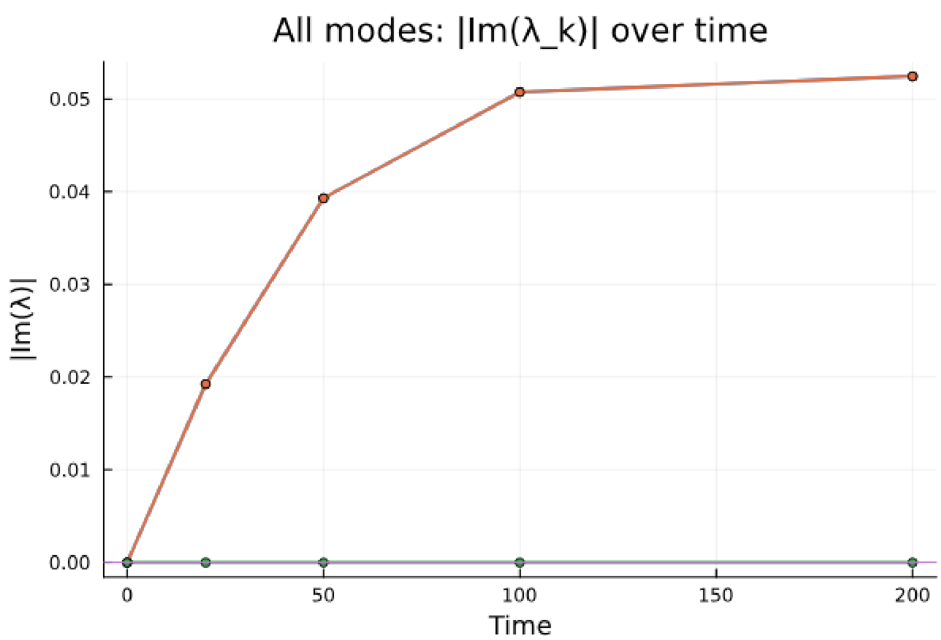

Figure 3: Absolute imaginary parts, $|\mathrm{Im}(\lambda_k)|$, of all eigenvalues over time for the SIRS model.

These plots show that the initially positive real eigenvalue is short-lived, after which the nonzero modes form a stable complex pair with negative real part. At the same time, the imaginary component grows away from zero and remains non-negligible at later times.

This indicates that oscillatory tendencies are not confined to a narrow transition point, but are instead a persistent feature of the local linearized dynamics once reinfection feedback becomes active.

7.3 Aggregated Parameter Elasticities Across All Eigenmodes

For each parameter and each reporting time, eigenvalue elasticities are computed for *every* eigenmode. To summarize parameter influence without discarding modal information, elasticities for that time are aggregated by taking the maximum absolute value across all modes:

$$\max_k |E_{k,j}|.$$

Figure 4 reports, over time, the aggregated elasticities across all parameters.

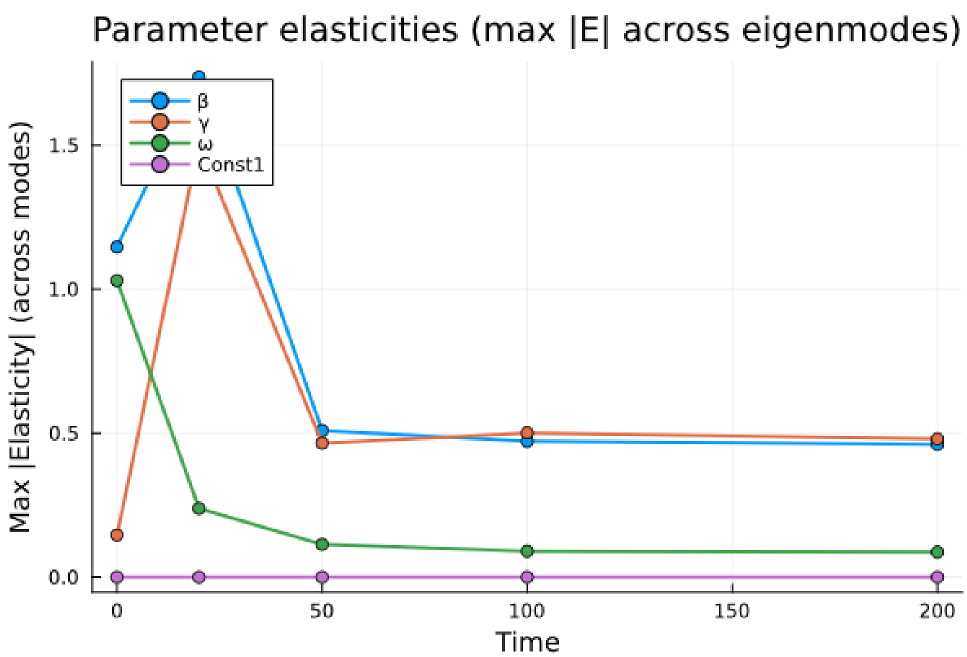

Figure 4: Maximum absolute per-parameter eigenvalue elasticity over time, aggregated across all eigenmodes for the SIRS model.

The results show a clear change in the structure of parameter influence relative to the SIR case. The infection rate β is the dominant elasticity driver during the initial and early transient phases. The recovery rate γ rises sharply after the initial time and becomes comparable to β, eventually overtaking it slightly at later reporting times. The waning-immunity parameter ω is non-negligible, especially at early times, but its influence remains smaller than that of β and γ after the initial transient. The constant term Const1 remains negligible throughout.

This is a substantively important result. In the SIR model, β and γ appeared nearly symmetric in the aggregated elasticity summaries. In the SIRS model, that symmetry is broken by the return flow from R to S. The addition of waning immunity alters the Jacobian structure sufficiently to separate the influence profiles of infection and recovery, while also introducing a third epidemiologically meaningful control parameter.

7.4 Sensitivity Dominance and Comparison with Elasticity Dominance

To complement the proportional interpretation provided by elasticities, we also compute the raw eigenvalue sensitivities

$$S_{k,j} = \frac{\partial \lambda_k}{\partial p_j}.$$

As with elasticities, sensitivities are aggregated across all eigenmodes using

$$\max_k |S_{k,j}|.$$

Figure 5 shows the aggregated sensitivity magnitudes.

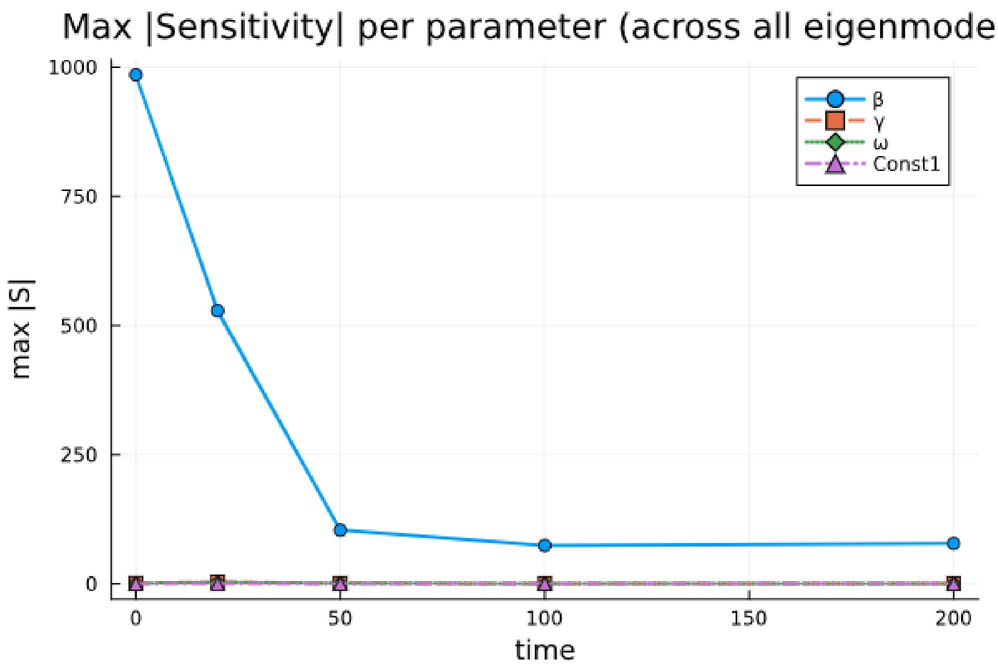

Figure 5: Maximum absolute per-parameter eigenvalue sensitivity over time, aggregated across all eigenmodes for the SIRS model.

In contrast to the elasticity results, sensitivity dominance remains with β throughout all reporting times. The aggregated sensitivity associated with β is substantially larger than that of γ or ω, especially during the initial regime. This indicates that, in absolute terms, the local spectrum is most responsive to perturbations in the infection rate even when proportional influence, as measured by elasticity, later shifts toward γ.

This difference between elasticity-based dominance and sensitivity-based dominance is analytically meaningful rather than contradictory. Elasticity dominance captures proportional leverage relative to the current parameter value and eigenvalue scale, whereas sensitivity dominance captures absolute spectral responsiveness. The SIRS case shows that these two notions of dominance need not coincide, particularly in models with persistent feedback structure and mixed real-complex modal behavior.

Figure 6 summarizes the dominant parameter over time under both criteria.

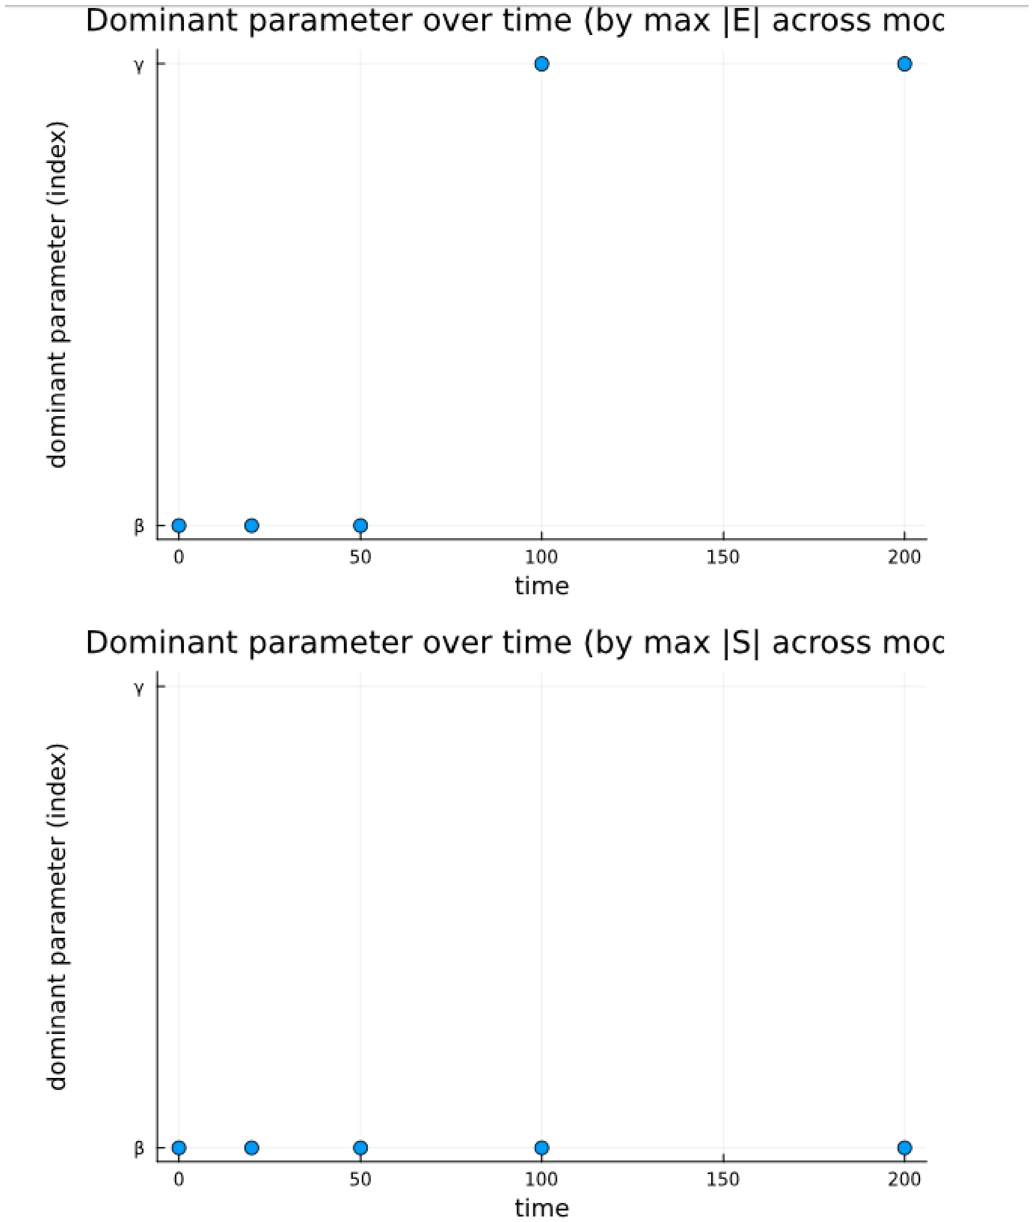

Figure 6: Dominant parameter over time in the SIRS model, identified using maximum absolute elasticity (top) and maximum absolute sensitivity (bottom) across all eigenmodes.

The figure confirms that elasticity-based dominance switches from β to γ at later times, while sensitivity-based dominance remains consistently with β. This provides a concrete example of the distinction emphasized in Zhang's framework: Different aggregation and normalization choices expose different but complementary aspects of structural leverage.

7.5 Summary Table and Interpretation

Table 1 summarizes the aggregated elasticity and sensitivity results at the selected reporting times.

Table 1: Aggregated parameter dominance measures for the SIRS model. All eigenmodes were retained, including complex modes.

Time	$\max\|E_\beta\|$	$\max\|E_\gamma\|$	$\max\|E_\omega\|$	$\max\|E_{\mathrm{Const1}}\|$	$\max\|S_\beta\|$	$\max\|S_\gamma\|$	$\max\|S_\omega\|$	Dom. by E	Dom. by S
0	1.14638	0.14641	1.02992	0.00000	985.081	1.00646	1.00080	β	β
20	1.73672	1.50786	0.23869	0.00000	528.402	3.67017	2.90485	β	β
50	0.50915	0.46561	0.11381	0.00000	103.717	0.75879	0.92737	β	β
100	0.47163	0.50125	0.09017	0.00000	74.2349	0.63117	0.56770	γ	β
200	0.46094	0.48034	0.08710	0.00000	78.1286	0.65134	0.59055	γ	β

Several patterns are evident. First, β dominates both elasticity and sensitivity at the initial time, reflecting its central role in determining early epidemic growth. Second, γ becomes nearly as influential as β during the transient phase and slightly exceeds it in the late-time elasticity summaries. Third, ω has a meaningful but secondary role: it is large enough to affect the spectrum, especially early on, yet does not become the leading parameter under either criterion. Finally, Const1 remains effectively inactive in the spectral summaries, indicating that the constant inflow term does not substantially shape the local modal structure of this particular configuration in the vicinity of the value of the parameter examined.

7.6 Interpretation and Theoretical Implications

The SIRS case study illustrates several core principles of parameter-based eigenvalue elasticity analysis.

- Parameter dominance is time-dependent and changes as the system moves from initial outbreak to reinfection-regulated dynamics.

- Retaining the full spectrum is essential, because the informative modal structure includes both a near-zero mode and a stable complex-conjugate pair.
- Elasticity-based dominance and sensitivity-based dominance need not agree; each captures a different notion of leverage.
- Adding waning immunity breaks the near-symmetry between infection and recovery observed in the simpler SIR case and introduces a distinct, interpretable role for ω.

Most importantly, the SIRS results show that Zhang's framework remains informative in a model whose local dynamics are not well described by a single real dominant mode. The presence of a persistent complex pair means that proportional parameter influence is distributed across oscillatory as well as decaying tendencies. Aggregation across all eigenmodes therefore provides a more faithful characterization of structural influence than any dominant-eigenvalue reduction would.

Overall, this case study supports Zhang's central claim: parameter-based eigenvalue sensitivity analysis is especially valuable in nonlinear, time-varying regimes where mode interaction, oscillatory structure, and changing dominance all coexist.

8 Conclusion

This project developed and implemented a rigorous, parameter-based eigenvalue elasticity analysis for nonlinear stock-and-flow models within a categorical modeling framework. The work is explicitly grounded in the PhD thesis of Qian Zhang and represents a faithful computational realization of her theoretical framework for eigenvalue sensitivity and elasticity with respect to parameters.

Unlike classical approaches that rely on a single dominant eigenvalue or on structural proxies such as feedback loops, the present framework retains the full eigen-spectrum at every point in time. Complex eigenvalues, near-zero modes, and secondary eigenmodes are treated as first-class analytical objects rather than numerical inconveniences. Parameter elasticities are computed for all eigenmodes and aggregated in a way that preserves structural sensitivity information even in oscillatory or transitional regimes.

Applied to the SIRS epidemiological model, the analysis reveals time-varying patterns of parameter dominance that cannot be recovered from dominant-eigenvalue analysis alone. Early growth, peak dynamics, and late decline exhibit qualitatively distinct sensitivity structures, with near-zero eigenvalue regimes showing the strongest sensitivity signals. These findings validate Zhang's claim that structural sensitivity is highest precisely where growth rates vanish or oscillatory behavior emerges.

The categorical `StockAndFlowF` representation plays a critical role in enabling this analysis. By separating structural specification from semantic interpretation, it allows eigenvalue sensitivities and elasticities to be derived systematically from model structure without *ad hoc* assumptions or manual differentiation. As a result, the framework is generic, extensible, and applicable to a wide class of nonlinear dynamical systems.

In summary, this work demonstrates that parameter-based eigenvalue elasticity analysis provides a robust and interpretable sensitivity framework precisely in regimes where traditional methods fail. By unifying Zhang's theoretical insights with categorical modeling and modern numerical tools, the project delivers a practical methodology for structural sensitivity analysis in complex dynamic systems.

Future work may extend this framework to higher-dimensional models, stochastic or hybrid dynamics, and alternative semantic interpretations, but the core conclusion remains: retaining all eigenmodes and respecting near-zero and oscillatory regimes is essential for meaningful sensitivity analysis in non-linear models.

References

[1] Li, X., Patterson, E., Mabry, P. L., & Osgood, N. D. (2025). Compositional System Dynamics: The Higher Mathematics Underlying System Dynamics Diagrams & Practice. *arXiv preprint arXiv:2509.18475*.

[2] Zhang, Q. (2013). *Eigenvalue Sensitivity Analysis and Its Application to Dynamic Systems*. MSc Thesis. Department of Computer Science, University of Saskatchewan

[3] Zhang, Q., & Osgood, N. (2010). Summary function elasticity analysis for an individual-based system dynamics model. *Proceedings of the 2010 Winter Simulation Conference*, 405–416.

[4] Fong, B., & Spivak, D. I. (2019). *Seven Sketches in Compositionality: An Invitation to Applied Category Theory*. Cambridge University Press.

[5] Kampmann, C. E. (2012). Feedback loop gains and system behavior. *System Dynamics Review*.

[6] Saleh, M., Oliva, R., Kampmann, C., & Davidsen, P. (2010). A comprehensive analytical approach for policy analysis of system dynamics models. *System Dynamics Review*.

[7] Guneralp, B. (2006). Towards coherent loop dominance analysis of system dynamics models. *System Dynamics Review*.

[8] Guneralp, B. (2006). *Exploring Structure-Behavior Relations in Nonlinear Dynamic Feedback Models*. PhD thesis, University of Illinois at Urbana-Champaign.

[9] Saleh, M. (2002). *The Characterization of Model Behavior and its Causal Foundation*. PhD thesis, University of Bergen, Bergen, Norway.

[10] Jamali, N., Zelko, J., & Osgood, N. D. (2025) ACT for Public Health: Current Practice, Recent Progress, and Future Opportunities. *Transactions on Applied Category Theory* (extended abstract).