A Nitrogen-Scaled POMP Model for Intraguild Predation Dynamics

Sijia Wang

Thesis Advisor: Prof. Edward L. Ionides

Research Mentor: Aaron J. Abkemeier

Abstract

This thesis analyzes time-series data from a tri-trophic chemostat system, as studied in Hiltunen et al. (2013). The original work proposes a deterministic ODE model to describe the population dynamics of the three species, but does not account for stochastic variability arising from ecological processes or measurement error.

In this thesis, we extend the deterministic ODE model into a state-space framework by formulating it as a Partially Observed Markov Process (POMP). In this formulation, the latent dynamics are represented as a stochastic dynamic system obtained by introducing stochastic perturbations to the original ODE system via an Euler discretization. A nitrogen-based scaling framework is employed to link latent states to observed data, and parameters are inferred via iterated filtering.

The main contributions of this work are: (i) a statistical reformulation of the model under a nitrogen-cycle interpretation that improves the interpretability of key ecological parameters, and (ii) a likelihood-based comparison highlighting the advantages of the mechanistic model over purely statistical approaches.

1 Introduction

Ecological communities exhibit temporal fluctuations driven by nonlinear interactions, environmental variability, and evolutionary processes. Chemostat systems are a controlled platform for studying these dynamics. Much work has investigated ecological dynamics in chemostat systems (Smith and Waltman (1995), Yoshida et al. (2003), Hiltunen et al. (2014)). This thesis examines the three-species food web studied by Hiltunen et al. (2013). The system exhibits intraguild predation, in which two consumers compete for a shared resource while one consumer also preys on the other. Here, rotifers and flagellates both consume algae, while rotifers also prey on flagellates.

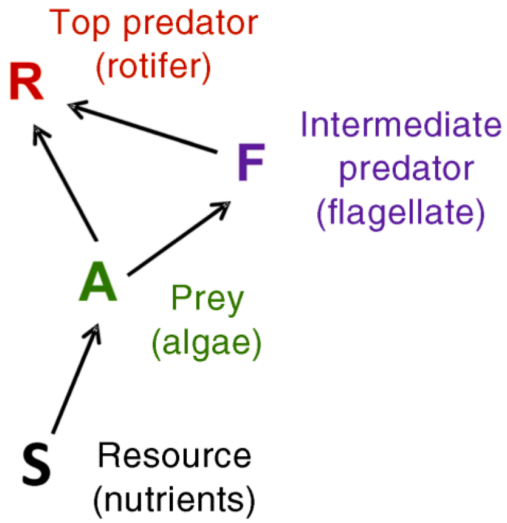

Figure 1: Structure of the three-species food web with intraguild predation, adapted from Hiltunen et al. (2013). Arrows represent uptake of the nutrient substrate (S) by the prey (the alga A, *Chlorella autotrophica*) and consumption of the prey by the intermediate predator (the flagellate F, *Oxyrrhis marina*) and the top predator (the rotifer R, *Brachionus plicatilis*).

A key challenge in analyzing ecological systems is that important processes are not directly observable. Assessing whether a mechanistic model adequately captures the structure of ecological time series is non-trivial, and diagnostic approaches such as those developed by Hooker and Ellner (2015) provide useful tools for model evaluation. Moreover, a mismatch often exists between the scale of latent ecological dynamics and that of observed data, making the specification of appropriate scaling relationships essential. In models with many ecological parameters, misspecification of the underlying mechanism can lead to identifiability issues in parameter estimation.

Partially Observed Markov Process (POMP) models provide a flexible statistical framework for representing such systems (E. L. Ionides, Bretó, and King (2006); Hooker and Ellner (2015)).

Modern plug-and-play inference algorithms, including iterated filtering (Edward L. Ionides et al. (2015)) and particle-based likelihood estimation, enable likelihood-based inference even when the process model is highly nonlinear. These methods are implemented in the `pomp` R package (King, Nguyen, and Ionides (2016)).

In this thesis, we develop a Partially Observed Markov Process (POMP) model for the tritrophic chemostat system studied by Hiltunen et al. (2013). A key contribution is a reformulation of the model under a nitrogen-cycle interpretation, assuming no evolutionary dynamics. We further demonstrate that the proposed formulation improves statistical performance through likelihood-based comparisons with baseline time-series models, while also yielding practically identifiable estimates for key ecological parameters.

2 Method

2.1 Model Structure

As described in Hiltunen et al. (2013), the chemostat experiments were conducted under controlled laboratory conditions, with nitrogen as the limiting nutrient. The system was maintained at a constant temperature (21°C) and salinity (35 g/L), and populations were sampled daily over a period sufficient to capture multiple oscillatory cycles.

The experimental setup motivates the formulation of the model on a nitrogen scale. Specifically, the state variables represent the amount of limiting substrate nitrogen contained in the substrate, algae, rotifers, and flagellates.

Accordingly, the model is formulated on two related scales: a nitrogen scale for the latent process and an observation scale for the measured abundances. Let S_n, A_n, R_n, and F_n denote the latent nitrogen quantities associated with the resource, algae, rotifers, and flagellates, respectively. To ensure positivity, the latent states are represented on the log scale throughout filtering and simulation.

The latent nitrogen states are mapped to the observation scale through scaling constants:

$$A_{\text{meas}} = A_{\text{sc}} A_n, \qquad R_{\text{meas}} = R_{\text{sc}} R_n, \qquad F_{\text{meas}} = F_{\text{sc}} F_n,$$

where A_{meas}, R_{meas}, and F_{meas} denote the conditional means of the observation process for algae, rotifers, and flagellates, respectively, and should not be confused with the observed data themselves.

For the substrate variable, the scaling is fixed at

$$S_{\text{sc}} = 1.$$

This choice is consistent with our definition of the nitrogen scale, in which the dissolved substrate is expressed in units of nitrogen concentration. Since S is not directly observed in the data, introducing an additional scaling parameter for S might lead to non-identifiability.

Appendix A of Hiltunen et al. (2013) suggests that one chemostat of nitrogen corresponds to approximately 2.5×10^6 algae cells/ml. Therefore, in the present model, the algae scaling constant is fixed at

$$A_{\text{sc}} = 2.5 \times 10^6.$$

The deterministic component of the latent process follows a non-evolutionary chemostat system adapted from Hiltunen et al. (2013):

$$\begin{aligned}
\frac{dS_n}{dt} &= \delta(1 - S_n) - \frac{rS_nA_n}{k_A + S_n}, \\
\frac{dA_n}{dt} &= A_n \left[\frac{rS_n}{k_A + S_n} - \frac{gR_n}{k_R + A_n + \alpha_F F_n} - \frac{hF_n}{1 + \alpha_A A_n} - \delta \right], \\
\frac{dR_n}{dt} &= R_n \left[\frac{gA_n}{k_R + A_n + \alpha_F F_n} + \frac{\eta F_n}{k_R + A_n + \alpha_F F_n} - \delta \right], \\
\frac{dF_n}{dt} &= F_n \left[\frac{hA_n}{1 + \alpha_A A_n} - \frac{\eta R_n}{k_R + A_n + \alpha_F F_n} - \delta \right] + I_F.
\end{aligned}$$

Detailed definitions of the ecological parameters are provided at the end of this section. Here, I briefly summarize the key mechanisms represented by the ODE system.

The substrate is continuously supplied through dilution at a rate of 0.25, while species interactions are governed by Type II functional responses as described by Holling (1959). Specifically, these include trophic interactions between rotifers and algae, flagellates and algae, and intraguild interactions between rotifers and flagellates.

As discussed in Hiltunen et al. (2013), flagellates may go extinct at low population levels during oscillatory cycles. To prevent this, a small external inflow of flagellates was introduced in the experiment. This is incorporated into the model through the inflow term I_F in the flagellate equation.

But a related units issue arises in the treatment of the flagellate inflow. The experimentally controlled quantity is the flagellate injection rate, approximately 10^4 cells per day into a 380 ml chemostat (Hiltunen et al. (2013)). Expressed on the measurement scale, this becomes

$$IF_{\text{meas}} = \frac{10^4}{380} \qquad \text{cells/ml/day}.$$

To maintain consistency with the nitrogen-scale process model, the inflow parameter in the latent system is defined by converting this measured inflow into nitrogen units:

$$IF = \frac{IF_{\text{meas}}}{F_{\text{sc}}}.$$

Under this formulation, F_{sc} controls the conversion between measurement units and nitrogen units for flagellates, while the experimentally known injection rate is treated as fixed.

To account for stochasticity in the system dynamics, process noise is introduced on the log scale of the latent states. After each Euler update of the deterministic system, Gaussian noise is added to the log-transformed states.

Specifically, for each state variable $X_n \in \{S_n, A_n, R_n, F_n\}$, the update is given by

$$\log X_n(t + \Delta t) = \log\Big(X_n(t) + \frac{dX_n}{dt}\Big|_t \Delta t \Big) + \varepsilon_X,$$

where $\frac{dX_n}{dt}$ is given by the corresponding right-hand side of the ODE system defined above.

where ε_X represents Gaussian process noise accounting for stochastic fluctuations in the latent dynamics, with mean zero and variance $\sigma^2_{\text{pro},X}\Delta t$:

$$\varepsilon_X \sim \mathcal{N}\big(0, \sigma^2_{\text{pro},X}\,\Delta t\big)\,.$$

The latent process is implemented via an Euler-discretized stochastic dynamic system with time step $\Delta t = 0.25$. Small positive lower bounds are imposed during state updating to ensure numerical stability.

The latent process is initialized using the day 40 measurements for algae, rotifers, and flagellates, mapped to the nitrogen scale through the corresponding scaling constants. The initial resource state is then set to balance the total nitrogen:

$$S_n(0) = 1 - A_n(0) - R_n(0) - F_n(0).$$

Conditional on the latent state, the observed abundances of algae, rotifers, and flagellates are assumed to be independent and lognormally distributed around the corresponding conditional means on the observation scale. This choice was found empirically to provide a reasonable fit to the observed abundances. Observations are collected at discrete time points corresponding to days 41–107.

$$\begin{aligned}
A_{\text{obs}}(t_n) \mid A_{\text{meas}}(t_n) &\sim \text{lognormal}(\log A_{\text{meas}}(t_n), \sigma_A^2),\\
R_{\text{obs}}(t_n) \mid R_{\text{meas}}(t_n) &\sim \text{lognormal}(\log R_{\text{meas}}(t_n), \sigma_R^2),\\
F_{\text{obs}}(t_n) \mid F_{\text{meas}}(t_n) &\sim \text{lognormal}(\log F_{\text{meas}}(t_n), \sigma_F^2).
\end{aligned}$$

Here, σ_A, σ_R, and σ_F denote the standard deviations of the measurement noise on the log scale, governing the variability of the observations around their conditional means.

Finally, the parameter set consists of:

(i) ecological parameters governing the deterministic dynamics,

(ii) measurement noise parameters,

(iii) process noise parameters,

(iv) scaling constants linking the nitrogen-scale states to the observation scale,

(v) initial state values.

Table 1: Model parameters, interpretations, and units.

Category	Symbol	Description	Units
Ecological	r	Substrate uptake rate	day^{-1}
Ecological	g	Rotifer grazing rate	day^{-1}
Ecological	h	Flagellate grazing rate	day^{-1}
Ecological	η	Intraguild predation rate	day^{-1}
Ecological	k_A	Half-saturation constant for substrate uptake	dimensionless
Ecological	k_R	Half-saturation constant for rotifer feeding / intraguild interaction	dimensionless
Ecological	α_A	Algal handling parameter in flagellate grazing	dimensionless
Ecological	α_F	Flagellate contribution to rotifer grazing denominator	dimensionless
Ecological	δ	Dilution rate	day^{-1}
Ecological	I_F	Flagellate inflow on the nitrogen scale	day^{-1}
Measurement noise	σ_A	Measurement noise for algae on the log scale	dimensionless
Measurement noise	σ_R	Measurement noise for rotifers on the log scale	dimensionless
Measurement noise	σ_F	Measurement noise for flagellates on the log scale	dimensionless
Process noise	$\sigma_{\text{pro},S}$	Process noise for substrate on the log scale	dimensionless
Process noise	$\sigma_{\text{pro},A}$	Process noise for algae on the log scale	dimensionless
Process noise	$\sigma_{\text{pro},R}$	Process noise for rotifers on the log scale	dimensionless
Process noise	$\sigma_{\text{pro},F}$	Process noise for flagellates on the log scale	dimensionless
Scaling	A_{sc}	Scaling constant mapping algae nitrogen state to observation scale	cells mL^{-1} per nitrogen unit
Scaling	R_{sc}	Scaling constant mapping rotifer nitrogen state to observation scale	individuals mL^{-1} per nitrogen unit
Scaling	F_{sc}	Scaling constant mapping flagellate nitrogen state to observation scale	cells mL^{-1} per nitrogen unit
Scaling	S_{sc}	Scaling constant for substrate	dimensionless
Initial condition	A_0	Initial observed algae abundance	cells mL^{-1}
Initial condition	R_0	Initial observed rotifer abundance	individuals mL^{-1}
Initial condition	F_0	Initial observed flagellate abundance	cells mL^{-1}

2.2 Parameter Inference

Parameter estimation is performed using iterated filtering (IF2; E. L. Ionides, Bretó, and King (2006), Edward L. Ionides et al. (2011)) in a two-stage procedure, implemented via the mif2 function in the R package pomp (King, Nguyen, and Ionides (2016)).

Iterated filtering is a stochastic optimization method (Edward L. Ionides et al. (2015)). The algorithm combines particle filtering (sequential Monte Carlo) with stochastic perturbations of parameters. The particle filter sequentially estimates the integrals in the prediction and filtering recursions, enabling likelihood-based inference in models with latent states.

Intuitively, parameters are treated as evolving stochastically across iterations, allowing the algorithm to explore the parameter space while repeatedly applying particle filtering to evaluate model fit. As the magnitude of the perturbations is gradually reduced, the parameter estimates concentrate near regions of high likelihood, providing an approximation to the maximum likelihood estimate.

In the first stage, a local search is conducted using six independent mif2 runs. The starting values for the local search are chosen empirically. Each run uses 1000 particles and 25 iterations, with a moderate cooling schedule for parameter perturbations. All runs are initialized at the same starting point but incorporate independent random perturbations, allowing exploration of nearby regions of the parameter space.

In the second stage, a global search is performed by generating 800 candidate parameter vectors. Each free parameter is independently sampled from a uniform interval centered at the best local estimate, with bounds given by one-fifth to five times that estimate. The dilution rate δ and the algae and resource scaling constants are held fixed throughout. Each candidate is then refined using an additional 40 mif2 iterations with 2000 particles and a moderate cooling schedule.

The log-likelihood of each fitted model is evaluated using repeated particle filtering with 5000 particles and 30 independent replicates. The reported log-likelihood is computed using logmeanexp, which provides a Monte Carlo estimate together with an associated standard error.

To improve computational efficiency, subsequent searches are initialized from parameter values in high-likelihood regions identified in earlier runs, rather than from the original starting point.

3 Results

To provide a baseline for comparison, the POMP model is evaluated against a standard time-series model, namely an ARMA (autoregressive–moving-average) model fitted to the log-transformed observations. The ARMA model is estimated using an automated model selection procedure based on the Akaike Information Criterion (AIC), as implemented in the forecast package developed by Hyndman and Khandakar (2008). Its predictive performance is evaluated based on how well it predicts each observation using information from previous time points (i.e., one-step-ahead prediction). To ensure a fair comparison, both models are evaluated using the same information set at each time point. Note that although the POMP model is defined on a latent nitrogen scale, the likelihood is computed on the observation scale through the measurement model. This ensures that both the POMP and ARMA models are compared based on their predictive performance for the observed data.

We assess species-specific model fit by examining the conditional log likelihood contributions at each time point for algae, flagellates, and rotifers.

The conditional likelihood of the observations at time t is given by

$$\begin{aligned} &p(A_{\text{obs}}(t), F_{\text{obs}}(t), R_{\text{obs}}(t) \mid A_{\text{obs}}(1{:}t-1), F_{\text{obs}}(1{:}t-1), R_{\text{obs}}(1{:}t-1)) \\ &= \int p(A_{\text{obs}}(t), F_{\text{obs}}(t), R_{\text{obs}}(t) \mid x_t)\, p(x_t \mid A_{\text{obs}}(1{:}t-1), F_{\text{obs}}(1{:}t-1), R_{\text{obs}}(1{:}t-1))\, dx_t. \end{aligned}$$

Although the three species are jointly modeled through a shared latent process in the POMP framework, the likelihood separates at the level of the measurement model due to the assumption of conditional independence across species. This follows from the measurement model, in which each observed abundance is generated by applying an independent lognormal observation error to its corresponding latent mean. This formulation reflects the idea that measurement errors arise from separate observation processes for each species, while ecological dependence is captured in the latent dynamics.

Specifically,

$$p(A_{\text{obs}}(t), F_{\text{obs}}(t), R_{\text{obs}}(t) \mid x_t) = p(A_{\text{obs}}(t) \mid x_t)\, p(F_{\text{obs}}(t) \mid x_t)\, p(R_{\text{obs}}(t) \mid x_t),$$

where each component follows a lognormal distribution centered at the corresponding latent mean.

Taking logarithms, the conditional log-likelihood can be written as

$$\log p(A_{\text{obs}}(t), F_{\text{obs}}(t), R_{\text{obs}}(t) \mid x_t) = \log p(A_{\text{obs}}(t) \mid x_t) + \log p(F_{\text{obs}}(t) \mid x_t) + \log p(R_{\text{obs}}(t) \mid x_t).$$

As a result, the total log likelihood at time t decomposes numerically as

$$\ell_t^{\text{total}} = \ell_t^{(A)} + \ell_t^{(F)} + \ell_t^{(R)}.$$

in this setting. This decomposition reflects the factorization of the measurement density under conditional independence, together with the additive structure of the log-likelihood. The latent ecological dynamics, however, remain fully coupled across species.

Table 2: Comparison of total conditional log-likelihoods.

Channel	POMP logLik	ARMA logLik	Difference
Algae	-833.0	-832.3	-0.7
Flagellates	-455.1	-461.8	6.8
Rotifers	-69.6	-157.6	88.0
Total	-1357.6	-1451.7	94.0

The table above summarizes the total conditional log-likelihood for each channel under the POMP and ARMA models. Overall, the POMP model improves the total conditional log-likelihood by +94.0 compared to the ARMA model, providing evidence in this dataset that incorporating mechanistic ecological structure can improve predictive performance. The improvement is particularly pronounced for the rotifer dynamics, where the POMP model achieves a substantially higher log-likelihood (+88.0). For flagellates, a moderate improvement (+6.8) is observed, while for algae the two models perform similarly, with a small disadvantage for the POMP model (-0.7).

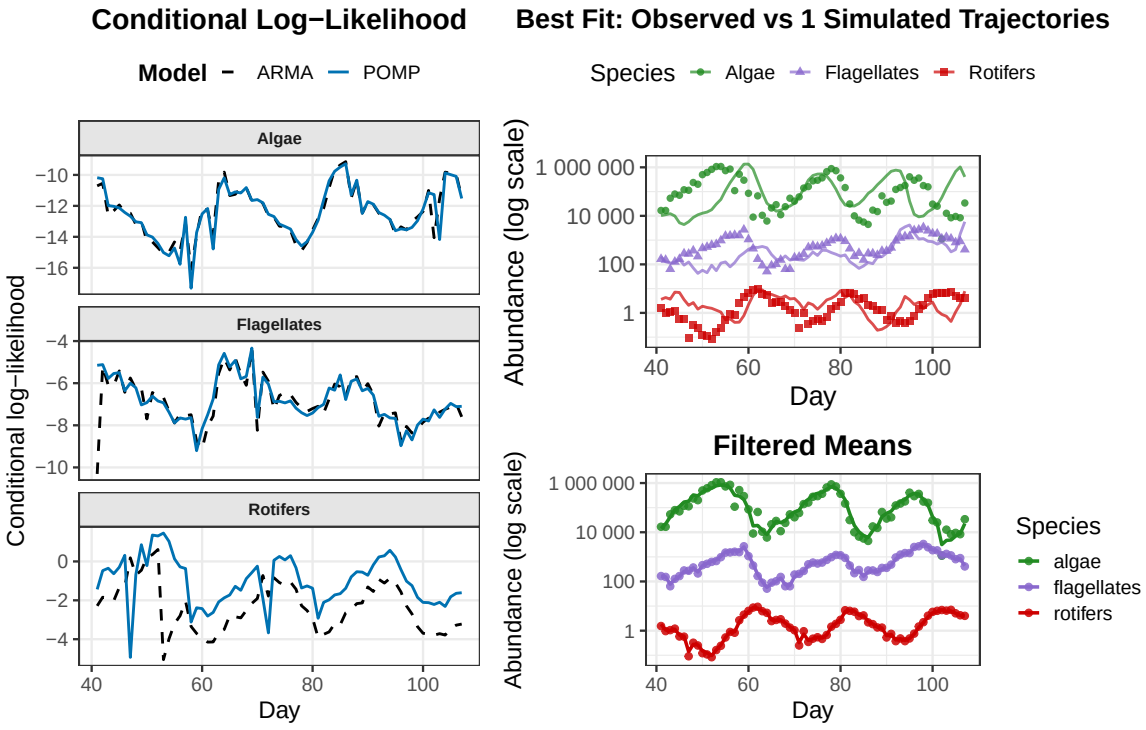

Figure 2: Conditional log-likelihood comparison (left), simulated trajectories (top right), and filtered means (bottom right).

As shown in the left panel of Figure 2, we compare the conditional log-likelihoods obtained from the ARMA and POMP models for each species. Although the overall temporal patterns are similar, the POMP model generally provides a better fit than the ARMA model in this case.

Table 3: Overall model comparison using conditional log-likelihood and AIC.

model	logLik	k	AIC
POMP	-1357.64	20	2755.27

model	logLik	k	AIC
ARMA	-1451.68	9	2921.35

The higher conditional log-likelihood suggests that the POMP model provides a better fit to the data. However, because it also involves more parameters than the ARMA benchmark, it is important to assess whether this improvement persists after accounting for model complexity. We therefore compare the models using the Akaike Information Criterion (AIC). The POMP model achieves a substantially lower AIC (2755.27) than the ARMA model (2921.35), indicating that the gain in log-likelihood more than offsets the penalty for increased model complexity.

This further supports the advantage of incorporating mechanistic ecological structure in modeling the system dynamics. As shown in the top-right panel of Figure 2, a representative simulated trajectory captures the broad magnitude and oscillatory patterns of the observed data. The filtered means in the bottom-right panel track the observations more closely, reflecting the model's ability to combine mechanistic dynamics with information from the data.

Furthermore, to more rigorously assess parameter identifiability, we conducted a profile likelihood analysis. This approach allows us to evaluate how well individual parameters are supported by the data while accounting for potential interactions with other parameters in the model. The table below summarizes the parameter definitions, maximum likelihood estimates, and corresponding 95% profile likelihood confidence intervals for the model.

Table 4: Parameter definitions, MLEs, and 95% profile likelihood confidence intervals for estimated parameters.

Parameter	Description	Value	CI
kA	Half-saturation constant	7.64	(7.27, 7.64)
kR	Prey half-saturation constant	7.1	(6.81, 7.1)
r	Substrate uptake rate	3.07	(1.87, 12)
g	Rotifer grazing rate	11.2	(9.19, 14.1)
h	Flagellate grazing rate	4.01	(1.44, 5.69)
alphaA	Algal handling parameter	2.75	(1.18, 4.12)
eta	Maximum grazing rate	5.33	(0.592, 5.33)
A_0	Initial algae	6.84e+03	(5.27e+03, 2.09e+04)
R_0	Initial rotifers	2.02	(1.38, 2.72)
F_0	Initial flagellates	135	(77.8, 176)
sigma_A	Measurement noise (algae)	0.491	(0.443, 0.628)
sigma_F	Measurement noise (flagellates)	0.143	(0.114, 0.168)
sigma_R	Measurement noise (rotifers)	0.371	(0.198, 0.578)

sigma_proS	Process noise (substrate)	0.231	(0.159, 0.384)
sigma_proA	Process noise (algae)	0.0134	(0.0103, 0.0672)
sigma_proF	Process noise (flagellates)	0.411	(0.307, 0.572)
sigma_proR	Process noise (rotifers)	0.607	(0.126, 0.797)
R_sc	Rotifer scaling constant	16.2	(12, 26.8)
F_sc	Flagellate scaling constant	8.62e+03	(4.02e+03, 3.62e+04)

The profile likelihood results further clarify the identifiability of model parameters. Several key ecological parameters, such as the half-saturation constants (k_A, k_R) and grazing rates (g, h), are reasonably well constrained and lie within biologically plausible ranges. But some functional responses, such as r and η, exhibit wider confidence intervals, indicating weaker identifiability. Initial conditions and certain scaling constants also remain weakly identified, reflecting their indirect connection to the observed measurements and the limited information available to constrain them.

Notably, both measurement and process noise parameters are estimated with relatively narrow confidence intervals, corresponding to low levels of stochastic variability. This may indicate that the model does not rely heavily on stochastic noise to explain the data, and instead that the mechanistic structure itself captures the main features of the system dynamics.

4 Discussion

The nitrogen-based formulation was not adopted at the initial stage of model development. Instead, the process model was initially specified on a standardized scale, with scaling constants implicitly determined by the maximum observed abundance of each species. We also explored an alternative formulation in which scaling is incorporated into the functional responses, allowing the process model to be expressed directly on the observation scale. While both approaches are practically workable, neither provides a fully coherent representation from the perspective of nutrient balance.

As noted in Hiltunen et al. (2013), nitrogen is the limiting nutrient in the system. The model is therefore naturally formulated on a nitrogen scale, where all state variables represent the amount of nitrogen contained in each compartment. The system is normalized so that the total nitrogen in the chemostat equals one unit at steady state. Under this formulation, all state variables can be interpreted as proportions of the total nitrogen, measured in "chemostat units of nitrogen." This perspective motivates the nitrogen-cycle formulation adopted in this thesis, which provides a more coherent and mechanistically interpretable representation of the underlying dynamics.

Several directions for future work remain. The dataset analyzed in this study was generated under experimental conditions designed to limit evolutionary change. Nevertheless, it remains

of interest to examine whether models incorporating prey evolutionary dynamics can provide a better statistical description of the observed time series. Rapid evolution can occur in such systems and influence population dynamics, as demonstrated by Yoshida et al. (2003). The tri-trophic eco-evolutionary chemostat system studied by Hiltunen et al. (2014)—comprising two algal genotypes and two predators—provides a natural setting for investigating how resource competition, predation, and rapid evolution jointly shape population trajectories.

While eco-evolutionary extensions are conceptually appealing, the present dataset does not provide sufficient information to reliably identify evolutionary processes. Incorporating algae genotype dynamics would likely require additional data or experimental design tailored to detect evolutionary change.

Another limitation of the current study is that the analysis is based on a single experimental dataset. Although the model captures key features of the observed dynamics, further validation across multiple systems would be necessary to assess the generality and robustness of the proposed framework.

5 Conclusion

In this thesis, we develop a Partially Observed Markov Process (POMP) model for a chemostat system with intraguild predation, building on the deterministic ODE framework of Hiltunen et al. (2013). By extending the original model into a state-space formulation, we incorporate stochastic variability arising from both ecological processes and measurement error, enabling likelihood-based inference for complex nonlinear dynamics. Empirical results indicate that the proposed model provides improved statistical performance relative to baseline time-series approaches, while yielding more interpretable parameter estimates.

A central contribution of this work is to demonstrate that integrating mechanistic models with time-series inference can provide reliable and ecologically interpretable insights from noisy observational data.

Acknowledgement

I would like to thank Professor Stephen Ellner and Professor Giles Hooker for their insightful suggestions. I also thank Kunyang He and Ziheng Wei for their helpful discussions.

References

Hiltunen, Teppo, Stephen P. Ellner, Giles Hooker, Laura E. Jones, and Nelson G. Hairston Jr. 2014. “Eco-Evolutionary Dynamics in a Three-Species Food Web with Intraguild

Predation: Intriguingly Complex." In *Advances in Ecological Research*, edited by Jordi Moya-Laraño, Jennifer Rowntree, and Guy Woodward, 50:41–73. Academic Press. https://doi.org/10.1016/B978-0-12-801374-8.00002-5.

Hiltunen, Teppo, L. E. Jones, Stephen P. Ellner, and Nelson G. Hairston. 2013. "Temporal Dynamics of a Simple Community with Intraguild Predation: An Experimental Test." *Ecology* 94 (4): 773–79.

Holling, C. S. 1959. "Some Characteristics of Simple Types of Predation and Parasitism." *The Canadian Entomologist* 91 (7): 385–98.

Hooker, Giles, and Stephen P. Ellner. 2015. "Goodness of Fit in Nonlinear Dynamics: Misspecified Rates or Misspecified States?" *The Annals of Applied Statistics* 9: 754–76.

Hyndman, Rob J., and Yeasmin Khandakar. 2008. "Automatic Time Series Forecasting: The Forecast Package for R." *Journal of Statistical Software* 27 (3): 1–22. https://doi.org/10.18637/jss.v027.i03.

Ionides, E. L., Carles Bretó, and Aaron A. King. 2006. "Inference for Nonlinear Dynamical Systems." *Proceedings of the National Academy of Sciences* 103 (49): 18438–43. https://doi.org/10.1073/pnas.0603181103.

Ionides, Edward L., Anindya Bhadra, Yves Atchadé, and Aaron A. King. 2011. "Iterated Filtering." *The Annals of Statistics* 39 (3): 1776–1802.

Ionides, Edward L., Dennis Nguyen, Yves Atchadé, Stilian Stoev, and Aaron A. King. 2015. "Inference for Dynamic and Latent Variable Models via Iterated, Perturbed Bayes Maps." *Proceedings of the National Academy of Sciences* 112 (3): 719–24. https://doi.org/10.1073/pnas.1410597112.

King, Aaron A., Dao Nguyen, and Edward L. Ionides. 2016. "Statistical Inference for Partially Observed Markov Processes via the r Package Pomp." *Journal of Statistical Software* 69 (12): 1–43. https://doi.org/10.18637/jss.v069.i12.

Smith, Hal L., and Paul Waltman. 1995. *The Theory of the Chemostat: Dynamics of Microbial Competition*. Cambridge: Cambridge University Press. https://doi.org/10.1017/CBO9780511530043.

Yoshida, Takehito, Laura E. Jones, Stephen P. Ellner, Gregor F. Fussmann, and Nelson G. Hairston Jr. 2003. "Rapid Evolution Drives Ecological Dynamics in a Predator–Prey System." *Nature* 424: 303–6. https://doi.org/10.1038/nature01767.