

# MECHANISTIC MODELS FOR PANEL DATA: ANALYSIS OF ECOLOGICAL EXPERIMENTS WITH FOUR INTERACTING SPECIES

BY BO YANG<sup>1,a</sup>, JESSE WHEELER<sup>2,b</sup>, MEGHAN A. DUFFY<sup>3,d</sup>, AARON A. KING<sup>3,e</sup> AND EDWARD IONIDES<sup>4,c</sup>

<sup>1</sup>DEPARTMENT OF MARKETING, COLUMBIA UNIVERSITY, [ybb@umich.edu](mailto:ybb@umich.edu)

<sup>2</sup>DEPARTMENT OF MATHEMATICS AND STATISTICS, IDAHO STATE UNIVERSITY, [jeswheel@umich.edu](mailto:jeswheel@umich.edu)

<sup>3</sup>DEPARTMENT OF ECOLOGY AND EVOLUTIONARY BIOLOGY, UNIVERSITY OF MICHIGAN

<sup>4</sup>DEPARTMENT OF STATISTICS, UNIVERSITY OF MICHIGAN, [ionides@umich.edu](mailto:ionides@umich.edu)

, [duffymeg@umich.edu](mailto:duffymeg@umich.edu); [kingaa@umich.edu](mailto:kingaa@umich.edu)

In an ecological context, panel data arise when time series measurements are made on a collection of ecological processes. Each process may correspond to a spatial location for field data, or to an experimental ecosystem in a designed experiment. Statistical models for ecological panel data should capture the high levels of nonlinearity, stochasticity, and measurement uncertainty inherent in ecological systems. Furthermore, the system dynamics may depend on unobservable variables. This study applies iterated particle filtering techniques to explore new possibilities for likelihood-based statistical analysis of these complex systems. We analyze data from a mesocosm experiment in which two species of the freshwater planktonic crustacean genus, *Daphnia*, coexist with an alga and a fungal parasite. Time series data were collected on replicated mesocosms under six treatment conditions. Iterated filtering enables maximization of the likelihood for scientifically motivated nonlinear partially observed Markov process models, providing access to standard likelihood-based methods for parameter estimation, confidence intervals, hypothesis testing, model selection and diagnostics. This toolbox allows scientists to propose and evaluate scientifically motivated stochastic dynamic models for panel data, constrained only by the requirement to write code to simulate from the model and to specify a measurement distribution describing how the system state is observed. Our iterated filtering framework provides a statistically efficient, likelihood-based solution, avoiding the information loss inherent when using the generalized method of moments or Approximate Bayesian Computing approaches. Existing full-information Bayesian tools are not applicable to panels of partially observed Markov process models with intractable transition densities.

---

*Keywords and phrases:* Panel Iterated Filtering, SIRJPF model, PanelPOMP, *Daphnia*.

**1. Introduction.** Biological dynamic systems often contain elements or processes that are difficult to observe, even under controlled experimental conditions. These unobservable, or latent, components are often crucial to understanding system dynamics, or are of scientific interest themselves. Additionally, both the measurements obtained from the systems in question and their underlying dynamic processes often exhibit substantial stochasticity. The combination of these two factors complicates data interpretation and analysis in ecological studies. Given these inherent properties of ecological systems, it is essential to use models that can account for key latent mechanisms and the stochastic nature of the system when conducting inference (Young, 1998; Bjørnstad and Grenfell, 2001).

Partially observed Markov process (POMP) models, also known as hidden Markov models or state space models (Bretó et al., 2009; Auger-Méthé et al., 2021), comprise a class of models suited to the description of such systems, which describe the latent process using a Markovian representation. A POMP model is mechanistic if it is based on scientific principles, which is usually equivalent to providing a causal description of the system that can explain the consequence of interventions. Statistical inferences on mechanistic models can provide insight into key latent mechanisms. Further, statistical tests between rival mechanistic hypotheses can inform qualitative understanding about suitable model structures. Thus, mechanistic models are invaluable tools for addressing challenges in ecology (Mouquet et al., 2015).

Collections of time series on related but disjoint systems are called panel data, also known as longitudinal data. Each component system is called a unit, and the units may be sites in a field study (e.g., plots, lakes) or independent microcosms (e.g., test tube experiments) or mesocosms (e.g., experimental ecosystems in buckets or small pools). Studying the entire panel may allow precision in statistical conclusions that is unattainable from a single unit. Alternatively, scientists may be interested in specific differences between the units, requiring the investigation of the entire panel of data. For either of these goals, investigation of panel data leads to the challenging statistical task of fitting a collection of related nonlinear or non-Gaussian vector-valued POMP models, called a PanelPOMP model (Bretó et al., 2020; Breto et al., 2025). In this article, we demonstrate how recent methodological advances can facilitate data analysis for such models in a multi-species system. We use data from a controlled experiment involving the population dynamics of two freshwater plankton species (*Daphnia dentifera* and *Daphnia lumholtzi*) together with an alga, *Ankistrodesmus falcatus*, as food and a fungal parasite, *Metschnikowia bicuspidata*, that can infect both *Daphnia* species. These data were collected by Searle et al. (2016a) to investigate how the competition between the North American native *D. dentifera* and the invasive *D. lumholtzi* is modified by their differing susceptibility to the parasite. Advances in data collection make this experiment representative of a growing class of situations where panel data is available on a complex multi-species system.

Our model builds upon the ordinary differential equation model of Searle et al. (2016a). We include dynamic stochasticity, leading to a stochastic differential equation model which provides an improved statistical fit to the data while also assisting identification of model misspecification (King et al., 2015). We carry out likelihood-based inference using panel iterated filtering (Bretó et al., 2020) which is a PanelPOMP extension of basic iterated filtering methodology (Ionides et al., 2015; King et al., 2016) for POMP models. Iterated filtering methods have a plug-and-play property (He et al., 2010) that they require a simulator for the dynamic model but do not need the ability to evaluate transition densities. The plug-and-play property permits practical consideration of

a broad class of PanelPOMP models, assisting the creative task of developing scientifically plausible models that simultaneously have good statistical fit. POMP models have long been advocated as an appropriate framework for ecological systems (Schnute, 1994; Buckland et al., 2004). A major practical limitation has been the computational difficulty in carrying out inference. When Laplace approximations are applicable, the Template Model Builder approach is available (Kristensen et al., 2016). Iterated filtering techniques apply to a more general class of POMP models, but the high-dimensional nature of PanelPOMP models adds additional methodological challenges. Full-likelihood Bayesian methods cannot readily accommodate the nonlinear, non-Gaussian mechanistic models considered here. Pseudomarginal Markov chain Monte Carlo methods are applicable in principle, but have limited practical applicability for this model class; see Section 2 of Häggström et al. (2026) for an extended discussion of this. We are not aware of any previous scientific study where nonlinear, non-Gaussian PanelPOMP models (with some parameters shared between units and some parameters specific to each unit) have been calibrated to multi-species panel data via maximum likelihood or pseudomarginal Bayesian inference.

A central question when using model-based statistical inference to investigate ecological dynamics is to discover which biological details are necessary to include in a model to provide an adequate statistical explanation for the data. Our approach to answering this question uses standard tools of likelihood-based inference. We compare the model likelihood to statistical benchmarks, such as linear or autoregressive models, seeking to continue improving a mechanistic model until it captures the predictability of the data with at least comparable skill to a simple non-mechanistic analysis (Wheeler et al., 2024; Li et al., 2024). In doing so, we ensure that the proposed dynamic model provides a plausible mechanistic description of the dynamic system, supported by evidence that the model also provides a suitable quantitative description of data that arise from the system.

To do this, we study residuals and conditional log-likelihood anomalies for each data point to identify relative weaknesses and strengths of different models. We calculate Akaike’s Information Criterion (AIC) to compare the overall fit of rival mechanistic models (Akaike, 1974). We compute profile likelihood confidence intervals, making appropriate adjustment for Monte Carlo uncertainty (Ionides et al., 2017; Ning et al., 2021). All these techniques are predicated on an ability to evaluate and maximize the log-likelihood function. This ability was unlocked by panel iterated filtering (Bretó et al., 2020) and we demonstrate empirically that this is further strengthened by a new marginalized panel iterated filtering algorithm. A contemporaneous investigation of marginalized panel iterated filtering from a theory and methodology perspective is provided by Wheeler et al. (2025). These algorithms are discussed in more detail in Section 4, after description of the experimental system in Section 2 and presentation of our models and their fitted parameter values in Section 3. Section 5 presents additional results showing that our stochastic model, that describes depletion of the algal food resource together with age-structured population dynamics, provides a better statistical explanation of the data than the previous deterministic model by Searle et al. (2016a) that supposed the dynamics are driven by a switch of reproductive strategies from the asexual production of juveniles to the sexual production of long-lasting dormant eggs. Our findings have consequences for the interpretation of this experiment and the planning of future experiments. Section 7 is a concluding discussion.

**2. A *Daphnia* mesocosm experiment.** *Daphnia* is a genus of filter-feeding zooplankton commonly known as “water fleas.” This diverse genus is a dominant grazer of algae and bacteria in freshwater lakes and rivers worldwide. *D. dentifera* is native to North American stratified lakes, whereas *D. lumholtzi* is an invasive species originating from Africa, Asia, and Australia. First observed in North America in the 1990s, *D. lumholtzi* has rapidly expanded across the United States, impacting native species and ecosystem dynamics (Havel and Hebert, 1993; Searle et al., 2016b). This leads to ecological interest in understanding competition between *D. dentifera* and *D. lumholtzi*, two species with high evolutionary divergence having their most recent common ancestor around 145 million years ago (Kotov and Taylor, 2011; Cornetti et al., 2019). Here, we investigate how this competition is affected by a parasite that infects both species. Searle et al. (2016a) hypothesized that the presence of a fast-growing invasive species, which has a high susceptibility to the native parasite *M. bicuspidata*, could amplify the parasite in the native population.

Female *Daphnia* reproduce asexually in favorable situations, in a process called *parthenogenesis* whereby females produce embryos that develop without fertilization. Maturation of juveniles to adults takes approximately five to ten days, depending on environmental conditions. In response to adverse environmental cues, *Daphnia* produce males, and females reproduce sexually to produce durable resting eggs called ephippia (Radzikowski et al., 2018; Searle et al., 2016b). The ephippia can remain viable for many years, awaiting the return of favorable conditions.

*Daphnia* become infected by the parasite by ingesting free-floating spores as they filter the water to extract food. The spores pass through the gut wall to infect the body cavity of the *Daphnia*, where the fungus grows until the body cavity is full of spores and the host dies, releasing the spores into the water. Infected *Daphnia* die within two to three weeks (Ebert, 2005), whereas healthy *Daphnia* can survive for many months. Infection leads to sterility and reduced feeding, and an opaque body cavity which can be readily distinguished from the transparent cavity of uninfected *Daphnia*.

Searle et al. (2016a) conducted an experiment involving a collection of mesocosms, each consisting of 15 liter of growth medium in a bucket, initialized with  $2.5 \times 10^8$  cells of green alga, *A. falcatus*, and with one or two species of *Daphnia*. Six treatments were carried out, comprised of three levels of the host (*D. dentifera* or *D. lumholtzi* or both) factored with two levels of the parasite (present or absent). The experimental units are replicates of the treatments; the design consisted of 8, 9 or 10 units for each treatment. The initial *Daphnia* population consisted of 45 adult females, which in two-host treatments consisted of 35 native *Daphnia* and 10 invasives. After a four day acclimatization period, spores of the fungal parasite, *M. bicuspidata*, were added at a concentration of 25 per mL if the mesocosm was assigned to the parasite exposure treatment. Every five days, a well-mixed sample of 1 liter was removed, and the species (native or invasive), infection status, sex, and age (juvenile or adult) were identified for each *Daphnia* using a microscope. A negligible number of infected juveniles were observed so henceforth we model infection only in adults. Algal and nutrient levels were supplemented twice per week, with medium added as necessary to maintain the volume. The mesocosms were kept inside, at  $23.3^\circ C$  and a 16-hour light, 8-hour dark cycle. Population densities and infection prevalence (if appropriate) were quantified every five days, beginning a week after the experiment started. Regular stirring of mesocosms ensured resuspension of algae and parasites. The experiment concluded after 52 days, totaling 10 sampling sessions.

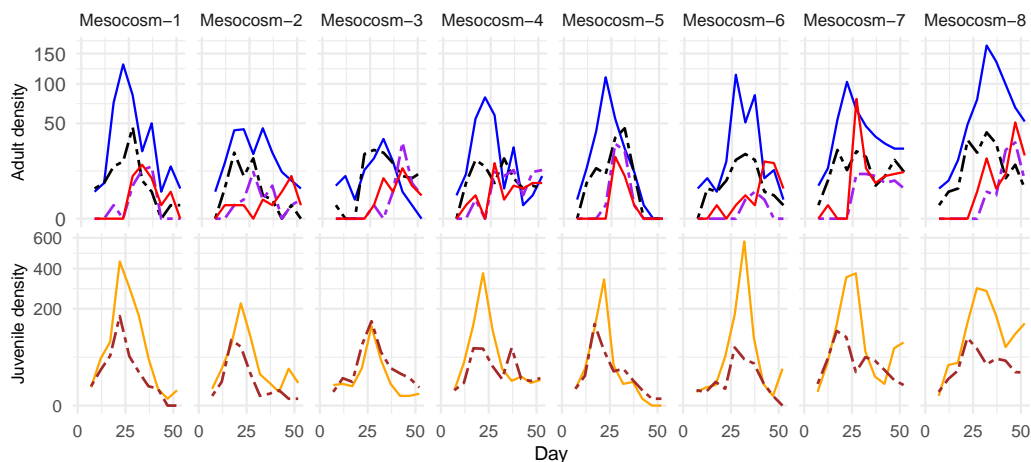


Figure 1: The experimental data. Density (Individuals/Liter) of *D. dentifera* (solid lines) and *D. lumholtzi* (dashed lines). The top panel shows adult female susceptibles (*D. dentifera*, blue; *D. lumholtzi*, black) and infecteds (*D. dentifera*, red; *D. lumholtzi*, purple). The bottom panel shows juvenile susceptibles (*D. dentifera*, orange; *D. lumholtzi*, brown). There were negligible infected juveniles and males. Columns are buckets corresponding to replications with same treatment setting.

Here, we focus on the most complex experimental treatment, namely, the four species treatment with both host species, the alga and the parasite. The simpler treatments are presented in Supplementary Sections S4, S5 and S6. Data from the four species treatment are displayed in Fig. 1. The females and juveniles show a population peak around day 25, typically followed by a collapse and a mild resurgence around day 50. The males arise predominantly later in the experiment as shown in Table S12. The hypothesis that motivated the model of Searle et al. (2016a) is that the arrival of the males, and the subsequent switch to the generation of ephippia by sexual reproduction, explains the decay of the other populations. This hypothesis is related to resource depletion, since shortage of food and high population density are both triggers for a change in reproductive strategy (Radzikowski et al., 2018; Searle et al., 2016b). However, resource depletion affects parthenogenetic reproduction as well influencing a switch to sexual reproduction. Therefore, we investigate the hypothesis that explicit modeling of resource depletion can help to explain the observed dynamics, via the introduction of a latent food compartment. We also hypothesize that the inclusion of age structure, via a juvenile stage, could help to explain the population dynamics even though very few juveniles show disease. Our paradigm is to assess models that represent known or hypothesized biological mechanisms in terms of their ability to explain observed data, quantified by likelihood penalized by model complexity. Our hypothesized model is supported so far as it is the best model for the available data according to this criterion.




Figure 2: A flow diagram for the SIRJPF2 model. The two species of *Daphnia* interact with each other indirectly through  $F$  (Food) and  $P$  (Parasite) compartments. Transition to the  $R$  state represents death or removal of any species.

**3. The Susceptible-Infected-Removed-Juvenile-Parasite-Food 2-species model (SIRJPF2).** Fig. 2 shows a flow diagram for a two-species model where the *Daphnia* are susceptible adult female ( $S$ ), infected adult female ( $I$ ), or juvenile ( $J$ ), with a superscript  $i$  for invasive and  $n$  for native. **Male *Daphnia* are excluded from this model because they are absent until late in the experiment, are physically smaller than females, and have low observed infection rates.** We keep track of the aquatic parasite spore density ( $P$ ) and the food density ( $F$ ). Removal (i.e., death, consumption or destructive sampling) for any species is described by transition to class  $R$ . Similar to many epidemiological models, this description of the dynamic system extends the basic SIR model (Weiss, 2013). Mathematically, the model is described by equations (1–7), a system of stochastic differential equations which we interpret in the Itô sense (Øksendal, 1998). Nonlinearity arises through product terms in the infection and consumption specifications. We use Gaussian noise, though multiplicative gamma noise can be employed to enforce non-negativity of rates (Bhadra et al., 2011). A version of SIRJPF2 with gamma noise, called SIRJPF2-Gamma, is presented in Section S9, though we found that a simpler Gaussian noise model is sufficient.

The nonlinearity and intrinsic stochasticity of the SIRPF2 model are characteristic features of dynamics in biological systems (Coulson et al., 2004; Wilkinson, 2018). In the following equations, the superscript  $k$  describes the species of *Daphnia*, taking value  $n$  for native (*D. dentifera*) and  $i$  for invasive (*D. lumholtzi*). The subscript  $u$  ranges in  $1 : U$ , with  $U = 8, 9, 10$  depending on the experiment. Most parameters are written with potential dependence on  $u$  to enable investigation of factors explaining differences between replicates. We will later establish that this dependence is not needed; random dynamic noise is sufficient to explain the observed differences between units. The experimentally determined sampling rate,  $\delta = 0.013$  per liter per day, and food replenishment rate,  $\mu = 0.37 \times 10^5$  cells per liter per day, are fixed parameter values that are shared between units based on the experimental protocol.

$$\begin{aligned}
 (1) \quad dS_u^k(t) &= \lambda_{J,u}^k J_u^k(t) dt - \{ \theta_{S,u}^k + p_u^k f_{S,u}^k P_u(t) + \delta \} S_u^k(t) dt + S_u^k(t) d\zeta_{S,u}^k, \\
 (2) \quad dI_u^k(t) &= p_u^k f_{S,u}^k S_u^k(t) P_u(t) dt - \{ \theta_{I,u}^k + \delta \} I_u^k(t) dt + I_u^k(t) d\zeta_{I,u}^k, \\
 (3) \quad dJ_u^k(t) &= r_u^k f_{S,u}^k F_u(t) S_u^k(t) dt - \{ \theta_{J,u}^k + \delta + \lambda_{J,u}^k \} J_u^k(t) dt + J_u^k(t) d\zeta_{J,u}^k, \\
 (4) \quad dP_u(t) &= \sum_{k \in \{n,i\}} \left( \beta_u^k \theta_{I,u}^k I_u^k(t) - f_{S,u}^k \{ S_u^k(t) + \xi_u I_u^k(t) \} P_u(t) \right) dt \\
 &\quad - \theta_{P,u} P_u(t) dt + P_u(t) d\zeta_{P,u},
 \end{aligned}$$

$$(5) \quad dF_u(t) = - \sum_{\mathbb{k} \in \{\mathbb{n}, \mathbb{i}\}} f_{S,u}^{\mathbb{k}} F_u(t) \left( S_u^{\mathbb{k}}(t) + \xi_{J,u} J_u^{\mathbb{k}}(t) + \xi_u I_u^{\mathbb{k}}(t) \right) dt + \mu dt + F_u(t) d\zeta_{F,u},$$

$$(6) \quad d\zeta_{S,u}^{\mathbb{k}} \sim N[0, (\sigma_{S,u}^{\mathbb{k}})^2 dt], \quad d\zeta_{I,u}^{\mathbb{k}} \sim N[0, (\sigma_{I,u}^{\mathbb{k}})^2 dt], \quad d\zeta_{J,u}^{\mathbb{k}} \sim N[0, (\sigma_{J,u}^{\mathbb{k}})^2 dt],$$

$$(7) \quad d\zeta_{F,u} \sim N[0, \sigma_{F,u}^2 dt], \quad d\zeta_{P,u} \sim N[0, \sigma_{P,u}^2 dt].$$

Here,  $N[\mu, \sigma^2]$  is the normal distribution with mean  $\mu$  and variance  $\sigma^2$ . Intraspecific competition among hosts is conceptualized via limitation of food and environment resources, with  $r_u^{\mathbb{k}}$  denoting the reproductive rate for species  $\mathbb{k} \in \{\mathbb{n}, \mathbb{i}\}$ . Host density growth, influenced by algal density, is posited to be linear, indicating a direct relationship between food availability and *Daphnia* reproduction. Mortality rates of susceptible hosts of species  $\mathbb{k}$  are denoted as  $\theta_{S,u}^{\mathbb{k}}$ . The model represents the dynamics of infection through *M. bicuspidata* ingestion by susceptible hosts. Specifically, susceptible hosts become infected at a rate depending on the density of parasite spores,  $P_u(t)$  and the host filtration rate,  $f_{S,u}^{\mathbb{k}}$ . This filtration rate is expected to be smaller for infected and juvenile individuals, modeled by rate factors  $\xi_u$  and  $\xi_{J,u}$  which are modeled as equal for both host species.

Infected hosts of species  $\mathbb{k}$  die at rate  $\theta_{I,u}^{\mathbb{k}}$ , at which point  $\beta_u^{\mathbb{k}}$  parasite spores are released. We do not measure the spores in this experiment, and so only the product  $\beta_u^{\mathbb{k}} p_u^{\mathbb{k}}$  is identified but not the separate values of  $\beta_u^{\mathbb{k}}$  and  $p_u^{\mathbb{k}}$ . We therefore fix  $\beta_u^{\mathbb{k}} = 3 \times 10^3$  for  $\mathbb{k} \in \{\mathbb{n}, \mathbb{i}\}$  and for all units (Penczykowski et al., 2014), though Searle et al. (2016a) estimated a slightly smaller value. Maturation and death of juvenile individuals are modeled to occur at rate  $\lambda_{J,u}^{\mathbb{k}}$  and  $\theta_{J,u}^{\mathbb{k}}$  respectively.

The model's effectiveness at statistically describing the variation in *Daphnia* population dynamics depends critically on its ability to account for overdispersion, a phenomenon which is overlooked in the common practice of using binomial and Poisson variation to add stochasticity to deterministic models (Bretó and Ionides, 2011). For the process model, this overdispersion is implemented by the noise parameters,  $\sigma_{S,u}^{\mathbb{k}}$ ,  $\sigma_{I,u}^{\mathbb{k}}$ ,  $\sigma_{J,u}^{\mathbb{k}}$ ,  $\sigma_{F,u}$ ,  $\sigma_{P,u}$ . For the measurement model, we obtain overdispersion by using a negative binomial model for the observed counts.

$$(8) \quad N_{S,u,n}^{\mathbb{k}} \sim \text{NBinomial}(S_u^{\mathbb{k}}(t_n), \tau_{S,u}^{\mathbb{k}}), \quad N_{I,u,n}^{\mathbb{k}} \sim \text{NBinomial}(I_u^{\mathbb{k}}(t_n), \tau_{I,u}^{\mathbb{k}})$$

where  $\text{NBinomial}(\mu, \tau)$  is the negative binomial distribution with mean  $\mu$  and variance  $\mu + \frac{\mu^2}{\tau}$ . Since the sampled volume is 1 liter, the densities  $S_u^{\mathbb{k}}(t_n)$  and  $I_u^{\mathbb{k}}(t_n)$  are on the same scale as the respective counts,  $N_{S,u,n}^{\mathbb{k}}$  and  $N_{I,u,n}^{\mathbb{k}}$ . We do not include the data on juveniles during the model fitting process, to keep our approach comparable to the investigation by Searle et al. (2016a). Therefore, these data are available for an out-of-sample model validation.

Simpler treatments with a single *Daphnia* species (SIRJPF) and without the parasite (SRJF and SRJF2) arise as special cases of SIRJPF2, with appropriate states and parameters set to zero. These experiment treatments and corresponding models are presented in Section S4, S5 and S6. By isolating the food variable in the absence of parasites and/or inter-species competition among *Daphnia*, these results provide additional evidence on how food resource availability or scarcity impacts the population dynamics of native and invasive species.

**4. Methodology for PanelPOMP models.** The applications of POMP models span various fields, including epidemiology (Mietchen et al., 2024; Fox et al., 2022; Wen

et al., 2024), ecology (Auger-Méthé et al., 2021; Marino et al., 2019; Blackwood et al., 2013), and finance (Bretó, 2014), among others. In epidemiological studies, POMP models are instrumental in estimating disease transmission rates and predicting future outbreaks by modeling the spread of infectious diseases when only a fraction of cases are observed (Subramanian et al., 2021). Similarly, POMP models help in understanding animal population dynamics and migration patterns based on limited or sporadic sighting data (Auger-Méthé et al., 2021). Widely used latent state estimation and parameter inference methods for nonlinear non-Gaussian POMP models build on the particle filter, also known as sequential Monte Carlo (Arulampalam et al., 2002). Particle filters involve simulating a large number of potential realizations of the latent process and updating these simulations as new observations become available, thereby approximating the conditional distribution of the state given the data. Various plug-and-play inference methods including iterated filtering, particle Markov chain Monte Carlo, and approximate Bayesian computing (ABC) are available using software such as pomp (King et al., 2016) or NIMBLE (de Valpine et al., 2017).

Algorithms and software for vector-valued POMP models could in principle be applied to PanelPOMP models, but in practice the high dimensionality of panel data requires modified algorithms. The panel iterated filter (PIF) algorithm of Bretó et al. (2020) enables plug-and-play likelihood-based inference for general PanelPOMP models. We start by setting up some general notation that we use to describe PIF, a variant of PIF called the marginalized panel iterated filter (MPIF), and our specific model for the mesocosm experiment.

Time series data are collected on  $U$  units, labeled  $1, 2, \dots, U$ . Unit  $u$  is observed at  $N_u$  measurement times, denoted by  $t_{u,1}, t_{u,2}, \dots, t_{u,N_u}$  which we write as  $t_{u,1:N_u}$ . For our example,  $N_u = N = 10$ , and  $t_{u,n} = t_n = 5n + 2$  days, though the algorithms and software used do not need regular observation times. The observation on unit  $u$  at time  $t_{u,n}$  is denoted by  $y_{u,n}^*$ , and in our example this corresponds to a vector of measured densities of *Daphnia* classified by species, life stage, and infection status. The data are modeled by a stochastic process  $\{Y_{u,n}, u \in 1:U, n \in 1:N_u\}$ . We suppose there is an unobserved continuous-time Markov process  $\{X_u(t), t_{u,0} \leq t \leq t_{N_u}\}$  such that  $Y_{u,n}$  is a noisy measurement of  $X_u(t_{u,n})$ . In our example,

$$(9) \quad X_u(t) = (S_u^n(t), I_u^n(t), J_u^n(t), S_u^i(t), I_u^i(t), J_u^i(t), P_u(t), F_u(t)),$$

$$(10) \quad Y_{u,n} = (N_{S,u,n}^n, N_{I,u,n}^n, N_{S,u,n}^i, N_{I,u,n}^i).$$

The time at which the latent process for unit  $u$  is initialized is  $t_{u,0}$ , and in our specific example  $t_{u,0} = t_0 = 0$ . We write  $X_{u,n} = X_u(t_{u,n})$ , and we note that the collection of these discrete-time latent state values is sufficient to describe the model for the data. Nevertheless, in many situations (including this case study) it is scientifically helpful to construct the latent process in continuous time.

The PanelPOMP model depends on a parameter vector,  $\theta$ , which we write as  $\theta = (\phi, \psi_{1:U})$  where  $\phi$  is the vector of parameters shared between all units and  $\psi_u$  is a vector of parameters specific to unit  $u$ . We can then define a general PanelPOMP model in terms of the transition density  $f_{X_{u,n}|X_{u,n-1}}(x_{u,n}|x_{u,n-1}; \phi, \psi_u)$ , the initial density  $f_{X_{u,0}}(x_{u,0}; \phi, \psi_u)$ , and the measurement density  $f_{Y_{u,n}|X_{u,n}}(y_{u,n}|x_{u,n}; \phi, \psi_u)$  of each unit  $u$  in the panel. For a model such as ours, the transition density is defined implicitly by the model equations. The plug-and-play property means that we never require explicit evaluation of this density, only the ability to obtain simulated trajectories satisfying equations (1–7). We do that using Euler’s method (Kloeden and Platen, 1999).

A defining feature of a PanelPOMP is that the latent processes are modeled as independent across units, so the data generating models for each unit are linked only via the shared parameter vector,  $\phi$ . The likelihood function can therefore be written as

$$(11) \quad \mathcal{L}(\theta) = \prod_{u=1}^U \int f_{Y_{u,1:N}|X_{u,1:N}}(y_{u,1:N}^* | x_{u,1:N}; \phi, \psi_u) f_{X_{u,1:N}}(x_{u,1:N}; \phi, \psi_u) dx_{u,1:N}.$$

Here, we are supposing the existence of conditional probability densities; these may be interpreted as probability mass functions if the latent process and/or the measurement model are discrete-valued. The likelihood function cannot generally be evaluated explicitly. However, techniques that let us approximately evaluate and maximize this likelihood, while making proper accommodation for the approximation error, let us employ likelihood-based inference tools for parameter estimation, confidence intervals, diagnosis of model misspecification, and model selection. Our subsequent investigation demonstrates all these aspects of data analysis.

Particle filters use a Monte Carlo swarm of latent state values, known as particles, to represent the conditional distribution of the latent variables given the data. Particles are updated according to the dynamics of the latent process and then resampled according to their likelihood given the data. This resampling is key to the particle filter’s efficacy, since it allows the algorithm to focus its computational effort on the most probable states of the latent process, as indicated by the observed data. Iterated particle filters maximize the likelihood for a POMP model by associating a parameter value with each particle. An analogy with Darwinian evolution explains heuristically how this happens. The parameter values are the heritable DNA belonging to each particle. The parameter perturbations are analogous to mutations, and the resampling is analogous to natural selection which preferentially favors the fittest particles, i.e., those with parameter values most consistent with the data. [Kitagawa \(1998\)](#) noticed that this favorable property of a particle filter, with perturbed parameters, can be empirically shown to move the population of these parameter values toward a region of increasing likelihood. [Ionides et al. \(2006\)](#) and [Ionides et al. \(2015\)](#) found ways to iterate this process that lead to concentration around the MLE. The panel extension of iterated filtering by [Bretó et al. \(2020\)](#) applies the same principle to PanelPOMP models. A flow diagram for a panel iterated filtering (PIF) algorithm is shown in [Fig. 3](#), and full details are given in [Section S2](#). At each iteration,  $m = 1:M$ , the magnitude of the perturbations is reduced, and we use geometric “cooling” with a reduction factor  $\rho < 1$ . By providing likelihood-based inference for a general class of nonlinear, non-Gaussian PanelPOMP models, PIF has the potential to facilitate discoveries in a wide range of scientific domains. This article provides a case study for such an application in the field of ecology. Parameter estimates and corresponding confidence intervals were obtained using PIF. A sufficiently large number of particles is needed to maintain particle diversity and achieve robust exploration of the parameter space. However, computational cost increases with the number of particles. We found that for this model and data,  $J = 500$  particles and cooling parameter  $\rho = 0.7^{1/50} = 0.993$  gave sufficient Monte Carlo accuracy and computational feasibility. The likelihood maximization via the PIF algorithm was conducted in three successive stages of  $M = 150$ ,  $M = 150$ , and  $M = 250$  iterations, respectively. This tempering approach has commonly been found useful for iterated filtering applications ([Bhadra et al., 2011](#)). After completing each stage, we selected the top 25% of parameter values ranked by their corresponding log-likelihoods and used these as the initial parameter swarm for the subsequent stage. Numerical experimentation under this setting provided log-likelihood estimates with standard error

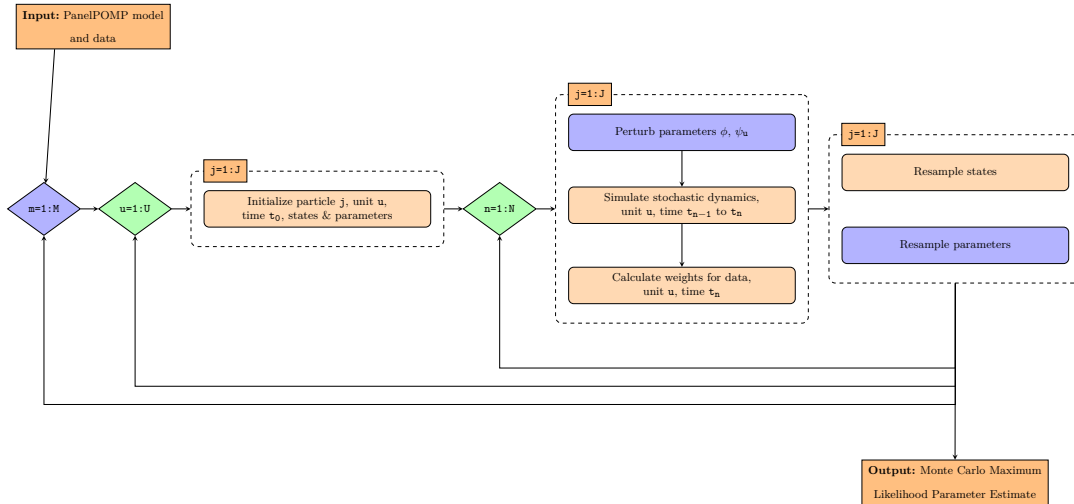


Figure 3: A flow diagram for panel iterated filtering. If the blue shaded blocks are omitted, this becomes a standard particle filter for each unit. Full pseudocode is in Section S2. In the PIF algorithm of Bretó et al. (2020), the parameter resampling step involves re-selecting the unit-specific parameters corresponding to particle  $j$  for units  $\tilde{u} \neq u$  as well as  $\psi_u$ . For the marginalized PIF algorithm, only  $\phi$  and  $\psi_u$  are reselected in this step.

0.35 and adequately supported parameter inference while keeping computations within the capabilities of one 36-core Linux machine.

In PanelPOMP models, it is necessary to decide which parameters should be shared across units, versus which should be unit-specific. This decision can be guided both by scientific considerations and by empirical evidence. As the number of units in the PanelPOMP model increases, the number of parameters increases linearly with the number of unit-specific parameters. Therefore, balancing model fit and parsimony becomes essential to avoid overly complex models. Here, we use AIC, defined to be twice the number of estimated parameters minus twice the maximized log-likelihood. AIC compensates for over-fitting by seeking to optimize out-of-sample forecast skill for a log-likelihood scoring rule (Akaike, 1974). Thus, AIC penalizes complexity only so far as it affects generalization error for prediction; parsimony has additional scientific benefits, but the scientific value of model simplicity cannot be determined by a purely statistical criterion. Various other information criteria have been proposed and used in ecological data analysis, in addition to likelihood ratio hypothesis tests (Johnson and Omland, 2004). However, AIC has earned widespread use due to its clear theoretical motivation and its applicability for evaluating collections of non-nested hypotheses, including comparing mechanistic models to non-mechanistic benchmark models.

Profile likelihood provides a useful approach for constructing confidence intervals in nonlinear stochastic systems (Simpson and Maclaren, 2023). Calculation of the profiles can validate that sufficient computational effort was conducted in order for the Monte Carlo optimization to converge: by increasing Monte Carlo replications, consistent convergence to the same maximum of the profile likelihood curve confirms adequate numerical effort for optimization. To obtain confidence intervals, we correct for the Monte Carlo uncertainty in the profile likelihood evaluations by using the Monte Carlo Adjusted Profile (MCAP) method of Ionides et al. (2017).

Parameter	Definition	Unit	Value	CI
$S^k$	Susceptible host density for species $k$	individual $\cdot L^{-1}$	Variable	
$I^k$	Infected host density for species $k$	individual $\cdot L^{-1}$	Variable	
$J^k$	Juvenile host density for species $k$	individual $\cdot L^{-1}$	Variable	
$F$	Alga density	$10^6 \cdot \text{cell} \cdot L^{-1}$	Variable	
$P$	Spore density	$10^3 \cdot \text{spore} \cdot L^{-1}$	Variable	
$r_i^n$	Birth rate of native juvenile	individual $\cdot 10^{-6} \cdot \text{cell}^{-1}$	4.08 $\cdot 10$	(2.64 $\cdot 10$ , 1.94 $\cdot 10^2$ )
$r_i^i$	Birth rate of invasive juvenile	individual $\cdot 10^{-6} \cdot \text{cell}^{-1}$	2.15 $\cdot 10^5$	(1.74 $\cdot 10^2$ , $\infty$ )
$f_{S_i}^n$	Native susceptible adult host filtering rate	L $\cdot$ individual $^{-1} \cdot$ day $^{-1}$	1.10 $\cdot 10^{-3}$	(4.21 $\cdot 10^{-4}$ , 1.64 $\cdot 10^{-3}$ )
$f_{S_i}^i$	Invasive susceptible adult host filtering rate	L $\cdot$ individual $^{-1} \cdot$ day $^{-1}$	2.42 $\cdot 10^{-7}$	(0.947 $\cdot 10^{-4}$ )
$p_i^n$	Number of native infections per spore	$10^{-3} \cdot$ individual $\cdot$ spore $^{-1}$	2.72 $\cdot 10^{-1}$	(1.54 $\cdot 10^{-1}$ , 7.28 $\cdot 10^{-1}$ )
$p_i^i$	Number of invasive infections per spore	$10^{-3} \cdot$ individual $\cdot$ spore $^{-1}$	1.34 $\cdot 10^3$	(1.13, $\infty$ )
$p_i^n \cdot f_{S_i}^n$	Effective infection rate of native adult hosts per-spore	$10^{-3} \cdot$ L $\cdot$ spore $^{-1} \cdot$ day $^{-1}$	2.99 $\cdot 10^{-4}$	(2.13 $\cdot 10^{-4}$ , 4.61 $\cdot 10^{-4}$ )
$p_i^i \cdot f_{S_i}^i$	Effective infection rate of invasive adult hosts per-spore	$10^{-3} \cdot$ L $\cdot$ spore $^{-1} \cdot$ day $^{-1}$	3.25 $\cdot 10^{-4}$	(2.44 $\cdot 10^{-4}$ , 4.87 $\cdot 10^{-4}$ )
$r_i^n \cdot f_{S_i}^n$	Effective birth rate of native juvenile	$10^{-6} \cdot$ L $\cdot$ cell $^{-1} \cdot$ day $^{-1}$	4.49 $\cdot 10^{-2}$	(3.83 $\cdot 10^{-2}$ , 8.53 $\cdot 10^{-2}$ )
$r_i^i \cdot f_{S_i}^i$	Effective birth rate of invasive juvenile	$10^{-6} \cdot$ L $\cdot$ cell $^{-1} \cdot$ day $^{-1}$	5.21 $\cdot 10^{-2}$	(4.47 $\cdot 10^{-2}$ , 6.91 $\cdot 10^{-2}$ )
$\theta_S^n$	Native susceptible adult host mortality rate	day $^{-1}$	8.48 $\cdot 10^{-4}$	(0.168 $\cdot 10^{-1}$ )
$\theta_S^i$	Invasive susceptible adult host mortality rate	day $^{-1}$	2.52 $\cdot 10^{-3}$	(0.583 $\cdot 10^{-3}$ )
$\theta_I^n$	Native infected adult host mortality rate	day $^{-1}$	5.84 $\cdot 10^{-1}$	(2.75 $\cdot 10^{-1}$ , 8.84 $\cdot 10^{-1}$ )
$\theta_I^i$	Invasive infected adult host mortality rate	day $^{-1}$	3.85 $\cdot 10^{-1}$	(2.21 $\cdot 10^{-1}$ , 5.92 $\cdot 10^{-1}$ )
$\theta_J^n$	Native juvenile mortality rate	day $^{-1}$	1.87 $\cdot 10^{-5}$	(0.166 $\cdot 10^{-2}$ )
$\theta_J^i$	Invasive juvenile mortality rate	day $^{-1}$	5.62 $\cdot 10^{-4}$	(0.179 $\cdot 10^{-2}$ )
$\theta_P$	Spore degradation rate	day $^{-1}$	9.48 $\cdot 10^{-4}$	(0.705 $\cdot 10^{-2}$ )
$\lambda_J^n$	Maturation rate of native juvenile	day $^{-1}$	1.00 $\cdot 10^{-1}$	
$\lambda_J^i$	Maturation rate of invasive juvenile	day $^{-1}$	1.00 $\cdot 10^{-1}$	
$\xi$	Ratio of infected host filtering rate	Unitless	2.22 $\cdot 10$	(1.01 $\cdot 10$ , 7.10 $\cdot 10$ )
$\xi_J$	Ratio of juvenile individual filtering rate	Unitless	1.00	
$\beta_i^n$	Spores produced per infected native individual	$10^3 \cdot$ spores $\cdot$ individual $^{-1} \cdot$ day $^{-1}$	3.00 $\cdot 10$	
$\beta_i^i$	Spores produced per infected invasive individual	$10^3 \cdot$ spores $\cdot$ individual $^{-1} \cdot$ day $^{-1}$	3.00 $\cdot 10$	
$\mu$	Alga refilling rate	$10^6 \cdot$ cell $\cdot L^{-1} \cdot$ day $^{-1}$	3.70 $\cdot 10^{-1}$	
$\delta$	Sampling rate	day $^{-1}$	1.30 $\cdot 10^{-2}$	
$\sigma_S^n$	Standard deviation of Brownian motion of susceptible native adult	$\sqrt{\text{individual} \cdot \text{day}^{-1}}$	0	
$\sigma_S^i$	Standard deviation of Brownian motion of susceptible invasive adult	$\sqrt{\text{individual} \cdot \text{day}^{-1}}$	0	
$\sigma_I^n$	Standard deviation of Brownian motion of infected native adult	$\sqrt{\text{individual} \cdot \text{day}^{-1}}$	2.93 $\cdot 10^{-4}$	(0, $\infty$ )
$\sigma_I^i$	Standard deviation of Brownian motion of infected invasive adult	$\sqrt{\text{individual} \cdot \text{day}^{-1}}$	1.73 $\cdot 10^{-7}$	(0, $\infty$ )
$\sigma_J^n$	Standard deviation of Brownian motion of native juvenile	$\sqrt{\text{individual} \cdot \text{day}^{-1}}$	2.84 $\cdot 10^{-1}$	(1.50 $\cdot 10^{-1}$ , 4.80 $\cdot 10^{-1}$ )
$\sigma_J^i$	Standard deviation of Brownian motion of invasive juvenile	$\sqrt{\text{individual} \cdot \text{day}^{-1}}$	3.02 $\cdot 10^{-1}$	(1.35 $\cdot 10^{-1}$ , 4.20 $\cdot 10^{-1}$ )
$\sigma_F$	Standard deviation of Brownian motion of alga	$\sqrt{\text{individual} \cdot \text{day}^{-1}}$	1.44 $\cdot 10^{-1}$	(5.41 $\cdot 10^{-2}$ , 2.33 $\cdot 10^{-1}$ )
$\sigma_P$	Standard deviation of Brownian motion of parasite	$\sqrt{\text{individual} \cdot \text{day}^{-1}}$	2.71 $\cdot 10^{-1}$	(4.41 $\cdot 10^{-2}$ , 4.20 $\cdot 10^{-1}$ )
$\tau_S^n$	Measurement dispersion for susceptible native adult	Unitless	4.10	(2.92, 7.65)
$\tau_S^i$	Measurement dispersion for susceptible invasive adult	Unitless	5.26	(2.66, 8.60)
$\tau_I^n$	Measurement dispersion for infected native adult	Unitless	9.02 $\cdot 10^{-1}$	(6.69 $\cdot 10^{-1}$ , 1.70)
$\tau_I^i$	Measurement dispersion for infected invasive adult	Unitless	1.39	(6.77 $\cdot 10^{-1}$ , 2.27)

TABLE 1

Variables and parameter definitions and estimates. The superscript  $n$  and  $i$  represents native (*D. dentifera*) and invasive (*D. lumholtzi*) respectively.

Model	Parameters	Max log-likelihood	AIC
SIRJPF2	26	-880.56	1809.12
SIRPF2	20	-891.80	1823.60
SIRJPF2-Gamma	26	-892.03	1836.06
Negative Binomial (Cubic)	24	-932.61	1913.21
Negative Binomial (Quadratic)	20	-943.79	1927.59
Negative Binomial (Linear)	16	-1009.25	2050.50
Searle et al. (2016)	19	-2483.49	5004.97

TABLE 2

Comparison of model fit for mechanistic models and non-mechanistic benchmarks.

**5. Results.** Parameter estimates and confidence intervals for the SIRJPF2 model are in Table 1, which also includes a brief description of the parameter mechanism and the corresponding measurement units. The log-likelihood, AIC and number of parameters for this model is compared against alternatives in Table 2. Differences in AIC values between models are interpretable only between alternative models for the same data, so we fit statistical benchmark models to calibrate these results. Specifically, we fit negative binomial regression models with linear, quadratic, and cubic dependence on time. Comparing a mechanistic model to a simple statistical benchmark model is a useful test of model specification, since it reveals quickly whether the mechanistic model is constructed in a way that is statistically compatible with the data (Wheeler et al., 2024; Li et al., 2024). The negative binomial models with polynomial (quadratic

or cubic) terms fit the dynamics better than a linear trend. Although these more flexible polynomial expansions do not represent latent biological processes directly, they better capture the non-linearity of the system, leading to a higher log-likelihood and a lower AIC than the linear model.

Table 2 includes the log-likelihood and AIC values for the fitted model and the deterministic ODE model of Searle et al. (2016a), which served as a starting point for our model. Although the mechanistic model of Searle et al. (2016a) was successfully used to provide scientific insight into the ecological dynamics of the experiment, we found it does not provide a competitive statistical fit to the data compared to simple alternatives. We also consider several variants of the SIRJPF2 model described here, namely a model that excludes the juvenile state (SIRPF2, described in Section S7.3), and a version of the SIRJPF2 model where the Gaussian noise has been replaced with Gamma noise (SIRJPF2-Gamma, see Section S9). The SIRPF2 model ignores the process of maturation of juvenile and instead highlights the infection process of the *Daphnia* adult populations. While this model does have an improved statistical fit compared to the deterministic ODE model of Searle et al. (2016a) and other benchmark models, it has weaker performance than the model that accounts for the effect of juveniles on population dynamics.

Our SIRJPF2 model quantitatively describes the observed data better than any of the alternatives, while simultaneously providing a plausible mechanistic description of the experimental system. Having assessed the fit of our model, we have some support for interpreting the estimated parameters of the model. It is hard to rule out the possibility that some biological details in the model are misspecified in ways that may lead to confounding with the modeled mechanisms and therefore bias in the causal interpretation of estimated parameters (Wheeler et al., 2024). With that caveat acknowledged, we proceed to interpret the model and parameters at face value.

There are some substantial differences in the observed trajectories between experimental replicates (Figs. 1, S5, S6, S14, S19, S20). This could result from small differences in the initialization of the replicates, or it could simply be a consequence of the natural variation in any biological system. In the context of our model, these hypotheses correspond to unit-specific parameters and stochastic dynamics, respectively. Here, the evidence suggests that stochastic dynamics are sufficient to explain the variation in the data. Table S1 in the supplement presents a comparison of model fit and complexity across various configurations of unit-specific parameters within the PanelPOMP framework for the SIRJPF2 model. These distinct configurations represent varying hypotheses explaining the differences across replicates. The models assessed include those with unit-specific parameters for combinations of mortality rates ( $\theta_S^n, \theta_S^i, \theta_I^n, \theta_I^i$ ), filter rates ( $f_S^n, f_S^i$ ), growth efficiency factors ( $r^n, r^i$ ), and spore infection rates ( $p^n, p^i$ ).

When parameters are shared between units, we obtain a substantial reduction in the number of parameters to be estimated. For field data, one may expect the need for unit-specific parameters describing structural differences between units. For example, in a PanelPOMP analysis of pre-vaccine polio transmission dynamics for state-level data in USA, Bretó et al. (2020) used unit-specific parameters to describe variations in the seasonality of transmission across units. Modeling decisions concerning shared and unit-specific parameters are testable hypotheses that are of scientific interest in the quest for a parsimonious understanding of the dynamics, and for understanding sources of variation.

In our selected model, the feeding rate of the invasive *Daphnia* species ( $f_S^i$ ) is not statistically significantly different from that of the native *Daphnia* species ( $f_S^n$ ). Also, the invasive species has a lower rate of infection per consumed spore ( $p^n < p^i$ ), consistent with Fig. 1C of Searle et al. (2016a). The native parameters—juvenile growth  $r^n$ , feeding rate  $f_S^n$  and infection probability  $p^n$  show full identifiability as shown in Fig. S1, whereas the corresponding invasive parameters  $r^i$ ,  $f_S^i$ ,  $p^i$  are only bounded on one side. This limited identifiability is attributable to (i) the lower density of *D. lumholtzi* which may increase stochastic effects; (ii) model complexity, since the parameters become identifiable in host-only experiments with simplified dynamics (Fig. S8); (iii) the product of these parameters governs the juvenile growth and infection process, which allows individual invasive coefficients to vary so long as their products remain fixed. To investigate this last explanation, Table 1 shows relatively precise identifiability of the effective infection rate ( $p^n f_S^n$ ,  $p^i f_S^i$ ), and the effective juvenile birth rate ( $r^n f_S^n$ ,  $r^i f_S^i$ ). At the same time, the identifiability of the measurement error parameter ( $\tau_S^n$ ,  $\tau_S^i$ ,  $\tau_I^n$ ,  $\tau_I^i$ ) explains variability that would otherwise inflate weakly identified process-noise terms. Our results indicate that stochasticity in the juvenile stage ( $\sigma_J^n$  and  $\sigma_J^i$ ), food ( $\sigma_F$ ) and parasite ( $\sigma_P$ ) dynamics are important model features, needed to adequately describe the data (see the corresponding confidence intervals in Table 1 and the profile likelihood plots in Fig. S1). That suggests the juvenile stage, food and parasite dynamics are particularly sensitive to unmodeled phenomena, such as lower fitness resulting from resource depletion, since unmodeled phenomena can be explained by the model as stochastic uncertainty. Noise in other stages is small, and often insignificantly different from zero, as shown by profile likelihood plots for  $\sigma_I^n$ ,  $\sigma_I^i$ . Since the profile for  $\sigma_S^n$ ,  $\sigma_S^i$  is flat and insignificantly different from zero, we fixed the  $\sigma_S^n$  and  $\sigma_S^i$  to be zero as constants. The proposed model is flexible enough to describe the predictability and uncertainty inherent in the system in various ways, and the data have come down in favor of  $\sigma_J^n$  and  $\sigma_J^i$ . This indicates that the model might be recalibrated after fixing some or all of the other noise intensities to zero, but there is no pressing need to do so. The noise term is helpful when describing other dynamics, for example, dynamics with two species of *Daphnia* and no parasite as shown in Supplement S5. Natural death rates, ( $\theta_S^n$ ,  $\theta_S^i$ ,  $\theta_J^n$  and  $\theta_J^i$ ), are poorly identified in the presence of the parasite since most death comes through the infection process. By contrast, these parameters are better identified in the treatments without the parasite (Fig. S15).

Alternative SIRJPF2 model specifications were fitted with various parameters taken to be unit-specific. The log-likelihood estimates and corresponding AIC values for these models are presented in Supplementary Table S1. These results suggest that introducing unit-specific parameters increased model complexity without significantly improving model performance. Similar results hold for the other treatments (Tables S4, S5, S7, S10, S11). Consequently, we choose to select the model with all shared parameters, as it offers a more parsimonious and robust representation of the system dynamics without substantially compromising the explanatory power. And the confidence intervals are obtained by the MCAP are shown in Fig. S1 for the SIRJPF2 model, and Figs. S7, S8, S15, S21, S22 for other treatments.

Our estimated value of the filtration rate factor for infected *Daphnia* is  $\xi > 1$ , which is superficially incompatible with previous empirical observations that infection leads to reduced filtration, especially as disease progresses (Searle et al., 2016a; Penczykowski et al., 2014). One way to reconcile these results is the possibility that *Daphnia* which filter faster are exposed to more spores and are consequently more likely to become infected. There are six genotypes in each experimental unit, and we only observe

population-level dynamics rather than individual-level data. When modeling these dynamics, our parameters of interest actually represent population-level means. In other words, given this specific combination of genotypes,  $\xi > 1$  indicates that at the population level, the average filtration rate of infected *Daphnia* is higher than that of susceptible *Daphnia*. However, the direction can be reversed at the individual level.  $\xi > 1$  could also be a result of reverse causation due to unmodeled heterogeneity: *Daphnia* with higher filtering rate are exposed to more pathogen spores and so filtering rate can causally affect infection status in addition to any causal effect of infection status on filtering rate. A subsequent investigation extending the model and measurements to include *Daphnia* size and genotype could help to clarify this.

As shown in Table 1, the estimated values of  $p^n$  and  $p^i$  differ significantly, suggesting that the parasite influences the proportion of individuals becoming infected within each species in distinct scale. This difference, in turn, affects the interactions between the two species.

**6. Model diagnostics.** Our best-fitting model includes explicit modeling of juveniles and their maturation, a feature not present in the model of Searle et al. (2016a). The extra delay in the adult dynamics introduced by the juvenile stage appears to be necessary for the model to provide a strong statistical fit to the data. In the original study, Searle et al. (2016a) collected juvenile data but did not include it as part of the data analysis. To make our analysis comparable to Searle et al. (2016a), we fitted to the same data, and therefore did not include a measurement model for juveniles. Instead, the juvenile data are available to use as model validation, which is done by comparing simulations from the fitted model to the unmodeled data in Fig. 4.

Knowing that the average body length of juvenile *Daphnia* is shorter than the average body size of adult *Daphnia* for both species, we expect that juvenile individuals will have a lower filtering rate. Yet, without fitting to juvenile data, it is unclear whether that question is identifiable from the model at hand. To explore this assumption, we compared the model using the value  $\xi_{J,u} = 1$  (juveniles and adults filter at the same rate) against the extreme value  $\xi_{J,u} = 0.001$  (juveniles filter much slower than adults).

Likelihood residual analysis is an effective tool for comparing alternative models by assessing differences in the conditional log-likelihood for each data point (Wheeler et al., 2024; Li et al., 2024). We maximized the log-likelihood subject to the fixed values of  $\xi_{J,u}$ , and we compared the conditional log-likelihood for each observation category at each time point for each fixed  $\xi_{J,u}$  value, as shown in Fig. S3(a). These results show that setting  $\xi_{J,u} = 1$  leads to only a slightly higher maximized log-likelihood ( $-881.13$ ) than  $\xi_{J,u} = 0.001$  ( $-882.24$ ), indicating that the parameter is weakly identified. Moreover, plotting the conditional log-likelihood for each state at each time, we observe that the log-likelihood components for susceptible and infected *D. dentifera* and *D. lumholtzi* are similar between the two parameter settings. These results suggest that both parameter values provide indistinguishable explanations of the observed data, which is a stronger conclusion than obtaining a similar total log-likelihood.

The choice of  $\xi_{J,u}$  nevertheless affects the interpretation of the fitted model: the maximum likelihood estimates constrained to different values of  $\xi_{J,u}$  lead to distinct behavior of juvenile *Daphnia* in simulations, as shown in Fig. S3(b). If  $\xi_{J,u}$  is set to be lower than is biologically plausible, simulated trajectories from the fitted model have far fewer juveniles than was observed, which is shown in Fig. S3(b) as well. This

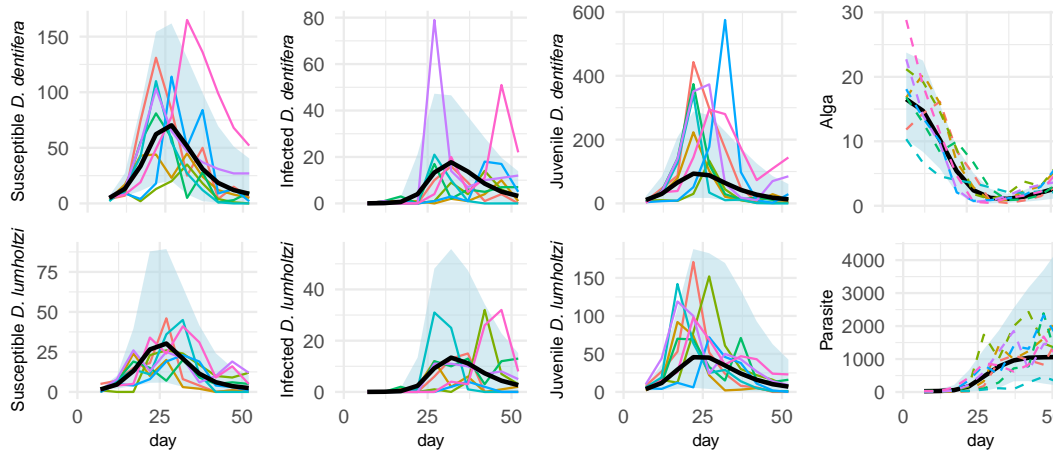


Figure 4: Simulated *Daphnia* densities (Individuals/Liter), alga density ( $10^6$ ·cells/Liter) and parasite density ( $10^3$ ·spores/Liter), over time (days), in the mixed species parasitized dynamics with parameters from Table 2. Blue bands are 95% pointwise confidence intervals; black lines show the mean; solid color lines are the experimental data. The parasite and food densities were not observed, and the dashed colored lines are simulations.

indicates model misspecification which could bias the estimates of  $r_u^k$ ,  $p_u^k$  and  $f_{S,u}^k$ . When a weakly identified parameter has consequences for the biological interpretation of the model, it is appropriate to fix that parameter at a biologically plausible value, and we therefore fixed  $\xi_{J,u} = 1$  for subsequent analysis, reasoning that larger juveniles should have filtering rates of the same order of magnitude as adults.

In addition to comparing model fit using AIC, we further validate the calibrated SIRJPF2 model via a simulation study. We generate 1000 simulations from the model and compare these simulations to the data (Fig. 4). Unit level simulation plots are shown in Fig. S2. The simulation bands presented in Fig. 4 closely align with the empirical data curves from the actual experiment, underscoring the model’s capability to replicate the observed dynamics. Notably, the model captures not only the trends but also the variability of the experimental data. Simulations from other models considered are included in the supplementary Figs. S9–S12, S16, S17, S23–S26.

The consistency of simulated data with the empirical data provides a qualitative indicator of the fit of the model that complements a comparison of log-likelihood values. Our simulations from the fitted SIRJPF2 model align with the observed data for the susceptible and infected species of *Daphnia*, and from the out-of fit juvenile data. This encourages us to explore the dynamics of latent states in the dynamic system by simulating the unobserved densities of *M. bicuspidata* and *A. falcatus* (see the third column of Fig. 4). The model predicts that the alga density decreases dramatically from its starting value, and then starts to rebound, due to resupply, once the *Daphnia* population has crashed. The parasite density is predicted to rise sharply during the population growth phase and then plateaus at a high level.

As mentioned in Section 3, high population density of *Daphnia* and limited food resources will prompt the production of male *Daphnia* leading to sexual reproduction

and the generation of diapausing eggs. That biological feature plays a large role in the model of Searle et al. (2016a), though Table 2 shows that our proposed mechanism for the role of resource depletion provides a better statistical explanation of the analyzed data. Our modeling and inference approach could be extended to include data on males and ephippial females, collected by Searle et al. (2016a) but not included in either their model or ours.

**7. Discussion.** Data science is having a transformative impact on ecological science (Mouquet et al., 2015) and yet too few studies undertake the task of reconciling ecological theory with time series data (Peacor et al., 2022). We have shown that our panel data methodology can allow full-information plug-and-play likelihood-based inference for the complex partially observed nonlinear stochastic dynamic models required to carry out this reconciliation. If the previous shortage of such data analyses is evidence for the previous lack of effective methods, PanelPOMP methodology should facilitate many previously intractable scientific investigations. Further advances in methodology, such as automatic differentiation for particle filters (Tan et al., 2024), will continue to increase the numerical tractability of PanelPOMP models.

Despite the complexity of our SIRJPF2 model, there are various ways in which further biological mechanisms could be incorporated. Our general plug-and-play framework can be further extended, and information criteria can be used to assess the importance of the proposed extension. For example, we could have incorporated the phenomenon that *Daphnia* with advanced infection have reduced feeding rate, or we could have allowed for the possibility that recently infected *Daphnia* can continue to reproduce. Additional residual analysis can search for patterns suggesting model improvements, following a similar approach to the residual analysis carried out here.

Although we did not calibrate the model using juvenile data, the addition of latent  $J_u^k$  states leads to an improved statistical fit. This could arise because the juvenile stage provides a delay in the *Daphnia* population growth dynamics, and because juveniles compete with adults for available food resources. Juveniles rarely exhibit disease, perhaps because of the latency in the parasite’s development within its host, but they nevertheless play a significant role in the infection dynamics. While we have shown the relevance of age-structured dynamics for understanding the data on adults, future work could identify the role of juveniles in the system dynamics more precisely. Adding juvenile data to the measurement model in the fitting procedure would enable higher statistical resolution on these dynamics. Our general POMP modeling and inference framework allows for the possibility of additional measurements or alternative model hypotheses.

Fitting a mechanistic model can provide a coherent explanation of system dynamics, but one must remain cautious about the possibility of confounding factors: mechanisms that are omitted from the model and which provide alternative (and perhaps scientifically pertinent) explanations of the observed phenomena (Li et al., 2024; Wheeler et al., 2024). We have an external test of our data by showing that it fits the juvenile data even though it was not fit to them, but this validation does not guarantee that all conclusions from the fitted model are correct. For example, while our results explain the decrease of *Daphnia* density by resource depletion, data on the exact food level were not available. The mechanism, although logical within the model’s framework, remains an unvalidated assumption in the theory rather than being directly supported by data.

Subsequent experimentation could substantiate or refute this prediction, and our data analysis methodology can assist such investigations.

Deterministic dynamic models continue to be widely used for ecological theory, despite the fact that inclusion of stochasticity can both improve model fit and enable diagnosis of model shortcomings (King et al., 2015). We demonstrate that including stochasticity can also lead to more parsimonious understanding of the system, since, in our case study, the dynamic stochasticity obviates a need for mesocosm-specific parameters to explain the difference between the experimental replicates.

This study has focused on the task of fitting models to understand the biological dynamics. This is closely related to the task of forecasting: fitting by maximum likelihood is equivalent to assessing one-step forecasting skill, and the AIC criterion is designed to evaluate out-of-sample predictive skill (Akaike, 1974). Within this framework, the goals of model-based biological understanding and forecasting are therefore closely aligned.

We have analyzed data collected in a controlled environment, raising the question of the applicability of our methods to field ecosystems. The applicability of our approach is contingent on having panel data of sufficient length and quality to inform interactions between species. Such panel datasets are available, for example, on pathogen-host ecosystems (Martinez-Bakker et al., 2015) and fisheries (Naiman et al., 2002). The flexibility of the PanelPOMP model class provides many opportunities for developing new models that integrate ecological theory with panel time series data.

**Significance Statement.** Collections of partially observed nonlinear stochastic dynamic systems arise in diverse scientific applications and raise methodologically challenging statistical inference problems. We consider this task in the context of analyzing panel time series data collected from ecological studies. We develop and demonstrate data analysis techniques that apply to a flexible model class, permitting scientists to obtain statistically efficient inference on models describing scientifically plausible mechanisms. For our experiment, a model with stochastic dynamics provides simpler interpretation than previous deterministic models: variation between experimental replicates can be described by dynamic noise rather than systematic differences between replicates.

**Acknowledgments.** Aaron A. King is additionally affiliated with the Center for the Study of Complex Systems, University of Michigan. This work was supported by National Science Foundation grant DMS-1761603.

## SUPPLEMENTARY MATERIAL

### Data Availability Statement

Code and data for the analysis is available on Github at <https://github.com/Megumi-ybb/Daphnia-ms> and is archived on Zenodo at <https://doi.org/10.5281/zenodo.15595669>.

### Supplementary Material

The Supplementary Material includes detailed specification of the Panel Iterated Filtering algorithm and the benchmark models. It also provides profile likelihoods, residual analyses, and diagnostics for the SIRJPF2, SIRJPF, and SRJF2 models, together with models for the additional dynamics discussed in the manuscript.

## REFERENCES

- Akaike, H. (1974). A new look at the statistical model identification. *IEEE Transactions on Automatic Control*, 19(6):716–723.
- Arulampalam, M. S., Maskell, S., Gordon, N., and Clapp, T. (2002). A tutorial on particle filters for online nonlinear/non-Gaussian Bayesian tracking. *IEEE Transactions on signal processing*, 50(2):174–188.
- Auger-Méthé, M., Newman, K., Cole, D., Empacher, F., Gryba, R., King, A. A., Leos-Barajas, V., Mills Flemming, J., Nielsen, A., Petris, G., and Thomas, L. (2021). A guide to state–space modeling of ecological time series. *Ecological Monographs*, 91(4):e01470.
- Bhadra, A., Ionides, E. L., Laneri, K., Pascual, M., Bouma, M., and Dhiman, R. C. (2011). Malaria in northwest India: Data analysis via partially observed stochastic differential equation models driven by Lévy noise. *Journal of the American Statistical Association*, 106(494):440–451.
- Bjørnstad, O. N. and Grenfell, B. T. (2001). Noisy clockwork: time series analysis of population fluctuations in animals. *Science*, 293(5530):638–643.
- Blackwood, J. C., Streicker, D. G., Altizer, S., and Rohani, P. (2013). Resolving the roles of immunity, pathogenesis, and immigration for rabies persistence in vampire bats. *Proceedings of the National Academy of Sciences*, 110(23):20837–20842.
- Bretó, C. (2014). On idiosyncratic stochasticity of financial leverage effects. *Statistics & Probability Letters*, 91:20–26.
- Bretó, C., He, D., Ionides, E. L., and King, A. A. (2009). Time series analysis via mechanistic models. *The Annals of Applied Statistics*, pages 319–348.
- Bretó, C. and Ionides, E. L. (2011). Compound Markov counting processes and their applications to modeling infinitesimally over-dispersed systems. *Stochastic Processes and their Applications*, 121(11):2571–2591.
- Bretó, C., Ionides, E. L., and King, A. A. (2020). Panel data analysis via mechanistic models. *Journal of the American Statistical Association*, 115(531):1178–1188.
- Breto, C., Wheeler, J., King, A. A., and Ionides, E. L. (2025). panelPomp: Analysis of panel data via partially observed Markov processes in R. *R Journal*, to appear.
- Buckland, S. T., Newman, K. B., Thomas, L., and Koesters, N. B. (2004). State-space models for the dynamics of wild animal populations. *Ecological Modelling*, 171(1-2):157–175.
- Cornetti, L., Fields, P. D., Van Damme, K., and Ebert, D. (2019). A fossil-calibrated phylogenomic analysis of *Daphnia* and the Daphniidae. *Molecular Phylogenetics and Evolution*, 137:250–262.
- Coulson, T., Rohani, P., and Pascual, M. (2004). Skeletons, noise and population growth: The end of an old debate? *Trends in Ecology and Evolution*, 19:359–364.
- de Valpine, P., Turek, D., Paciorek, C. J., Anderson-Bergman, C., Lang, D. T., and Bodik, R. (2017). Programming with models: writing statistical algorithms for general model structures with NIMBLE. *Journal of Computational and Graphical Statistics*, 26(2):403–413.
- Ebert, D. (2005). *Ecology, Epidemiology, and Evolution of Parasitism in Daphnia*. National Library of Medicine.
- Fox, S. J., Lachmann, M., Tec, M., Pasco, R., Woody, S., Du, Z., Wang, X., Ingle, T. A., Javan, E., Dahan, M., Gaither, K., Escott, M. E., Adler, S. I., Johnston, S. C., Scott, J. G., and Meyers, L. A. (2022). Real-time pandemic surveillance using hospital admissions and mobility data. *Proceedings of the National Academy of Sciences*, 119(7):e2111870119.
- Häggström, H., Persson, S., Cvijovic, M., and Picchini, U. (2026). Simulation-based inference for stochastic nonlinear mixed-effects models with applications in systems biology. *Statistics and Computing*, 36:99.
- Havel, J. E. and Hebert, P. D. (1993). *Daphnia lumholtzi* in North America: Another exotic zooplankton. *Limnology and Oceanography*, 38(8):1823–1827.
- He, D., Ionides, E. L., and King, A. A. (2010). Plug-and-play inference for disease dynamics: Measles in large and small towns as a case study. *Journal of the Royal Society Interface*, 7:271–283.
- Ionides, E. L., Bretó, C., and King, A. A. (2006). Inference for nonlinear dynamical systems. *Proceedings of the National Academy of Sciences*, 103(49):18438–18443.
- Ionides, E. L., Breto, C., Park, J., Smith, R. A., and King, A. A. (2017). Monte Carlo profile confidence intervals for dynamic systems. *Journal of the Royal Society Interface*, 14:1–10.
- Ionides, E. L., Nguyen, D., Atchadé, Y., Stoev, S., and King, A. A. (2015). Inference for dynamic and latent variable models via iterated, perturbed Bayes maps. *Proceedings of the National Academy of Sciences*, 112(3):719–724.
- Johnson, J. B. and Omland, K. S. (2004). Model selection in ecology and evolution. *Trends in Ecology & Evolution*, 19(2):101–108.

- King, A. A., Domenech de Cellès, M., Magpantay, F. M. G., and Rohani, P. (2015). Avoidable errors in the modelling of outbreaks of emerging pathogens, with special reference to Ebola. *Proceedings of the Royal Society of London, Series B*, 282(1806):20150347.
- King, A. A., Nguyen, D., and Ionides, E. L. (2016). Statistical inference for partially observed Markov processes via the R package pomp. *Journal of Statistical Software*, 69(12).
- Kitagawa, G. (1998). A self-organising state-space model. *Journal of the American Statistical Association*, 93:1203–1215.
- Kloeden, P. E. and Platen, E. (1999). *Numerical Solution of Stochastic Differential Equations*. Springer, New York, 3rd edition.
- Kotov, A. A. and Taylor, D. J. (2011). Mesozoic fossils (> 145 Mya) suggest the antiquity of the subgenera of *Daphnia* and their coevolution with chaoborid predators. *BMC Evolutionary Biology*, 11:1–9.
- Kristensen, K., Nielsen, A., Berg, C. W., Skaug, H., and Bell, B. M. (2016). TMB: Automatic differentiation and Laplace approximation. *Journal of Statistical Software*, 70:1–21.
- Li, J., Ionides, E. L., King, A. A., Pascual, M., and Ning, N. (2024). Inference on spatiotemporal dynamics for coupled biological populations. *Journal of the Royal Society Interface*, 21(216):20240217.
- Marino, J. A., Peacor, S. D., Bunnell, D. B., Vanderploeg, H. A., Pothoven, S. A., Elgin, A. K., Bence, J. R., Jiao, J., and Ionides, E. L. (2019). Evaluating consumptive and nonconsumptive predator effects on prey density using field times series data. *Ecology*, 100(3):e02583.
- Martinez-Bakker, M., King, A. A., and Rohani, P. (2015). Unraveling the transmission ecology of polio. *PLoS Biology*, 13(6):e1002172.
- Mietchen, M. S., Clancey, E., McMichael, C., and Lofgren, E. T. (2024). Estimating sars-cov-2 transmission parameters between coinciding outbreaks in a university population and the surrounding community. *medRxiv*, pages 2024–01.
- Mouquet, N., Lagadeuc, Y., Devictor, V., Doyen, L., Duputié, A., Eveillard, D., Faure, D., Garnier, E., Gimenez, O., Huneman, P., Jabot, F., Jarne, P., Joly, D., Julliard, R., Kéfi, S., Kergoat, G. J., Lavorel, S., Le Gall, L., Meslin, L., Morand, S., Morin, X., Morlon, H., Pinay, G., Pradel, R., Schurr, F. M., Thuiller, W., and Loreau, M. (2015). Predictive ecology in a changing world. *Journal of Applied Ecology*, 52(5):1293–1310.
- Naiman, R. J., Bilby, R. E., Schindler, D. E., and Helfield, J. M. (2002). Pacific salmon, nutrients, and the dynamics of freshwater and riparian ecosystems. *Ecosystems*, 5(4):399–417.
- Ning, N., Ionides, E. L., and Ritov, Y. (2021). Scalable Monte Carlo inference and rescaled local asymptotic normality. *Bernoulli*, 27:2532–2555.
- Øksendal, B. (1998). *Stochastic Differential Equations*. Springer, New York, 5th edition.
- Peacor, S. D., Dorn, N. J., Smith, J. A., Peckham, N. E., Cherry, M. J., Sheriff, M. J., and Kimbro, D. L. (2022). A skewed literature: Few studies evaluate the contribution of predation-risk effects to natural field patterns. *Ecology Letters*, 25(9):2048–2061.
- Penczykowski, R. M., Lemanski, B. C., Sieg, R. D., Hall, S. R., Housley Ochs, J., Kubanek, J., and Duffy, M. A. (2014). Poor resource quality lowers transmission potential by changing foraging behaviour. *Functional Ecology*, 28(5):1245–1255.
- Radzikowski, J., Krupińska, K., and Ślusarczyk, M. (2018). Different thermal stimuli initiate hatching of daphnia diapausing eggs originating from lakes and temporary waters. *Limnology*, 19:81–88.
- Schnute, J. T. (1994). A general framework for developing sequential fisheries models. *Canadian Journal of Fisheries and Aquatic Sciences*, 51(8):1676–1688.
- Searle, C. L., Cortez, M. H., Hunsberger, K. K., Grippi, D. C., Oleksy, I. A., Shaw, C. L., de la Serna, S. B., Lash, C. L., Dhir, K. L., and Duffy, M. A. (2016a). Population density, not host competence, drives patterns of disease in an invaded community. *The American Naturalist*, 188(5):554–566.
- Searle, C. L., Shaw, C. L., Hunsberger, K. K., Prado, M., and Duffy, M. A. (2016b). Salinization decreases population densities of the freshwater crustacean, *daphnia dentifera*. *Hydrobiologia*, 770:165–172.
- Simpson, M. J. and Maclaren, O. J. (2023). Profile-wise analysis: a profile likelihood-based workflow for identifiability analysis, estimation, and prediction with mechanistic mathematical models. *PLoS Computational Biology*, 19(9):e1011515.
- Subramanian, R., He, Q., and Pascual, M. (2021). Quantifying asymptomatic infection and transmission of COVID-19 in New York City using observed cases, serology, and testing capacity. *Proceedings of the National Academy of Sciences*, 118(9).
- Tan, K., Hooker, G., and Ionides, E. L. (2024). Accelerated inference for partially observed Markov processes using automatic differentiation. *arXiv:2407.03085*.
- Weiss, H. H. (2013). The SIR model and the foundations of public health. *Materials Mathematics*, pages 1–17.

- Wen, L., Yin, Y., Li, Q., Peng, Z., and He, D. (2024). Modeling the co-circulation of influenza and COVID-19 in Hong Kong, China. *Advances in Continuous and Discrete Models*, 2024(1):32.
- Wheeler, J., Abkemeier, A. J., and Ionides, E. L. (2025). Iterating marginalized Bayes maps for likelihood maximization. *In preparation*.
- Wheeler, J., Rosengart, A., Zhuxun, J., Tan, K. H. E., Treutle, N., and Ionides, E. L. (2024). Informing policy via dynamic models: Cholera in Haiti. *PLOS Computational Biology*, 20:e1012032.
- Wilkinson, D. J. (2018). *Stochastic Modelling for Systems Biology*. CRC press.
- Young, P. (1998). Data-based mechanistic modelling of environmental, ecological, economic and engineering systems. *Environmental Modelling & Software*, 13(2):105–122.