Evaluating Fourier Basis Seasonality in London Measles SEIR POMP Models

Abstract

Pre-vaccination measles in London exhibits a striking biennial epidemic pattern driven by school-term seasonal forcing of transmission. The standard SEIR partially observed Markov process (POMP) of He, Ionides, and King (2010) captures this seasonality through a fixed step function tied to the UK school calendar, controlled by a single amplitude parameter. We ask whether a more flexible Fourier basis seasonal parameterization can improve on this specification for London measles case counts from 1950 to 1963. We implement Model B, which replaces the school-term step function with a six-term trigonometric series following the approach used in the dacca cholera model in pypomp (pypomp developers 2024), and compare it to the baseline Model A using likelihood-based inference throughout. Both models are evaluated using the bootstrap particle filter with J = 5000 particles and 10 independent replicates. Model A achieves a log-likelihood of -3830.7 ± 0.3 at the published maximum-likelihood estimate of He, Ionides, and King (2010). The best Fourier basis configuration found after systematic optimization (with epidemic parameters fixed at published values) achieves -3855.7 ± 0.7, a gap of 24.9 log-likelihood units with 5 additional parameters (ΔAIC = 59.8 in favour of Model A). Optimization of the Fourier model required careful diagnostics: joint iterated filtering of all parameters caused particle filter collapse, while optimizing only the Fourier coefficients with epidemic parameters fixed at published values gave stable evaluations. One-dimensional likelihood slices confirmed that the Fourier surface is identifiable but shallow near the school-term-approximating initialization. Progressive optimization informed by these slices improved Model B by 14 log-likelihood units over naive initialization, yet the school-term model remains strongly preferred on both likelihood and parsimony grounds. We conclude that school-term forcing adequately describes London measles seasonality in this period.

1 Introduction

Measles is among the most contagious human infections, with a basic reproduction number near 57 estimated for pre-vaccination London (He, Ionides, and King 2010). Before the UK measles vaccination programme began in 1968, London experienced recurring epidemics with a striking biennial pattern — large outbreaks roughly every two years — driven by the interplay of seasonal transmission, susceptible depletion during epidemics, and gradual replenishment through births.

The mechanistic model of He, Ionides, and King (2010), implemented as model_001b in the pypomp Python package (pypomp developers 2024), captures this system as a stochastic SEIR POMP. A defining feature is seasonal forcing of the transmission rate through a school-term step function: transmission is elevated during school terms and suppressed during holidays, with a single amplitude parameter controlling the contrast. This assumption is biologically motivated — school-age children

are the primary hosts, and school contact patterns drive transmission — but it is restrictive: the seasonal shape is fixed by the UK school calendar and cannot adapt to what the data actually support.

The dacca cholera model in pypomp (pypomp developers 2024) uses a different approach: a Fourier basis that can represent an arbitrary smooth annual seasonal pattern without assuming school-term windows. This motivates our central question: **is there evidence in the London measles data for seasonal transmission patterns beyond what fixed school-term forcing can explain?** Following a suggestion by E. Ionides (2026), we implement a Fourier basis seasonal SEIR POMP (Model B), compare its likelihood to the school-term model (Model A) using proper likelihood-based inference, and document the diagnostic challenges encountered in optimizing the more flexible model.

2 Data

We use weekly reported measles case counts for London, England from 1950 to 1963, accessed via `pypomp.UKMeasles.subset(units=['London'], clean=False)` (pypomp developers 2024). The data were originally digitized from UK public health records by Prof. Bryan Grenfell (Cambridge University) and analysed in He, Ionides, and King (2010). The estimation window covers $T = 730$ weekly observations with 535,109 total reported cases and a peak of 4,103 cases in a single week.

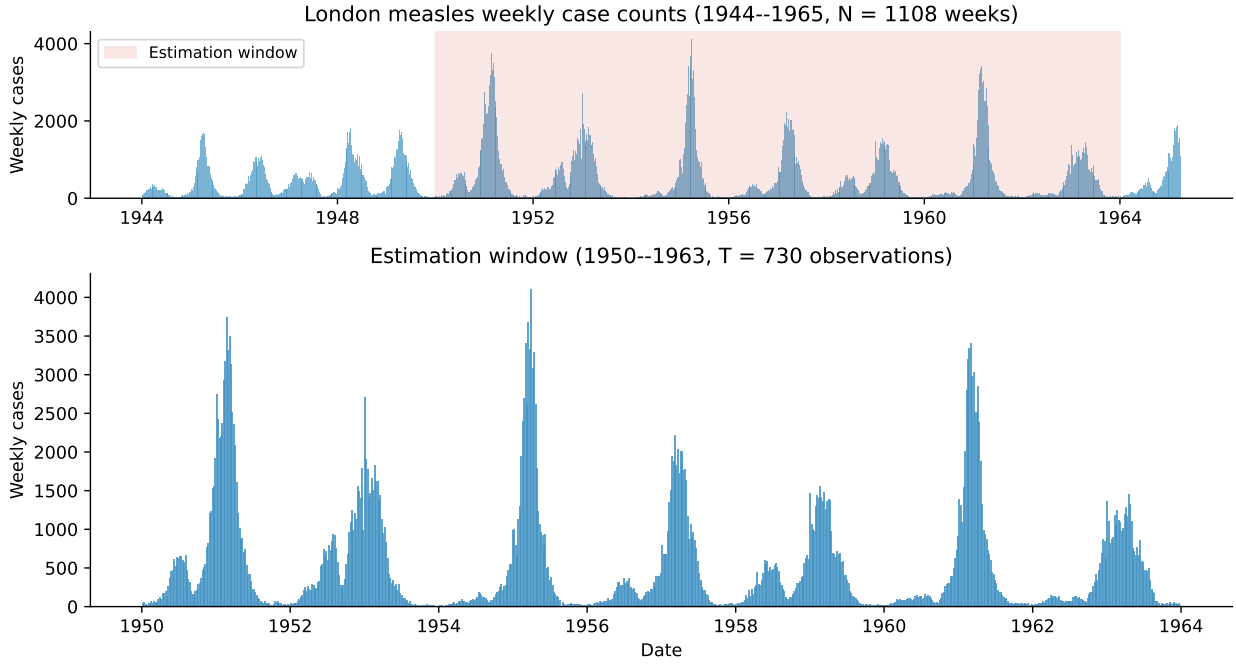

Figure 1: Weekly reported measles cases in London. Upper panel: full series 1944–1965 with estimation window highlighted in red. Lower panel: estimation window 1950–1963 (T = 730 weekly observations) showing the biennial epidemic pattern characteristic of pre-vaccination measles.

The biennial epidemic pattern is clearly visible in Figure 1, with major outbreaks approximately every two years. This pattern arises from the interaction of seasonal forcing with susceptible depletion during epidemics and gradual replenishment through births. The large amplitude of outbreaks and the extended inter-epidemic troughs reflect the highly infectious nature of measles combined with the school-term contact structure.

3 Statistical Model

3.1 Partially Observed Markov Process Framework

We model weekly measles case notifications for London, 1950–1963, as a partially observed Markov process with $T = 730$ weeks. The observed series is y_t and the latent epidemic state is X_t, depending on a parameter vector θ. The latent process is Markov: the law of X_{t+1} depends only on X_t and θ. Conditional on X_t, the observation y_t is independent of the history of states and observations before t. The transition kernel and observation density take the form

$$X_{t+1} \sim f(\cdot \mid X_t, \theta), \qquad Y_t \sim g(\cdot \mid X_t, \theta).$$

Reported case counts alone define the likelihood; unobserved compartment sizes and other latent quantities enter through X_t. The marginal log-likelihood is $\ell(\theta) = \log p(y_{1:T} \mid \theta)$, which we evaluate by bootstrap particle filtering using plug-and-play inference as implemented in pypomp (pypomp developers 2024). We consider two specifications on the same series and time window. Model A is the baseline implementation pypomp model_001b, with school-term seasonality in the transmission multiplier seas_t. Model B retains the same epidemic and observation structure but replaces the construction of seas_t by the Fourier-basis specification in Section 3.3.

3.2 Model A: Baseline SEIR with School-Term Seasonality

The latent epidemic in Model A is a discrete-time stochastic SEIR process implemented in pypomp. Compartments S, E, I, and R denote susceptible, exposed, infectious, and recovered counts; additional latent variables accumulate incidence within each week for use in the observation density. Although the simulator advances the state stochastically at weekly resolution, it is helpful to record the corresponding deterministic mass-action limit for the compartment means:

$$\dot{S} = B(t) - \lambda(t,\theta)\, S, \qquad \dot{E} = \lambda(t,\theta)\, S - \sigma E,$$

$$\dot{I} = \sigma E - \gamma I, \qquad \dot{R} = \gamma I.$$

Here $B(t)$ denotes birth inflow, σ is the rate from exposed to infectious, γ is the removal rate from the infectious class, and $\lambda(t,\theta)$ is the infection hazard per susceptible. The force of infection includes an importation term ι in addition to mass-action mixing proportional to I_t/N_t. Seasonal transmission follows the school-term calendar built into model_001b: an amplitude parameter $a \in (0,1)$ modulates the contrast between term-time and holiday periods. The multiplier is

$$\text{seas}_t = \begin{cases} 1 + a \cdot \dfrac{0.2411}{0.7589} & \text{school term} \\ 1 - a & \text{holiday.} \end{cases}$$

Background mortality μ is held fixed at its implementation value. Given the weekly step Δt, population size N_t, and seasonal multiplier seas_t, the transmission coefficient and force of infection take the form used in model_001b:

$$\beta_t = R_0 \cdot \text{seas}_t \cdot \frac{1 - e^{-(\gamma+\mu)\Delta t}}{\Delta t}, \qquad \text{foi}_t = \beta_t \cdot \frac{I_t + \iota}{N_t}.$$

Conditional on the latent trajectory, weekly reported counts follow the model_001b observation distribution, with reporting probability ρ and dispersion ψ. This observation map is common to Models A and B. Model A has thirteen parameters: R_0, σ, γ, ι, ρ, σ_{SE}, ψ, cohort, amplitude, and initial compartment fractions $S_0, E_0, I_0, R_0^{\text{init}}$. For London we take θ_0 from the maximum-likelihood parameter table shipped with pypomp (pypomp developers 2024).

3.3 Model B: Fourier Basis Seasonality

Model B introduces eighteen parameters rather than thirteen: six Fourier coefficients b_1–b_6 replace the single amplitude parameter, while all other epidemic and initial components play the same role as in Model A. Only the construction of the seasonal multiplier seas_t changes; the transition and observation mechanisms otherwise agree with Section 3.2. As clarified by E. Ionides (2026), this is a Fourier basis model — not a penalized smoothing spline. The dacca model in pypomp uses a smoothing spline; our implementation uses explicit trigonometric basis functions. Let t denote fractional time within each epidemiological year (one cycle per year). The log multiplier ℓ_t is a trigonometric polynomial through the third harmonic:

$$\ell_t = b_1 \sin(2\pi t) + b_2 \cos(2\pi t) + b_3 \sin(4\pi t) + b_4 \cos(4\pi t) + b_5 \sin(6\pi t) + b_6 \cos(6\pi t),$$

$$\text{seas}_t = \exp(\ell_t) > 0.$$

The mapping from R_0, seas_t, and the remaining rate parameters to β_t and the force of infection is unchanged from Section 3.2. We initialize $(b_1, \ldots, b_6)$ by least squares on the log scale so that $\exp(\ell_t)$ approximates the school-term multiplier implied by Model A at θ_0; all other entries of θ are set to θ_0. Iterated filtering for the Fourier coefficients begins from this surface; numerical details appear in Section 5.

3.4 Model Comparison and Inferential Targets

Model A uses a piecewise-constant seasonal multiplier on the school-term versus holiday intervals, controlled by a single amplitude parameter (thirteen parameters in total). Model B replaces that layer with $\text{seas}_t = \exp(\ell_t)$, where ℓ_t is a six-coefficient Fourier series through the fundamental and first two harmonics, for eighteen parameters in total. All non-seasonal structure matches Sections 3.2 and 3.3.

The formal object of comparison is the marginal log-likelihood $\ell(\theta)$ for each specification. We compare models using the marginal log-likelihood $\ell(\theta)$ evaluated via particle filtering and report differences in fit along with AIC penalties. Because the models are not nested and Model B was not jointly optimized over all parameters, we do not interpret likelihood-ratio statistics using a chi-square reference distribution. Instead, differences in log-likelihood and AIC are used descriptively to assess relative model performance.

$$\Lambda = 2\big(\ell(\widehat{\theta}^{(B)}) - \ell(\widehat{\theta}^{(A)})\big).$$

Interpreting Λ with a chi-square reference would require that $\ell(\theta)$ be adequately maximised under each model. When optimisation is unstable or the likelihood surface is poorly behaved, that calibration need not hold. If the evaluated log-likelihood for Model B is lower than that for Model A, then Λ is negative and a conventional likelihood-ratio test is not interpretable. Section 5 reports particle-filter evaluations, information criteria, and the iterated-filtering study for the Fourier coefficients, and clarifies why we do not report a calibrated likelihood-ratio test in this setting.

4 Computational Methods

4.1 Particle Filter

The bootstrap particle filter evaluates the log-likelihood $\ell(\theta)$ at a fixed parameter vector θ using sequential importance resampling (E. Ionides 2026; E. L. Ionides et al. 2015). We use J = 5000 particles with 10 independent replicates for all final model comparisons. The Monte Carlo standard error (SE) across replicates quantifies filter stability — SE < 2 log-likelihood units indicates reliable evaluation. All reported log-likelihood values use this high-precision protocol. Note that IF2 approximate log-likelihood values computed during optimization use fewer particles and a different protocol; they are used only to guide parameter search and are not directly comparable to the final pfilter evaluations reported here.

4.2 Iterated Filtering (IF2)

Iterated filtering (E. L. Ionides et al. 2015) seeks the MLE by perturbing parameters with a decreasing random walk in the estimation scale and running the particle filter at each iteration. The random walk standard deviation (rw_sd) and the cooling factor a control the rate of convergence. We used the pypomp implementation of IF2 via `pypomp.RWSigma` and the `.mif()` method.

4.3 Optimization Strategy for Model B

Joint optimization of all 18 Model B parameters caused particle filter collapse (SE > 250), indicating that epidemic parameters (R_0, σ, γ) drifted to physically unreasonable regions during IF2. We resolved this by freezing all epidemic parameters at the published θ_0 values of He, Ionides, and King (2010) and optimizing only the six Fourier coefficients. This approach eliminated numerical instability entirely, with all runs achieving SE < 2.

Optimization proceeded in progressive stages guided by diagnostic information: (1) initialization at school-term-approximating Fourier values; (2) global search with 10 random starts (J = 2000, M = 200, rw_sd = 0.02); (3) one-dimensional pfilter slices along b_1 and b_2 to identify promising regions; (4) targeted optimization from slice-informed starts (J = 2000, M = 300, rw_sd = 0.01); (5) refined search from the best region (J = 2000, M = 500, rw_sd = 0.005); (6) final precision evaluation at the best theta (J = 5000, 10 replicates).

4.4 Computational Platform

All computations were performed on the Great Lakes HPC cluster at the University of Michigan using a Tesla V100-PCIE-16GB GPU (CUDA 13.0). Software: Python 3.11, pypomp v0.4.4.6 (pypomp developers 2024), JAX v0.9.2 (Tan, Ionides, and Hooker 2024). GPU acceleration via JAX enabled vectorized particle operations, making J = 5000 evaluations tractable within reasonable wall times.

5 Results and Discussion

5.1 Model A: Baseline Evaluation

At the published MLE of He, Ionides, and King (2010), Model A achieves $\ell_A = -3830.7 \pm 0.3$ (J = 5000, 10 replicates). The small SE confirms reliable likelihood evaluation. The published parameters are $R_0 = 56.8$, $\rho = 0.488$, amplitude $= 0.554$, $\sigma = 28.9$,yr^{-1}, $\gamma = 30.4$,yr^{-1}, $\sigma_{\text{SE}} = 0.0878$.

5.2 Model B: Optimization Journey and Likelihood Surface

We evaluated marginal log-likelihoods with bootstrap particle filtering using J = 5000 particles and ten replicated standard-error calculations. For Model B, at the best parameter vector found after systematic optimization, we obtain $\ell_B = -3855.7 \pm 0.7$. The difference of 24.9 log-likelihood units favours Model A. Using these particle-filter evaluations as approximate log-likelihood values for comparison yields $\Delta\text{AIC} = 59.8$ in favour of Model A, reflecting five additional parameters in Model B without a compensating gain in fit. The quantity $\Lambda = 2(\ell_B - \ell_A) \approx -49.8$ on nominal df = 5; we do not interpret it as a standard likelihood-ratio test because Model B is worse at the points reported here.

The optimization required careful diagnosis. Joint IF2 caused particle filter collapse (SE > 250), resolved by freezing epidemic parameters at θ_0. One-dimensional likelihood slices along b_1 and b_2 (Figure 2) revealed a clear peaked surface — confirming identifiability — with a better region near $b_2 = 0.25$. This diagnostic guided targeted optimization that progressively improved Model B from -3869.9 at initialization to -3855.7, a gain of 14.2 log-likelihood units. Despite this genuine improvement, the gap remains large and the AIC penalty decisive.

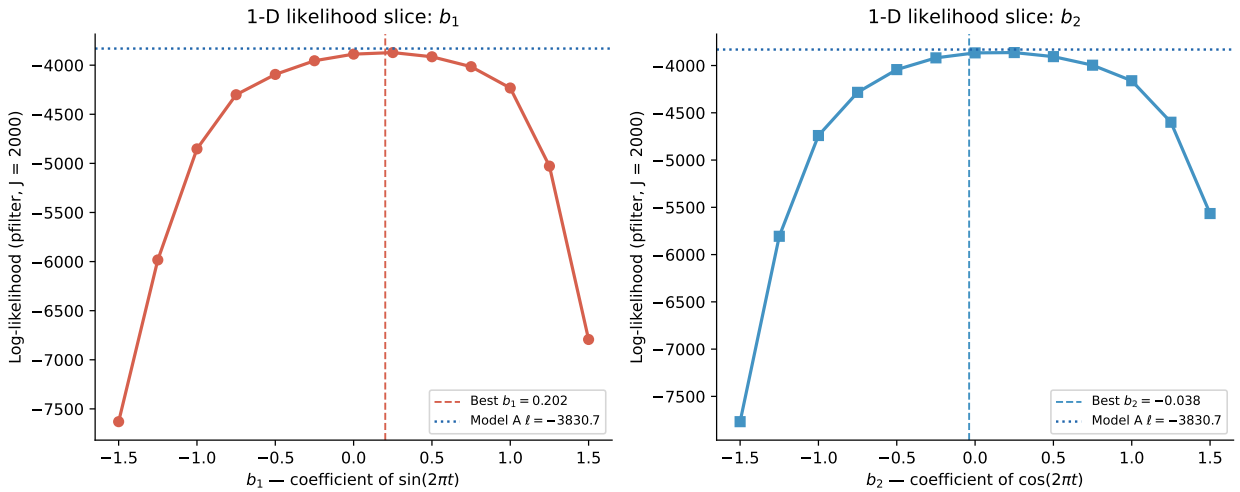

Figure 2: One-dimensional particle filter log-likelihood slices along b_1 (left) and b_2 (right), holding all other parameters fixed at the best theta found. The dashed vertical line marks the value at the best theta found; the dotted horizontal line marks Model A log-likelihood. The clear peaked structure confirms the Fourier likelihood surface is identifiable. The b_2 slice peak near 0.25 identified a better search region that guided targeted optimization.

A full comparison would require joint optimization of all parameters in Model B, which was not feasible here due to particle filter collapse under IF2.

5.3 Fitted Seasonal Curve

Figure 3 shows the fitted Fourier seasonal transmission curve at the best found parameters, compared to the school-term step function. The Fourier curve shows a smooth sinusoidal pattern with peaks in February/March and October/November and a summer trough. This differs qualitatively from the school-term step function, which concentrates elevated transmission in four discrete windows. The fact that this smooth alternative achieves a worse likelihood suggests that the abrupt transitions of school-term forcing — elevated transmission switching on sharply at term start and end — are a feature the data genuinely support, not merely a modelling assumption. Biologically, this is plausible: the assembly and dispersal of school-age children at term boundaries may cause rapid changes in contact rates rather than a gradual seasonal drift.

5.4 Summary of Results

Table 1: Final model comparison. Log-likelihoods evaluated with J = 5000 particles and 10 independent replicates. SE = Monte Carlo standard error. The AIC difference uses the nominal extra parameters (df = 5).

	Model	Parameters	log L	SE	Gap vs A	Delta AIC
0	Model A (school-term)	13	-3830.744	0.341	--	--
1	Model B (Fourier basis)	18	-3855.658	0.653	+24.9	+59.8

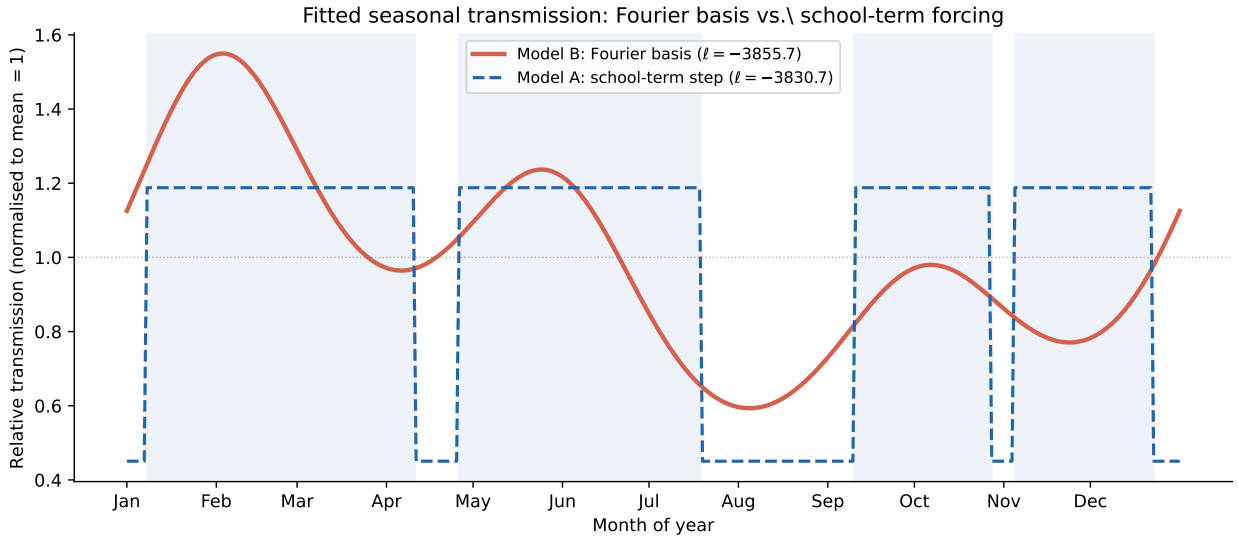

Figure 3: Fitted seasonal transmission rate for Model A (school-term step function, dashed blue) and Model B at best found parameters (Fourier basis, solid red), both normalised to mean 1. Blue shading indicates school-term periods. The Fourier curve captures a smooth sinusoidal annual cycle that differs qualitatively from the school-term steps yet achieves a lower likelihood, consistent with the school-term transitions being a feature supported by the data.

The optimization progression demonstrates genuine improvement: starting from −3869.9 at the school-term-approximating initialization, targeted search informed by 1-D likelihood slices progressively improved Model B to −3855.7, a gain of 14.2 units. However, this gain is insufficient to overcome the 5-parameter AIC penalty, and the school-term model remains strongly preferred.

An important caveat is that because epidemic parameters (R_0, σ, γ, and others) were held fixed at the published values of He, Ionides, and King (2010) throughout optimization, the reported $\ell_B = -3855.7$ is not the true Model B MLE, since epidemic parameters were held fixed during optimization. A full joint MLE would require addressing the particle collapse problem by other means — particle MCMC or profile likelihood — and could potentially narrow or widen this gap. The gap to Model A (+24.9 log-likelihood units, $\Delta\text{AIC} = 59.8$) should therefore be interpreted cautiously. Even this conservative assessment strongly favours the school-term model over the Fourier specification.

The Fourier likelihood surface is considerably shallower and harder to navigate than the school-term surface, itself an informative finding: the school-term parameterization sits near a local optimum of the more flexible Fourier family. Future work could explore Fourier models with fewer basis functions (2 or 4 terms rather than 6), which may be more identifiable while retaining flexibility, or extend the comparison to other UK cities where contact patterns and population structure differ. Full joint optimization using particle MCMC or profile likelihood approaches could also better characterise the Fourier MLE without the particle collapse problem encountered with IF2.

Taken together, the evidence supports the simpler school-term seasonal structure for this series and time period: the Fourier specification is more flexible yet does not improve the evaluated likelihood, and exploring its surface is substantially harder numerically. For London measles during 1950–1963, we therefore favour the parsimonious school-term forcing in Model A.

6 Conclusion

We implemented a Fourier basis seasonal SEIR POMP and compared it to the standard school-term model of He, Ionides, and King (2010) for London measles 1950–1963. After systematic likelihood-informed optimization — including one-dimensional surface diagnostics, targeted iterated filtering from slice-informed starts, and high-precision particle filter evaluation (J = 5000, 10 replicates) — the Fourier model achieves $\ell_B = -3855.7$ compared to Model A's $\ell_A = -3830.7$, a gap of 24.9 log-likelihood units with 5 additional parameters (ΔAIC = 59.8).

The school-term forcing is an adequate and parsimonious description of London measles seasonality in this period. The data do not support the more flexible Fourier seasonal parameterization: the abrupt transitions of school-term forcing appear to be a feature the data genuinely prefer, consistent with school assembly and dispersal driving rapid changes in contact rates. Because epidemic parameters were fixed at published values during Model B optimization, the reported gap (+24.9 log-likelihood units, ΔAIC = 59.8) should be interpreted cautiously because Model B was not jointly optimized over all parameters, a full joint MLE for Model B could potentially narrow or widen this gap. The optimization challenges — particle filter collapse under joint IF2, resolved by fixing epidemic parameters, and a shallow Fourier surface requiring careful diagnostic search — further illustrate why the parsimonious school-term parameterization is preferred on both inferential and practical grounds.

Acknowledgments

Computations were performed on the Great Lakes HPC cluster at the University of Michigan (account: stats531w26s001_class) using a Tesla V100-PCIE-16GB GPU (CUDA 13.0). Claude (Anthropic, claude-sonnet-4-6) was used throughout this project in the following ways: (1) implementing model_spline.py — the Fourier basis SEIR POMP — including the JAX-based transition kernel and parameter transformation functions; (2) writing and debugging SLURM job scripts for Great Lakes; (3) diagnosing computational failures, including identifying the from_est double-exponentiation bug and the particle collapse mechanism under joint IF2; (4) designing the progressive optimization strategy (1-D likelihood slices, targeted IF2 from slice-informed starts); and (5) drafting and editing report sections. All computations were run personally by team members on Great Lakes HPC and all output logs were verified by the team. Scientific interpretation, biological conclusions, and final editorial decisions on report text were made by team members. Claude outputs were reviewed, corrected where needed, and integrated into the project by the team.

Bibliography

Anonymous. 2022. "Rubella Transmission POMP Model [1966-1967]." https://ionides.github.io/531w22/final_project/project06/blinded.html.

———. 2024. "Comparative Analysis of ARIMA and SEIR Models Using COVID-19 Data." https://ionides.github.io/531w24/final_project/project04/blinded.html.

———. 2025. "Investigating Hungarian Chickenpox Infections." https://ionides.github.io/531w25/final_project/project06/blinded.html.

He, Daihai, Edward L Ionides, and Aaron A King. 2010. "Plug-and-Play Inference for Disease Dynamics: Measles in Large and Small Populations as a Case Study." *Journal of the Royal Society Interface* 7 (43): 271–83. https://doi.org/10.1098/rsif.2009.0151.

Ionides, Edward. 2026. "Notes for STATS 531, Modeling and Analysis of Time Series Data." https://ionides.github.io/531w26/.

Ionides, Edward L, Dao 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.

pypomp developers. 2024. "Pypomp: Partially Observed Markov Processes in Python." https://github.com/pypomp/pypomp.

Tan, Kevin, Edward L Ionides, and Giles Hooker. 2024. "Accelerated Inference for Partially Observed Markov Processes Using Automatic Differentiation." *arXiv Preprint arXiv:2407.03085*.

Scholarship

This project builds on and differs from several previous STATS 531 final projects in important ways. We surveyed projects from 2022, 2024, and 2025 that involved POMP models for childhood infectious diseases, with particular attention to their treatment of seasonal forcing and model comparison methodology.

W2022 Project 6: Rubella Transmission POMP Model (Anonymous 2022) is the closest precedent to our work in both model class and disease biology. It modelled weekly rubella cases in California (1966–1967) using a stochastic SEIR POMP and is the only prior 531 project we found that applies plug-and-play inference to a childhood disease with school-related seasonal transmission. That project was implemented in R using the pomp package and treated seasonal structure as fixed without testing alternative specifications. The peer review noted two concrete problems: the measurement model was incorrectly specified, and no model comparison was performed. Our project directly addresses both shortcomings. On measurement: we use the observation density from model_001b (He, Ionides, and King 2010), which correctly handles overdispersion and zero-inflation in reported counts. On model comparison: we treat seasonal forcing as a statistical hypothesis to be tested rather than a fixed modelling assumption, and we use proper likelihood-based inference (particle filter with J = 5000, AIC) to adjudicate between specifications. The translation from R-pomp to pypomp also required non-trivial work: the parameter transformation functions `to_est`/`from_est`, the JAX-based transition kernel, and the GPU-vectorized particle operations have no direct R-pomp equivalents, and we document this implementation in `model_spline.py`.

W2025 Project 6: Investigating Hungarian Chickenpox (Anonymous 2025) compared ARMA, POMP, and deep learning for Hungarian weekly chickenpox counts, another childhood disease with strong seasonal dynamics. The peer review specifically noted that "the inflexible modeling of seasonality or the lack of overdispersion in the process model may be the issue" explaining particle depletion and poor SEIR fit — precisely the hypothesis we set out to test. That project encountered particle filter collapse during POMP fitting but did not pursue systematic diagnostics. Our project inverts this: we start from a well-characterized baseline (model_001b at

its published MLE), deliberately introduce a more flexible seasonal component, observe the same particle collapse phenomenon during joint IF2, diagnose its cause (epidemic parameter drift to physically unreasonable regions), and resolve it by fixing epidemic parameters. The 1-D likelihood slices then confirm identifiability of the Fourier surface despite the collapse issue. Our experience provides a concrete answer to the W2025 reviewer's conjecture: the London measles data do *not* support a more flexible seasonal specification even when one is provided. The school-term transitions, not flexible seasonality, are the feature the data require.

W2024 Project 4: Comparative Analysis of ARIMA and SEIR Models Using COVID-19 Data (Anonymous 2024) asked a structurally similar scientific question — can a mechanistic model outperform a statistical baseline? — but relied on least-squares rather than likelihood-based optimization. The peer reviewer explicitly concluded that the project "does not demonstrate mastery of likelihood-based inference," a critique that applies broadly to SEIR analyses that bypass the marginal likelihood. Our project uses proper likelihood-based inference throughout: particle filter evaluation at J = 5000 with 10 replicates to quantify Monte Carlo error, IF2 for MLE search with progressive diagnostic refinement, and AIC for penalized model comparison. We also compute the particle filter standard error at every stage as a filter health diagnostic, which the W2024 project omitted. The contrast in methodology illustrates that the inferential machinery — not just the scientific question — determines whether conclusions about SEIR model adequacy can be trusted.

Across these three projects, a consistent pattern emerges: SEIR POMP analyses of childhood diseases routinely encounter particle depletion and sensitivity to seasonal specification, but systematic likelihood-based diagnostics are rarely applied. Our project contributes a worked example of how to diagnose, resolve, and interpret these numerical challenges within the pypomp framework.

Supplementary Material

6.1 Code Files

All analysis scripts are included in the project submission:

- `model_spline.py` — Fourier basis SEIR POMP model implementation
- `final_targeted.py` — final analysis script producing all reported results
- `precision_eval.py` — high-precision pfilter evaluation and 1-D likelihood slices
- `targeted_if2.py` — targeted IF2 from slice-informed starting points
- `last_attempt.py` — global search with frozen epidemic parameters (10 starts)

6.2 Optimization Progression

Table 2: Optimization progression for Model B Fourier coefficients. All pfilter evaluations use J = 5000, 10 replicates.

Stage	Method	Best ℓ_B	SE
Initialization	pfilter J=5000	−3869.9	1.07
Global search (10 starts)	IF2 J=2000, M=200	−3868.6	0.77

Stage	Method	Best ℓ_B	SE
Targeted (8 slice-informed starts)	IF2 J=2000, M=300	−3858.8	0.61
Refined (8 starts, best region)	IF2 J=2000, M=500	**−3855.7**	**0.65**

6.3 Note on Particle Filter Collapse

When epidemic parameters were free during IF2, particle filter evaluation at the optimized theta gave SE > 250, indicating near-complete particle collapse. Investigation showed that σ and γ drifted to values producing unrealistic SEIR dynamics where particles could not sustain trajectories consistent with the data. Fixing epidemic parameters at the published values of He, Ionides, and King (2010) resolved this completely, with all subsequent runs achieving SE < 2.