# STATS 531 final project


### Abstract

This project develops a mechanistic SEIR model with a latent serotype switching state to analyze dengue fever outbreaks in Taiwan during 2014–2016. Weekly confirmed case counts are modeled using iterated filtering (IF2) within the POMP framework, implemented in Python via pypomp and JAX. The model incorporates seasonal transmission via covariates, serotype specific transmission effects, and a negative binomial model. A two-stage estimation procedure consisting of a global search over 60 random starts and a local refinement successfully identifies a high likelihood parameter region, with the serotype switching rate omega estimated via MCAP profile maximum likelihood estimate at 0.0392 (95% CI: 0.0123, 0.0434). Probe diagnostics reveal systematic model misfit, particularly in epidemic peak magnitude and overall trend, suggesting the model over predicts outbreak size. Comparison against ARMA and GARCH benchmarks shows the fitted SEIR particle filter achieves a modestly higher log-likelihood than the best GARCH competitor, though the margin is narrow. These results suggest that while the latent serotype mechanism provides a biologically interpretable framework for dengue fever outbreaks, further model refinement is needed to adequately capture the scale and structure of consecutive outbreaks.


## 1 Introduction

### 1.1 Background

Dengue fever is a mosquito transmitted disease and is not transmissible between humans, but rather can be transmitted from mosquito to mosquito directly or between two mosquitoes via an infected human (Strickler 2018). It comes in four serotype strains: DENV1-4. DENV-2 is considered to be the most associated with severe dengue cases (Dinkar and Singh 2026). This disease is more common in tropical climates like Taiwan where there were significant outbreaks in 2014, 2015, and 2023 (Taiwan Centers for Disease Control 2024). It may become more pervasive with environmental changes due to climate change. Since Taiwan is an advanced economy with a robust health system, the death rate due to dengue fever is consistently below 1% (Wei, Shu, and Hung 2016), but learning to model the disease is useful as it could possibly worsen in years to come. This research focuses on the back to back outbreaks of 2014 and 2015.

### 1.2 Research Questions

We are interested in examining if serotype can effectively simulate dengue outbreaks and how often a dominant serotype switches in a given time frame. To do this, compare simulations from our SEIR model against the observed data of the outbreaks to see if it can reasonably approximate using the disease serotype as a latent state. Also, assessing the serotype switch parameter via profile likelihood is also desired. To further examine the effectiveness of our mechanistic model, we compare against the simpler benchmarks of ARMA and GARCH models.

## 2 Data Aggregation and Interval Selection

Our dataset comes from Kaggle (taweilo 2024). We aggregate cases on a weekly level to reflect dengue fever's recovery period of three to eight days (Ligon 2005). This also reduces some of the noise compared with aggregating on a daily interval. The 2014 and 2015 outbreaks are of greater interest to us due to how close together they are. With the possibility of dengue fever becoming more and more common with time, back to back outbreaks could become more frequent and it is important to be prepared. Simply modeling single outbreak every couple of years may not be sufficient for the epidemiological reality in the future.

```
First week: 2013-12-30 00:00:00
Last week: 2016-05-30 00:00:00
Number of weeks: 127
```

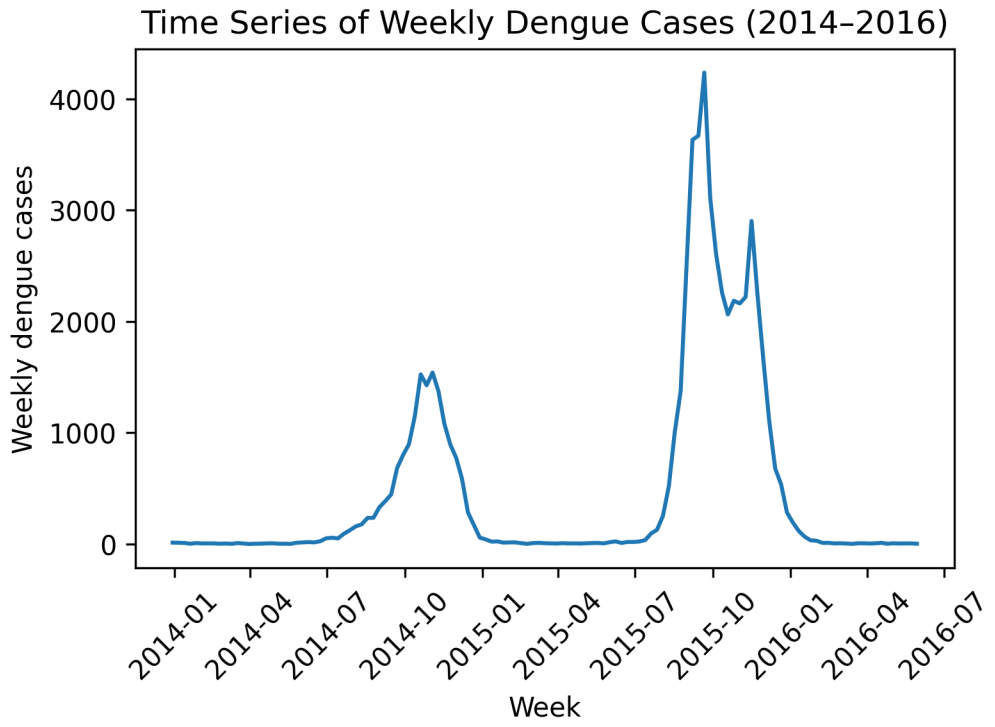

# 3 Model Setup

## 3.1 Parameters

SIR model is common in epidemiological studies. Since dengue is not transmitted directly between humans but between mosquitoes (Strickler 2018), we include $E(t)$ to control the infection time between mosquitoes. The SEIR model we use is based on the simpler SEIR presented in the lectures (Ionides 2026). Our SEIR included a latent serotype state $K(t)$ and a time variant transmission rate $\beta(t)$ to include effects of seasonality. Specifically, the latent states are:

$$X(t) = (S(t), E(t), I(t), R(t), H(t), K(t))$$

where $S(t), E(t), I(t), R(t)$ are the standard susceptible, exposed, infectious, and recovered counts, respectively. $H(t)$ is an accumulator counting new infections in week $t$. $K(t) \in \{1, 2, 3, 4\}$ is a latent indicator representing serotype. The relationship $S(t) + E(t) + I(t) + R(t) \leq N(t)$ holds for all $t$. Also, our serotype state evolves as a Markov chain at each time step.

$$P(K_{t+1} = j \mid K_t = i) = \begin{cases} 1 - \omega, & \text{if } j = i \\ \omega/3, & \text{if } j \neq i \end{cases}$$

where $\omega \in (0, 1)$ is the weekly rate of serotype switching. We assume equal switching probability, that no serotype is inherently more likely to appear than another. Smaller values of $\omega$ suggest longer dominating periods of one serotype, while larger values would indicated more frequent switching of the dominant serotype.

We use a time dependent baseline transmission rate as dengue fever transmission is driven by mosquito populations, which are seasonal since they depend highly on rainfall and temperature patterns. The transmission rate is given by a seasonal component as well as a serotype specific effect.

$$\log \beta(t) = b_0 + b_1 \sin\left(\frac{2\pi(t+\phi)}{52.18}\right) + b_2 \cos\left(\frac{2\pi(t+\phi)}{52.18}\right) + b_3 \sin\left(\frac{4\pi(t+\phi)}{52.18}\right) + b_4 \cos\left(\frac{4\pi(t+\phi)}{52.18}\right) + \kappa_{K_t}$$

where $t$ is the week of the year, $\phi = 20$ is a phase offset chosen to align the seasonal cycle with the observed timings of the outbreaks, 52.18 is the average number of weeks in a year, and $\kappa_k$ is a serotype specific effect with $\kappa_1 = 0$ as the reference level.

Given that the outbreaks span 2.5 years, we allow time variation in $N(t)$ to account for population change in Taiwan (United Nations, Department of Economic and Social Affairs, Population Division 2024). Specifically, the sampling population varies by year:

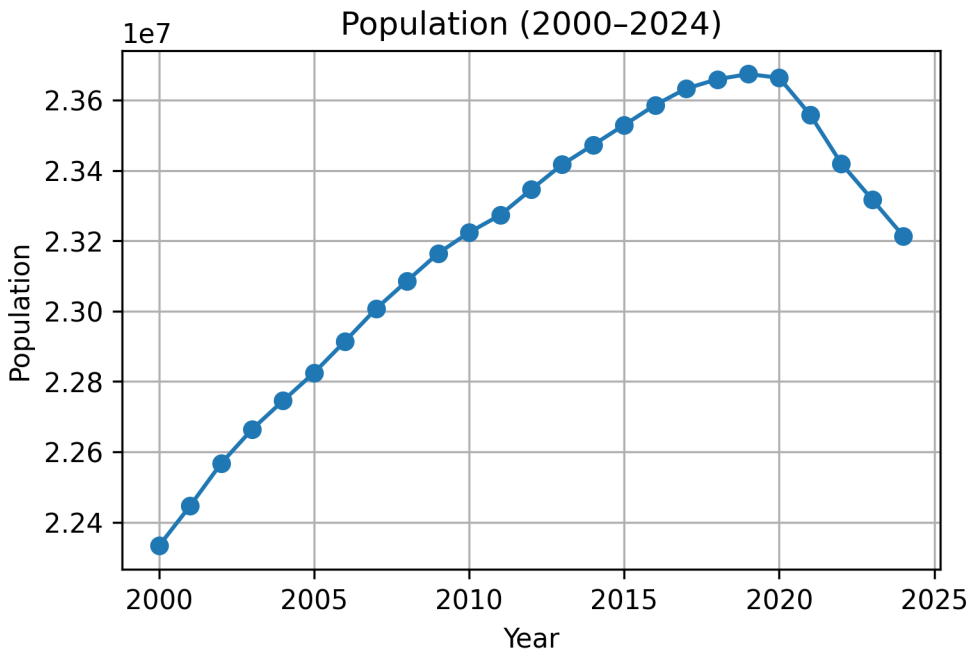


In this research, the effective sample size is $ESS = N(t) \times 0.01$.

Stochastic transition equations are drawn over each daily sub-step $\Delta t = 1/7$ week as:

$$
\begin{aligned}
\Delta N_{SE} &\sim \text{Binomial}\left(S(t), 1 - e^{-\beta_t I_t / N_t \Delta t}\right) \\
\Delta N_{EI} &\sim \text{Binomial}\left(E(t), 1 - e^{-\mu_{EI} \Delta t}\right) \\
\Delta N_{I,R+D} &\sim \text{Binomial}\left(I(t), 1 - e^{-(\mu_{IR} + \mu_{ID}) \Delta t}\right) \\
\Delta N_{IR} &\sim \text{Binomial}\left(\Delta N_{I,out}, \frac{\mu_{IR}}{\mu_{IR} + \mu_{ID}}\right) \\
\Delta N_{ID} &= \Delta N_{I,R+D} - \Delta N_{IR}
\end{aligned}
$$

where $\mu_{EI}$ is the transmission rate from $E(t)$ to $I(t)$, $\mu_{IR}$ is the recovery rate, and $\mu_{ID}$ is the mortality rate from dengue. Under the assumption that recovery and dengue mortality are independent exponential processes, the probability that an individual leaving $I(t)$ recovers rather than dies is given by $\mu_{IR}/(\mu_{IR} + \mu_{ID})$. The state update then is given by:

$$
\begin{aligned}
S_{t+1} &= S_t - \Delta N_{SE} \\
E_{t+1} &= E_t + \Delta N_{SE} - \Delta N_{EI} \\
I_{t+1} &= I_t + \Delta N_{EI} + \iota_t - \Delta N_{IR} - \Delta N_{ID} \\
R_{t+1} &= R_t + \Delta N_{IR} \\
H_{t+1} &= \Delta N_{EI}
\end{aligned}
$$

where $\iota_t = \iota_{\text{const}} + \text{imports}_t$ is the weekly flow of infected individuals coming in from outside the modeled population.

At $t = 0$, the susceptible category is initialized as $S_0 = \eta N_0$ where $\eta \in (0, 1)$ is the initial proportion of the population that is susceptible. It is treated as a free parameter. This allows the model to take into account prior immunity which is important in the context of Taiwan, as dengue has a long history in the region (Taiwan Centers for Disease Control 2024). $E_0$ and $I_0$ are also treated as free parameters, and $R_0 = \max(N_0 - S_0 - E_0 - I_0, 0)$. The latent serotype is initialized as serotype DENV-1 ($K_0 = 1$).

Finally, reported weekly case counts $Y_t$ are assumed to follow a negative binomial distrubution conditional on the latent state following from (Ionides 2026)

$$Y_t \mid X_t \sim \text{NegBin}(\mu_t, k)$$

where

$$\mu_t = \max(\rho \times (H + 0.01 \times I), 0.01)$$

with $\rho \in (0, 1)$ as the reporting probability and $k$ the overdispersion parameter. The $0.01 \times I_t$ term is a floor for when $H_t = 0$ to prevent degeneracy.

## 3.2 Simulation

Based on plug and play model simulations, serotype appears to be able to simulate the dengue outbreaks effectively. The $\kappa$ values are derived from the data as serotypes one and two are equally dominant in quantity compared to serotypes three and four, which are far less common. The values of $\kappa_3$ and $\kappa_4$ are exponentially transformed proportionally to the baseline of $\kappa_1 = \kappa_2 = 0$. The simulations are as follow:

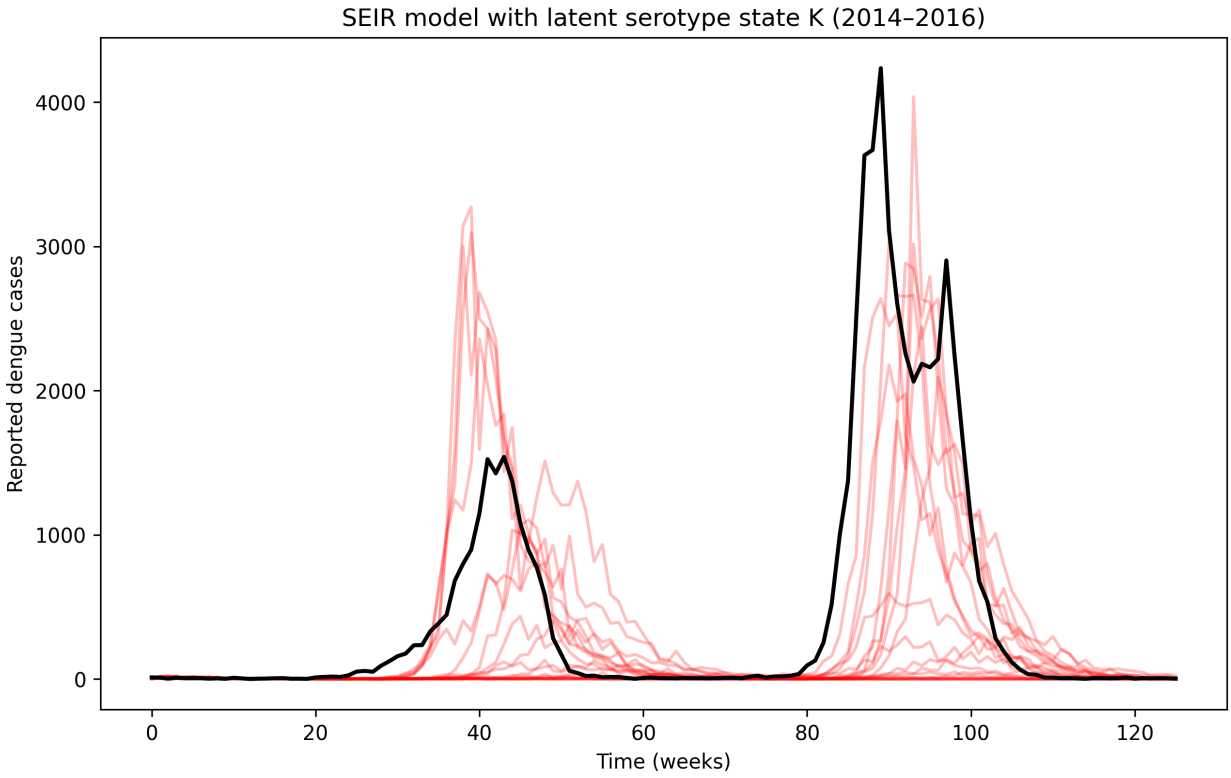


# 4 Particle Filter

## 4.1 Local and Global Search of Parameters

Our baseline log-likelihood before filtering is:

```
Baseline log-likelihood: -6966.50
```

The local search shows log-likelihood chains converging quite rapidly, within the first 20 iterations and stabilizing relatively near zero, given the scale. As for specific parameters, $\eta$, $\mu_{EI}$, and $\mu_{ID}$ appear the hardest to identify. The rest show some degree of convergence to a tighter range of values. The parameters involved in $\beta(t)$ largely converge and can be well identified based on this search. $\kappa_1$ and $\kappa_2$ are fixed to represent already known rates of prevalence for serotypes one and two, the main serotypes of interest. The log-likelihood obtained from the local search is:

```
Best final IF2 logLik = -569.20
Mean final IF2 logLik = -597.13
```

Dengue latent-serotype SEIR — local IF2 convergence traces

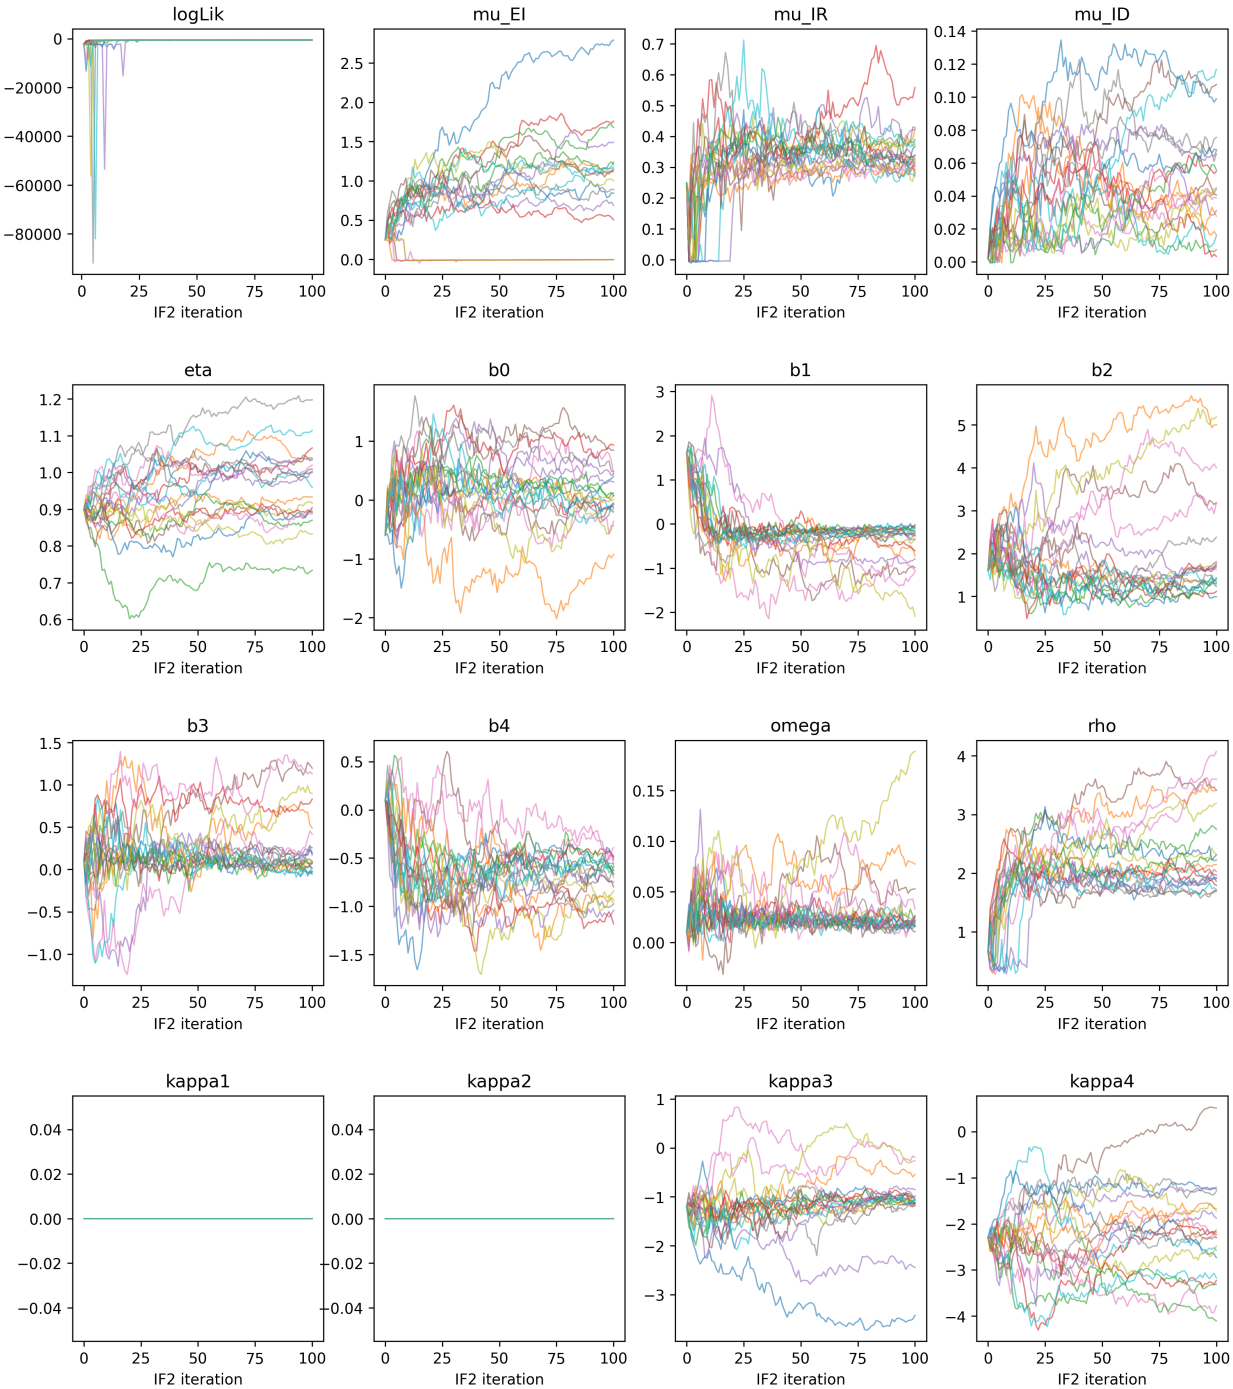

The global search has noisier log-likelihood with less convincing convergence. Chains span a wide range and are volatile over 100 iterations. $\mu_{EI}$, $\mu_{ID}$, $\eta$, and $\rho$ are difficult to identify.

Dengue latent-serotype SEIR — global IF2 convergence traces

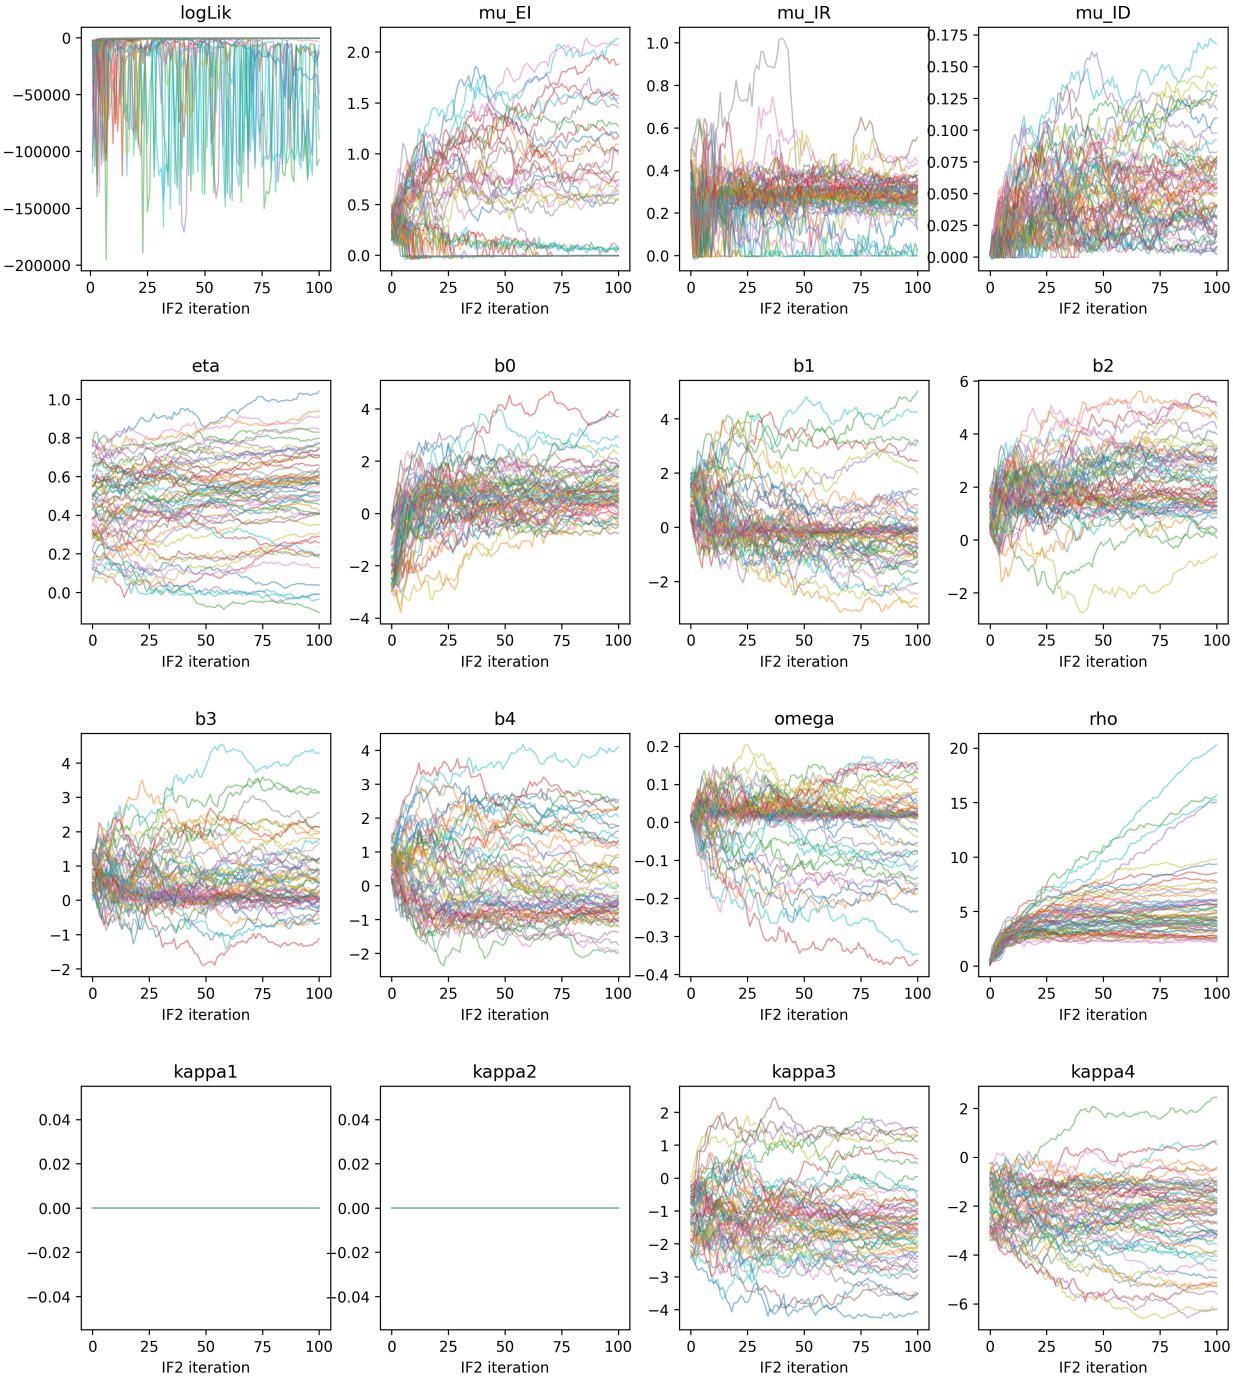

The best global search parameters and log-likelihood are:

| | |
|---|---|
| mu_EI | 2.081499 |
| mu_IR | 0.337101 |
| mu_ID | 0.054756 |
| kappa1 | 0.000000 |
| kappa2 | 0.000000 |
| kappa3 | -0.945812 |
| kappa4 | -2.981178 |
| I0 | 400.000000 |
| E0 | 400.000000 |
| K0 | 0.000000 |
| eta | 0.912184 |
| b0 | -0.046065 |
| b1 | -0.093949 |
| b2 | 1.111901 |
| b3 | 0.013224 |
| b4 | -0.483489 |
| iota | 0.030000 |
| omega | 0.017890 |
| rho | 2.216487 |
| k | 40.000000 |
| loglik | -566.404665 |

From the scatter plots, $\mu_{EI}$ against other parameters have bimodal scatters, with one group very near zero and the other taking higher values. $\rho$ has a very left skewed histogram, indicating a low reporting rate. $\omega$ is also clustered near zero which could be a sign of some difficulty with identifiability.

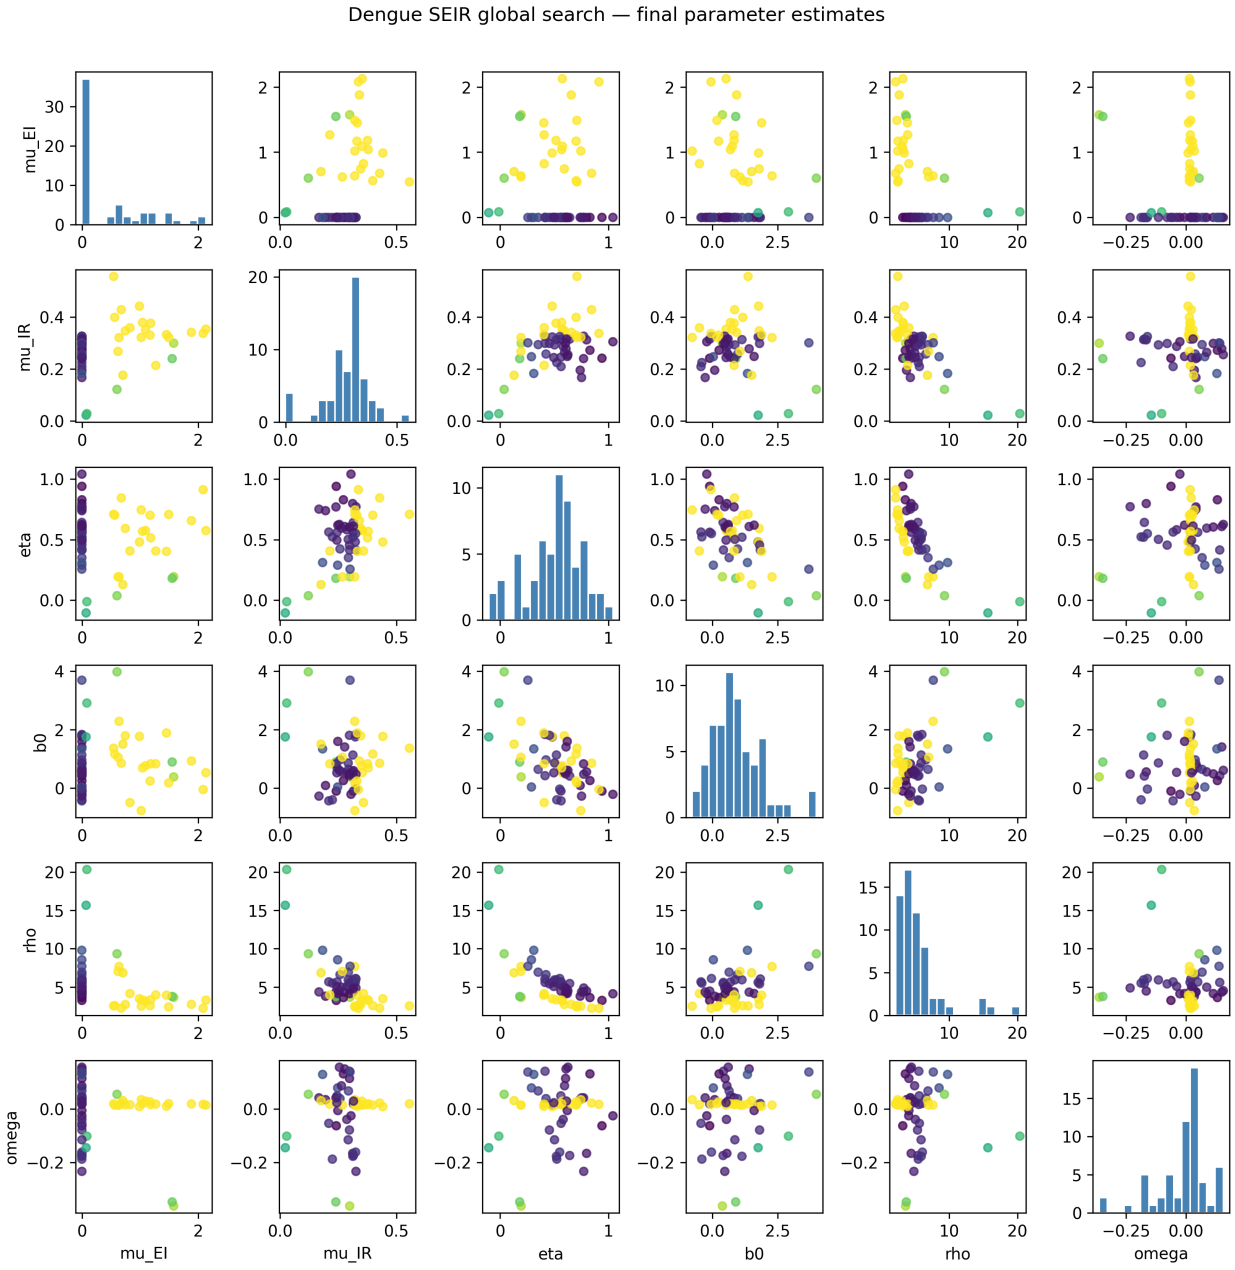


The high variance in the results of the global search still indicates some uncertainty. Our local search obtained better convergence results likely thanks to refined starting values. The higher run level improved some chain stability, although the variance in the log likelihood traces in the global search remains a challenge due to occasional small $H(t)$.

## 4.2 Profile Likelihood for $\omega$

We take a profile over $\omega$ as it is a parameter of interest in the serotype analysis. The profile likelihood plot for $\omega$ shows relatively higher log-likelihood between 0.02 and 0.04. It does not have

a strong shape, suggesting that $\omega$ is difficult to identify. The non-identifiability of $\omega$ is consistent with the results of our local and global parameter search, where it converged to small values with some variability between chains.

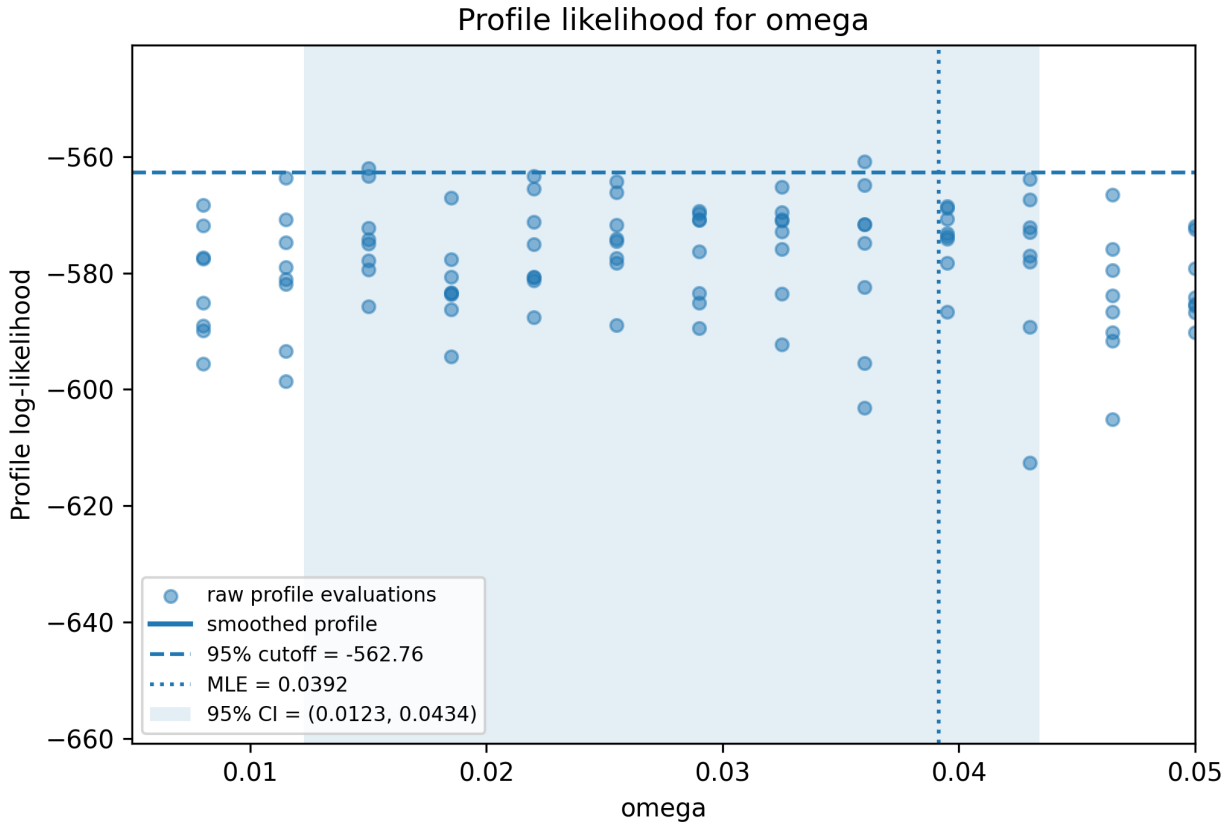


```
MCAP omega MLE: 0.03916016016016016
MCAP omega CI : (0.012281281281281281, 0.04342742742742743)
```

The profile suggests around 4% chance per week of switching serotype, or an average dominant period of 25 weeks. The lower bound of the interval corresponds to a switch every 80 weeks, and the upper bound to a switch every 23 weeks.

# 5 Model Diagnostics

## 5.1 Probes Distribution

The growth rate probe suggests that our parameter region is somewhat far from the parameter region that would actually generate the data. With the histogram centered on negative values and the target statistic being positive. The residual standard deviation probe shows that the model appears to underestimate the noise in the data. It is also right skewed compared to the true value,

although not terribly far. Epidemic peaks created by our simulations are also too large compared simulated data.

```
Probe results:
growth_rate  data=    0.0117  sim_mean=   -0.0435  p~=  0.148
residual_sd  data=    2.4394  sim_mean=    2.2591  p~=  0.460
ACF1         data=    0.9589  sim_mean=    0.9703  p~=  0.396
peak_size    data= 4235.0000  sim_mean=13221.2640  p~=  0.188
```

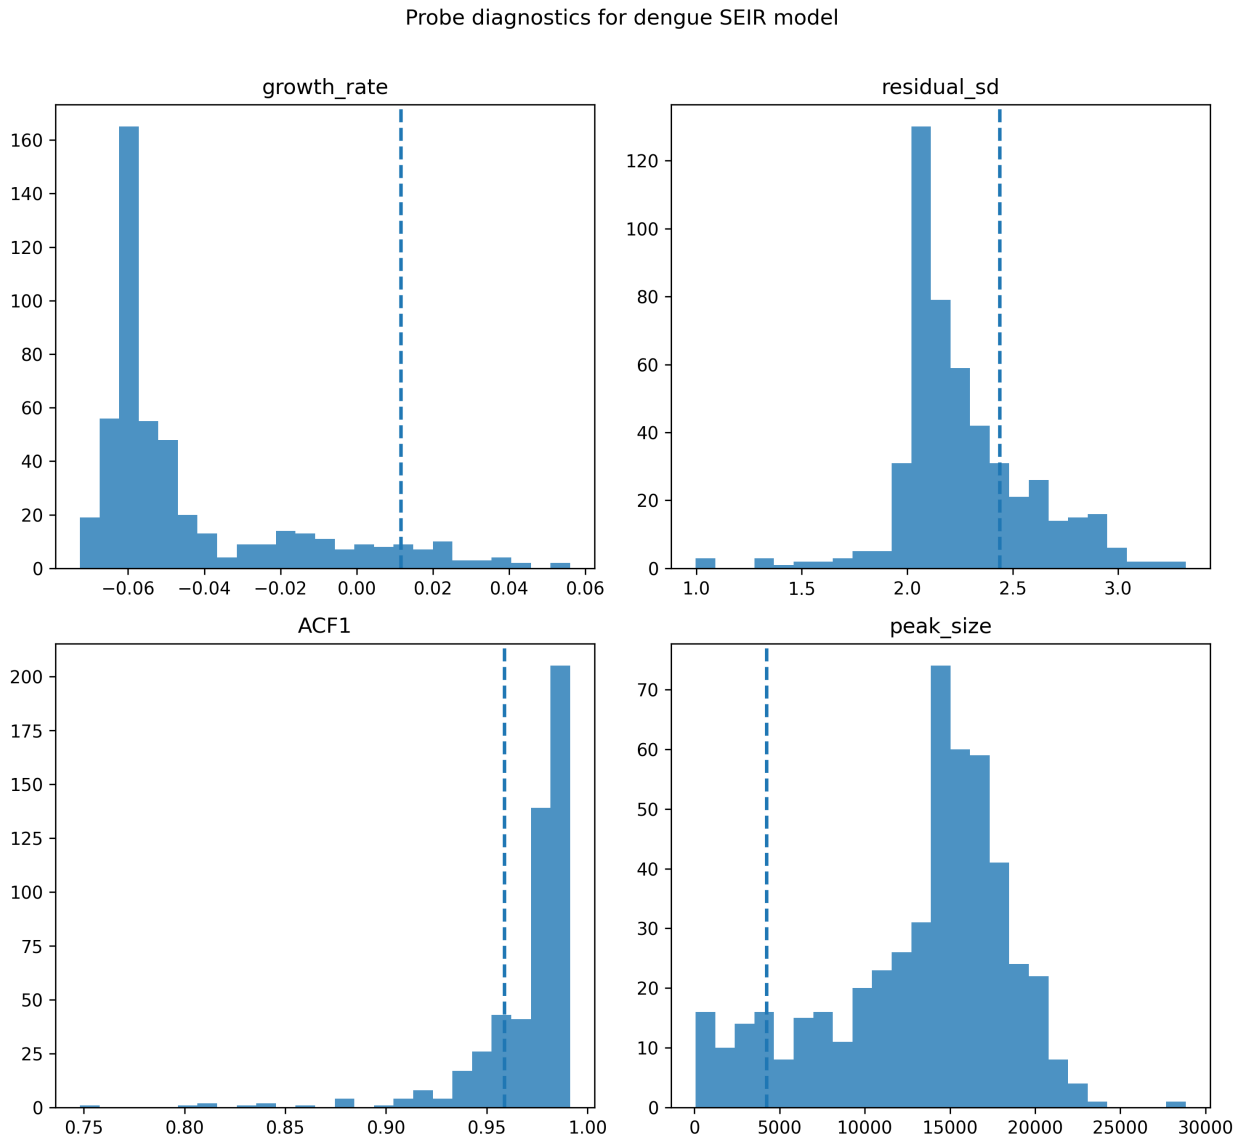

## 5.2 Comparison

We fit a number of simpler ARMA and GARCH models to compare against our SEIR particle filter. Log-likelihood results from a selection of these models (ARMA at (1,0), (1,1), (2,0), (2,1), and (2,2) and constant-GARCH(1,1), zero-GARCH(1,1), and AR(1)-GARCH(1,1)) show that our particle filter only performs slightly better than the benchmark results, although provides some improvement on the log-likelihood results from the selections. The addition of the serotype latent state in our SEIR particle filter does not appear to be as helpful in capturing the dynamics of the twin 2014 and 2015 outbreaks as we had hoped. While the GARCH models perform reasonably well without the inclusion of additional epidemiological components like serotype or population demographics in the SEIR particle filter.

Likelihood comparison:

| model | loglik | n_params | aic | bic |
|---|---|---|---|---|
| SEIR (particle filter) | -566.404665 | 20 | NaN | NaN |
| AR(1)-GARCH(1,1) on raw counts | -606.337768 | 5 | 1222.675535 | 1236.817104 |
| AR(1)-GARCH(1,1) on diff raw counts | -647.583751 | 5 | 1305.167502 | 1319.268910 |
| Zero-GARCH(1,1) on diff raw counts | -651.317169 | 3 | 1308.634338 | 1317.119280 |
| Constant-GARCH(1,1) on raw counts | -658.192198 | 4 | 1324.384395 | 1335.729523 |
| ARMA(2,2) on raw counts | -848.466523 | 6 | 1708.933046 | 1725.950737 |
| ARMA(2,1) on raw counts | -854.027986 | 5 | 1718.055972 | 1732.237382 |
| ARMA(2,0) on raw counts | -856.628976 | 4 | 1721.257953 | 1732.603080 |
| ARMA(1,1) on raw counts | -862.596297 | 4 | 1733.192595 | 1744.537723 |
| ARMA(1,0) on raw counts | -871.722768 | 3 | 1749.445537 | 1757.954383 |

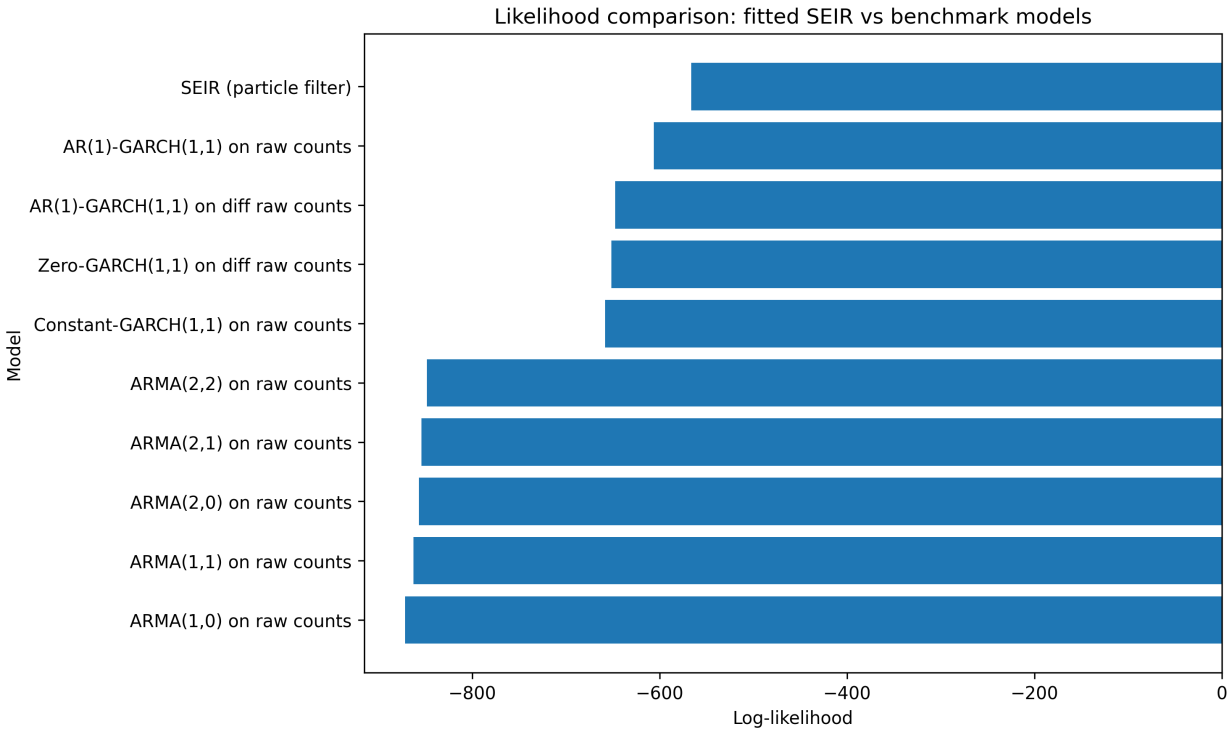


## 6 Conclusion

Our model has some significant challenges. The difficulty in the log-likelihood convergence for certain parameters during the global search suggests that further use could have been made of more profile likelihoods to confirm weak identification, or trying some kind of reformulation of the model. Another problem is the degenerate results occasionally produced by the model when $H(t) = 0$. Due to this numerical problem, we introduced a clamp on $\mu$ in our `dmeas` and `rmeas` of $\mu = \max(\rho(H(t) + 0.01I(t)), 0.01)$ where 0.01 is a floor for numerical stability. This guarantees that an observation stays well defined even in weeks where $H(t) = 0$, but at the cost of approximation to the negative binomial likelihood in lower transmission periods. This construction could be creating unreliable log-likelihood values.

This SEIR model was unable to model the data particularly well. The parameters given to us by the local and global parameters searches did not reflect the magnitude of the epidemic peaks well.

## References Statement

ChatGPT and Claude were used in the coding process. No other sources beyond those mentioned were used.

## Bibliography

Dinkar, Anju, and Jitendra Singh. 2026. "Dengue Serotypes and Their Severity Correlation: A Hospital-Based Observational Study." *Enfermedades Infecciosas y Microbiología Clínica (En-*

*glish Edition)* 44 (2). https://doi.org/10.1016/j.eimce.2026.503074.

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

Ligon, B. Lee. 2005. "Dengue Fever and Dengue Hemorrhagic Fever: A Review of the History, Transmission, Treatment, and Prevention." *Seminars in Pediatric Infectious Diseases* 16 (1): 60–65. https://doi.org/10.1053/j.spid.2004.09.013.

Strickler, Robin. 2018. "Mosquito-to-Mosquito Infections Keep Dengue Circulating." Cornell Chronicle. https://news.cornell.edu/stories/2018/10/mosquito-mosquito-infections-keep-dengue-circulating.

Taiwan Centers for Disease Control. 2024. "Dengue Fever." Centers for Disease Control, R.O.C. (Taiwan); https://www.cdc.gov.tw/En/Category/ListContent/bg0g_VU_Ysrgkes_KRUDgQ?uaid=9_Oq7OYHa-l8B05iUwyVvQ.

taweilo. 2024. "Taiwan Dengue Daily Confirmed Cases 1998–2024." Kaggle. https://www.kaggle.com/datasets/taweilo/taiwan-dengue-daily-confirmed-cases-1998-2024.

United Nations, Department of Economic and Social Affairs, Population Division. 2024. "World Population Prospects 2024." Standard Projections. United Nations. https://population.un.org/wpp/.

Wei, Hsin-Yi, Pei-Yun Shu, and Min-Nan Hung. 2016. "Characteristics and Risk Factors for Fatality in Patients with Dengue Hemorrhagic Fever, Taiwan, 2014." *The American Journal of Tropical Medicine and Hygiene* 95 (2): 322–27. https://doi.org/10.4269/ajtmh.15-0905.