Stochastic Epidemic Modelling

Evie Miller[1] **James Neill**[1]
[1]STOR-i Centre for Doctoral Training, Lancaster University

Introduction

Stochastic epidemic modelling uses statistical models to understand how infectious pathogens might spread within a population.

The standard stochastic epidemic model (the SIR model) and extensions of it are used to simulate an outbreak. If we have infection time data, Markov chain Monte Carlo (MCMC) methods can be applied to infer the rates of infection and recovery.

As part of the TRACS-Liverpool project (an initiative focused on tracking antimicrobial resistance in care settings across Liverpool), a model has been designed for various hospital and care settings, specifically shown here for Aintree ward 17B.

SIR Model

The population is divided into three categories:

- **Susceptible (S)** – Individuals who are not yet infected.
- **Infective (I)** – Individuals who are currently infected with the pathogen and can transmit it to others.
- **Recovered (R)** – Individuals previously infected with the pathogen, but no longer transmit it, and are now immune.

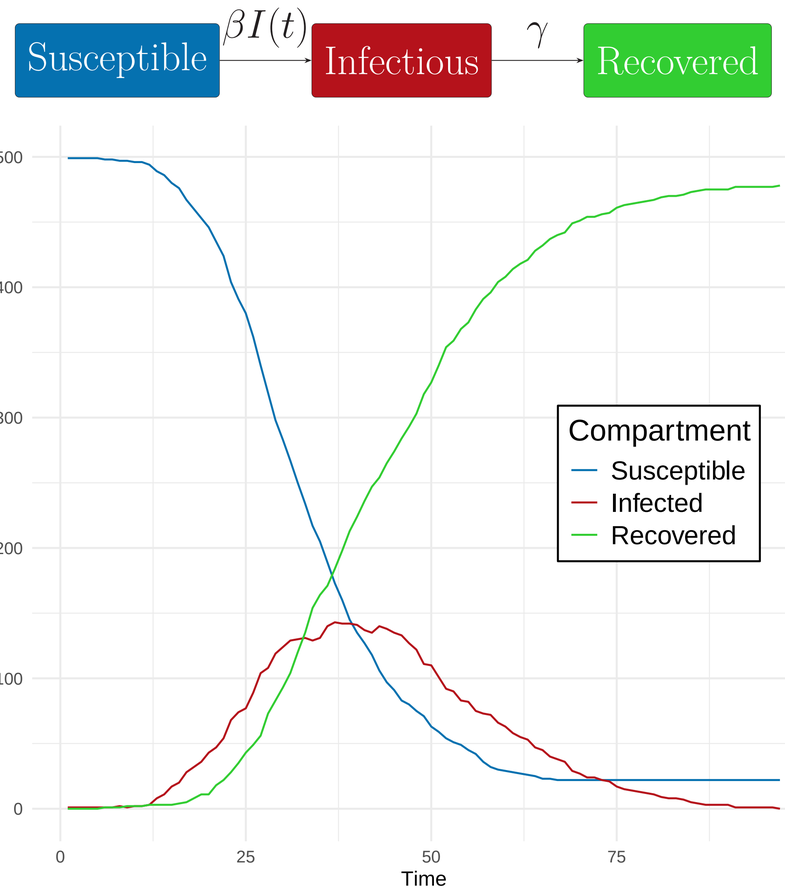

Stochastic SIR Model with Individual Tracking

Infection Rate (β): Determines how easily the disease spreads. The probability of a susceptible individual getting infected at time t is given by: $\mathrm{P}(\text{Infection}) = 1 - \exp(-\beta \times I[t-1])$.

Recovery Rate (γ): Determines recovery time of infected individuals. The constant probability of recovery: $\mathrm{P}(\text{Recovery}) = 1 - \exp(-\gamma)$

At each time step, the model updates the state of each individual. Binomial probabilities randomly determine if they become infected or recover, and the dynamics of the population over time is produced.

Extensions of the SIR Model - SEIRS

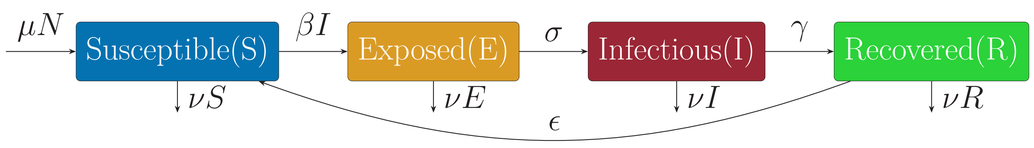

The SEIRS model accounts for susceptibility, exposure, infection, recovery, and the potential for immunity loss, along with population changes due to births and deaths. Transitions between stages are governed by the rates outlined above.

Inference with Random Walk Metropolis MCMC

To estimate β and γ, Random Walk Metropolis is applied:

- **Proposal:** New β and γ values proposed from normal distributions centered at current values, with standard deviations λ_β & λ_γ.
- **Acceptance:** A proposed value is accepted based on:

$$\log(\alpha_\theta) = \log \Pi(\theta_{\text{proposed}}|X) - \log \Pi(\theta_{\text{current}}|X)$$

where θ is β or γ, and $\Pi(\theta|X)$ is the posterior distribution of θ.

- **Adaptive Tuning:** Every 100 iterations, proposed standard deviations λ_β & λ_γ are adjusted to maintain a 30% acceptance rate.
- **Trace Plots and Histograms:** The evolution of β and γ over 2000 iterations is visualised using trace plots, and their distributions are displayed with histograms.

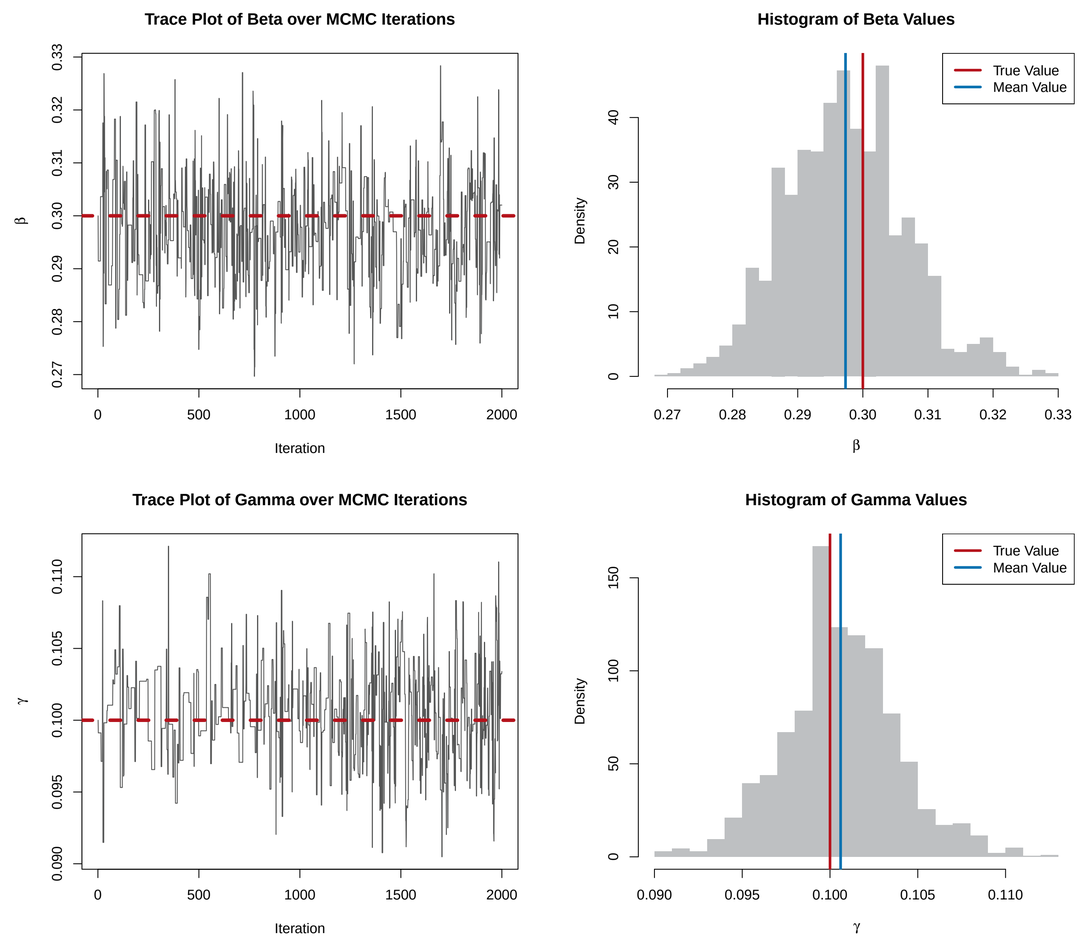

The Random Walk Metropolis MCMC efficiently estimates stochastic SIR model parameters by exploring the parameter space and adaptively tuning the proposal distribution. This MCMC is conditional on full knowledge of infection and recovery times.

Hospital Ward Simulation

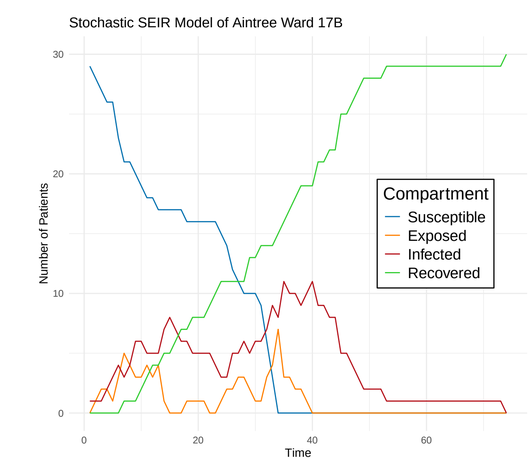

SEIR Model of Aintree ward 17B assuming 20 patients, room allocation is random.

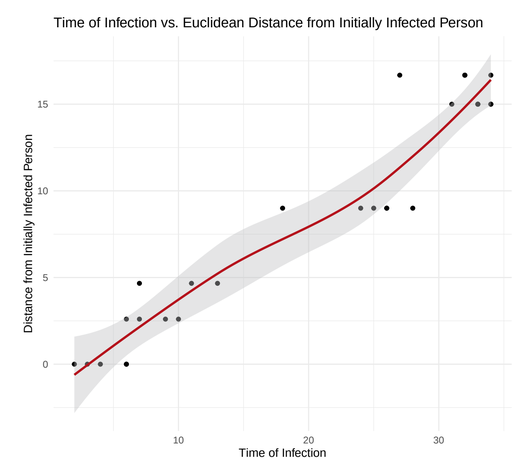

Time of infection against Euclidean distance from room of initially infected individual.

The model for Aintree ward 17B incorporates the distance between rooms using an inverse square law to simulate infection spread between individuals. To handle cases where the distance is zero, the infection rate, β, is set higher at **0.3** for patients in the **same room** compared to **0.1** for those in **different rooms**. As shown in the distance-time plot, the model behaves as expected, with patients further apart becoming infected later. This suggests that the model can simulate an epidemic in a ward like this when real data is available.

e.miller2@lancaster.ac.uk