A New Approach for Parameter Estimation in Complex Epidemiological Models

Chris Pooley (BioSS), Andrea Doeschl-Wilson (Roslin) and Glenn Marion (BioSS)

Epidemic Models: Insights from the RAMP Project
March 24th 2022

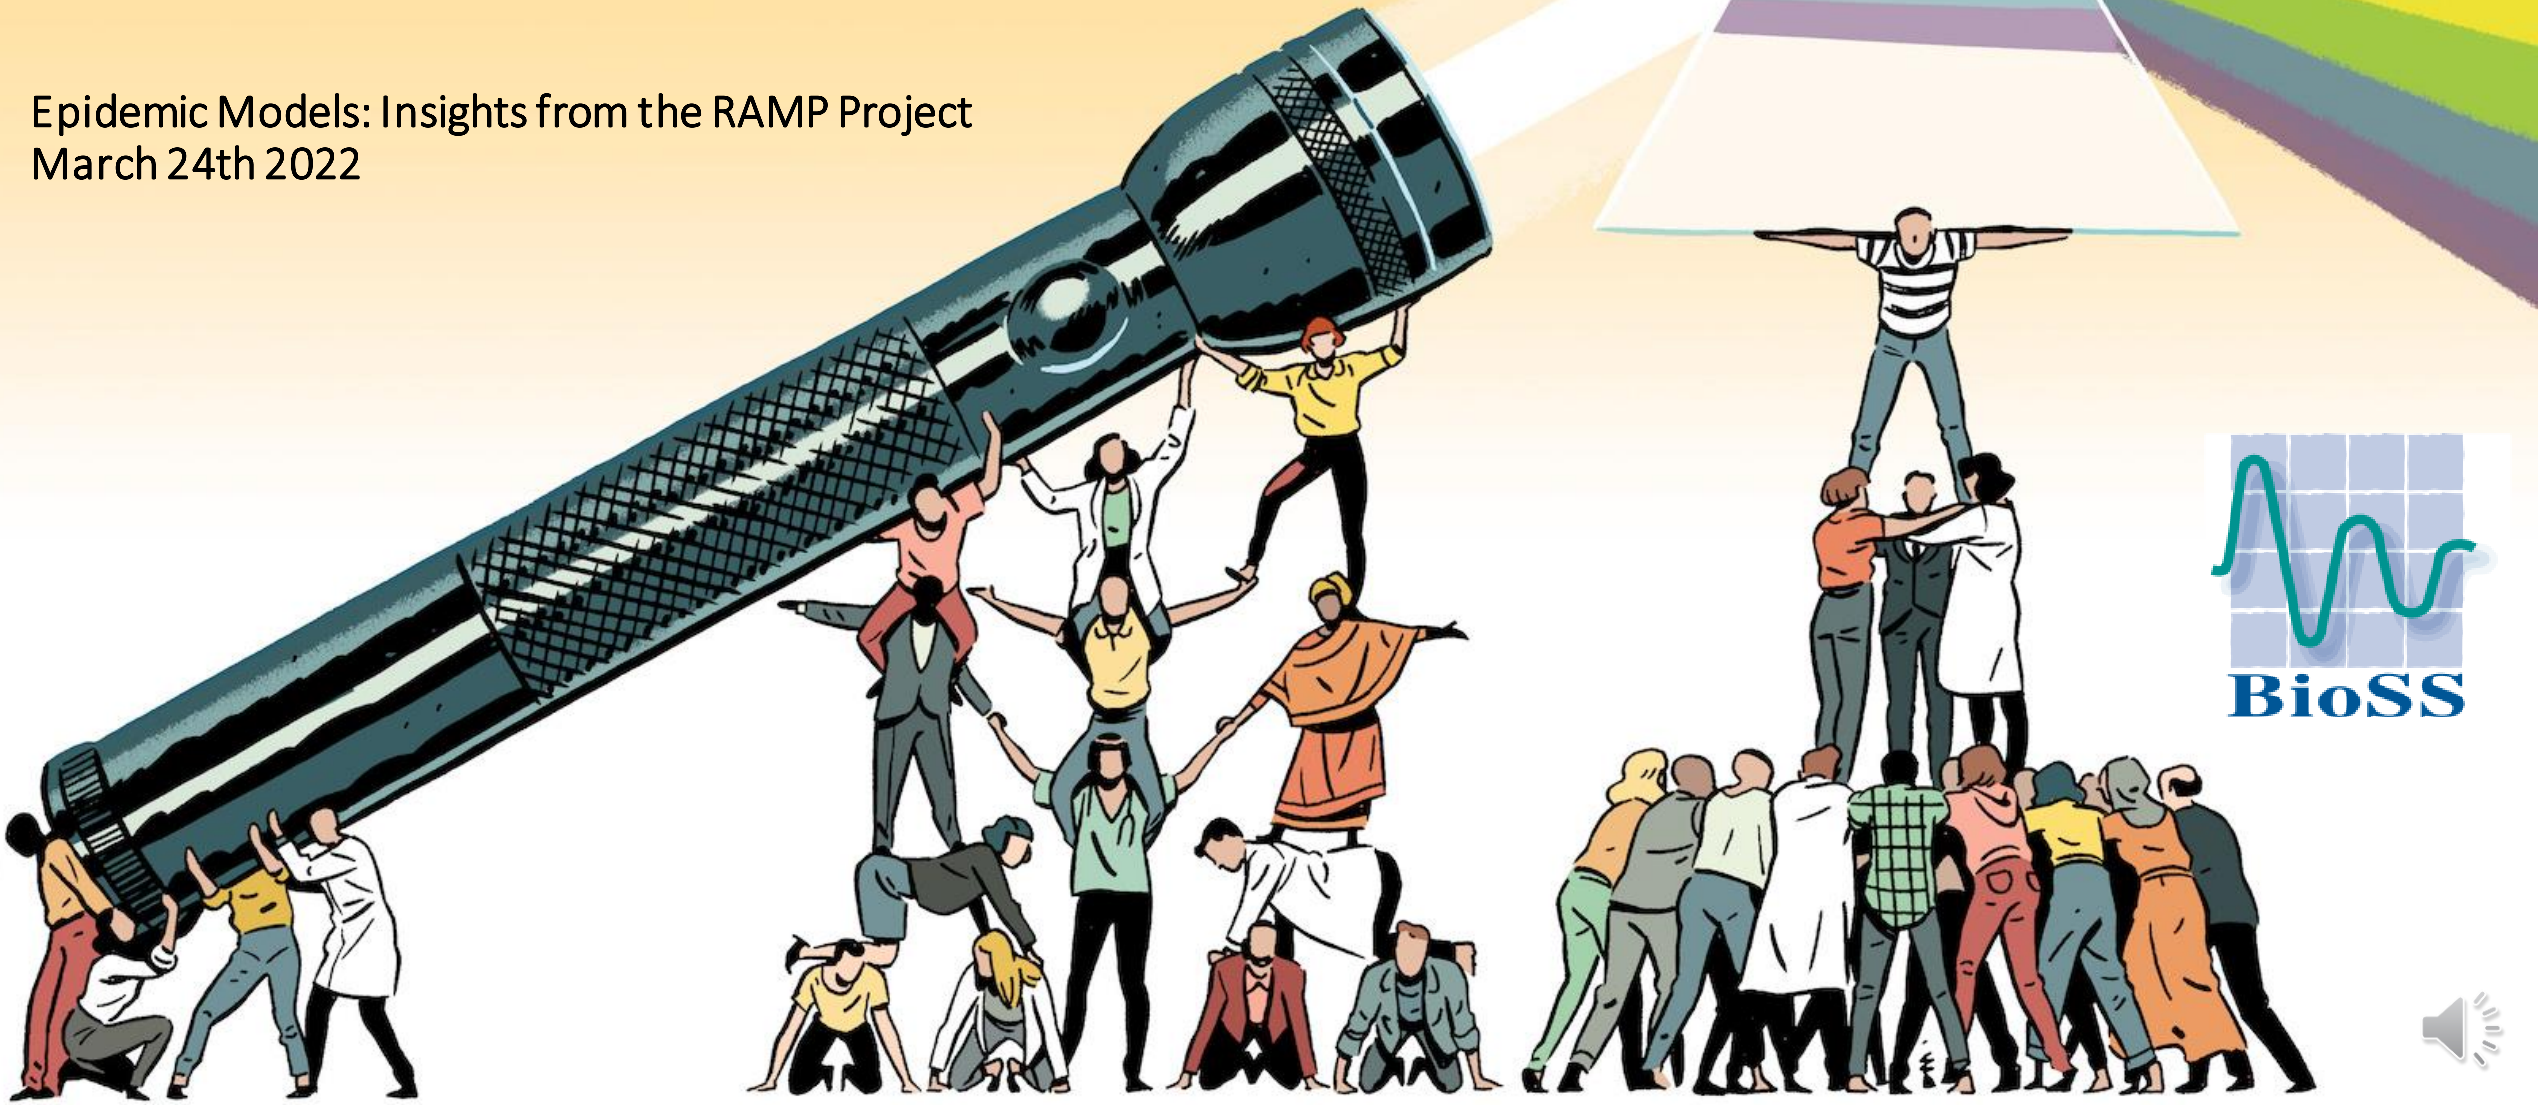

BioSS

My background

- Physics → Genetics → Epidemiology

- Understanding genetics of disease spread in animals
  - Selective breeding programs to help combat disease

- Bayesian methods for disease transmission experiments
  - Typically data from 100s of animals/fish
  - Individual-based models

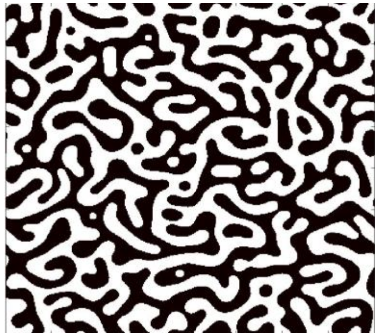

Motivation

- Develop mathematical models for COVID-19
  - Understanding
  - Looking into the future
  - Test control strategies before applying them in reality

- Models only useful if well parameterised

- Bayesian inference
  - Accounts for uncertainty in prior and data
  - Output model parameters and system behaviour INCLUDING uncertainty
  - Provides best evidence-based advice for policy

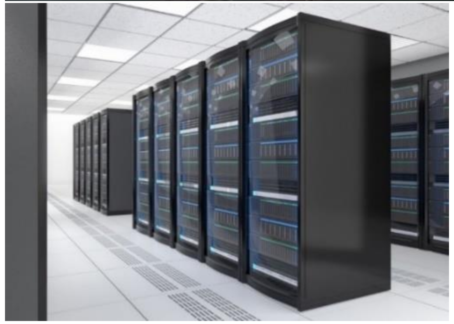


Aim

- Develop models and Bayesian methods to analyse publically available COVID-19 data

- Ideally account for:
  - Differences in disease severity
  - Spatial variation
  - Demographic variation (*e.g*. age, sex)
  - Time variation in external force of infection
  - Government interventions
  - Human behavioural change
  - Different virus variants
  - Vaccination effects
  - Biases in the data

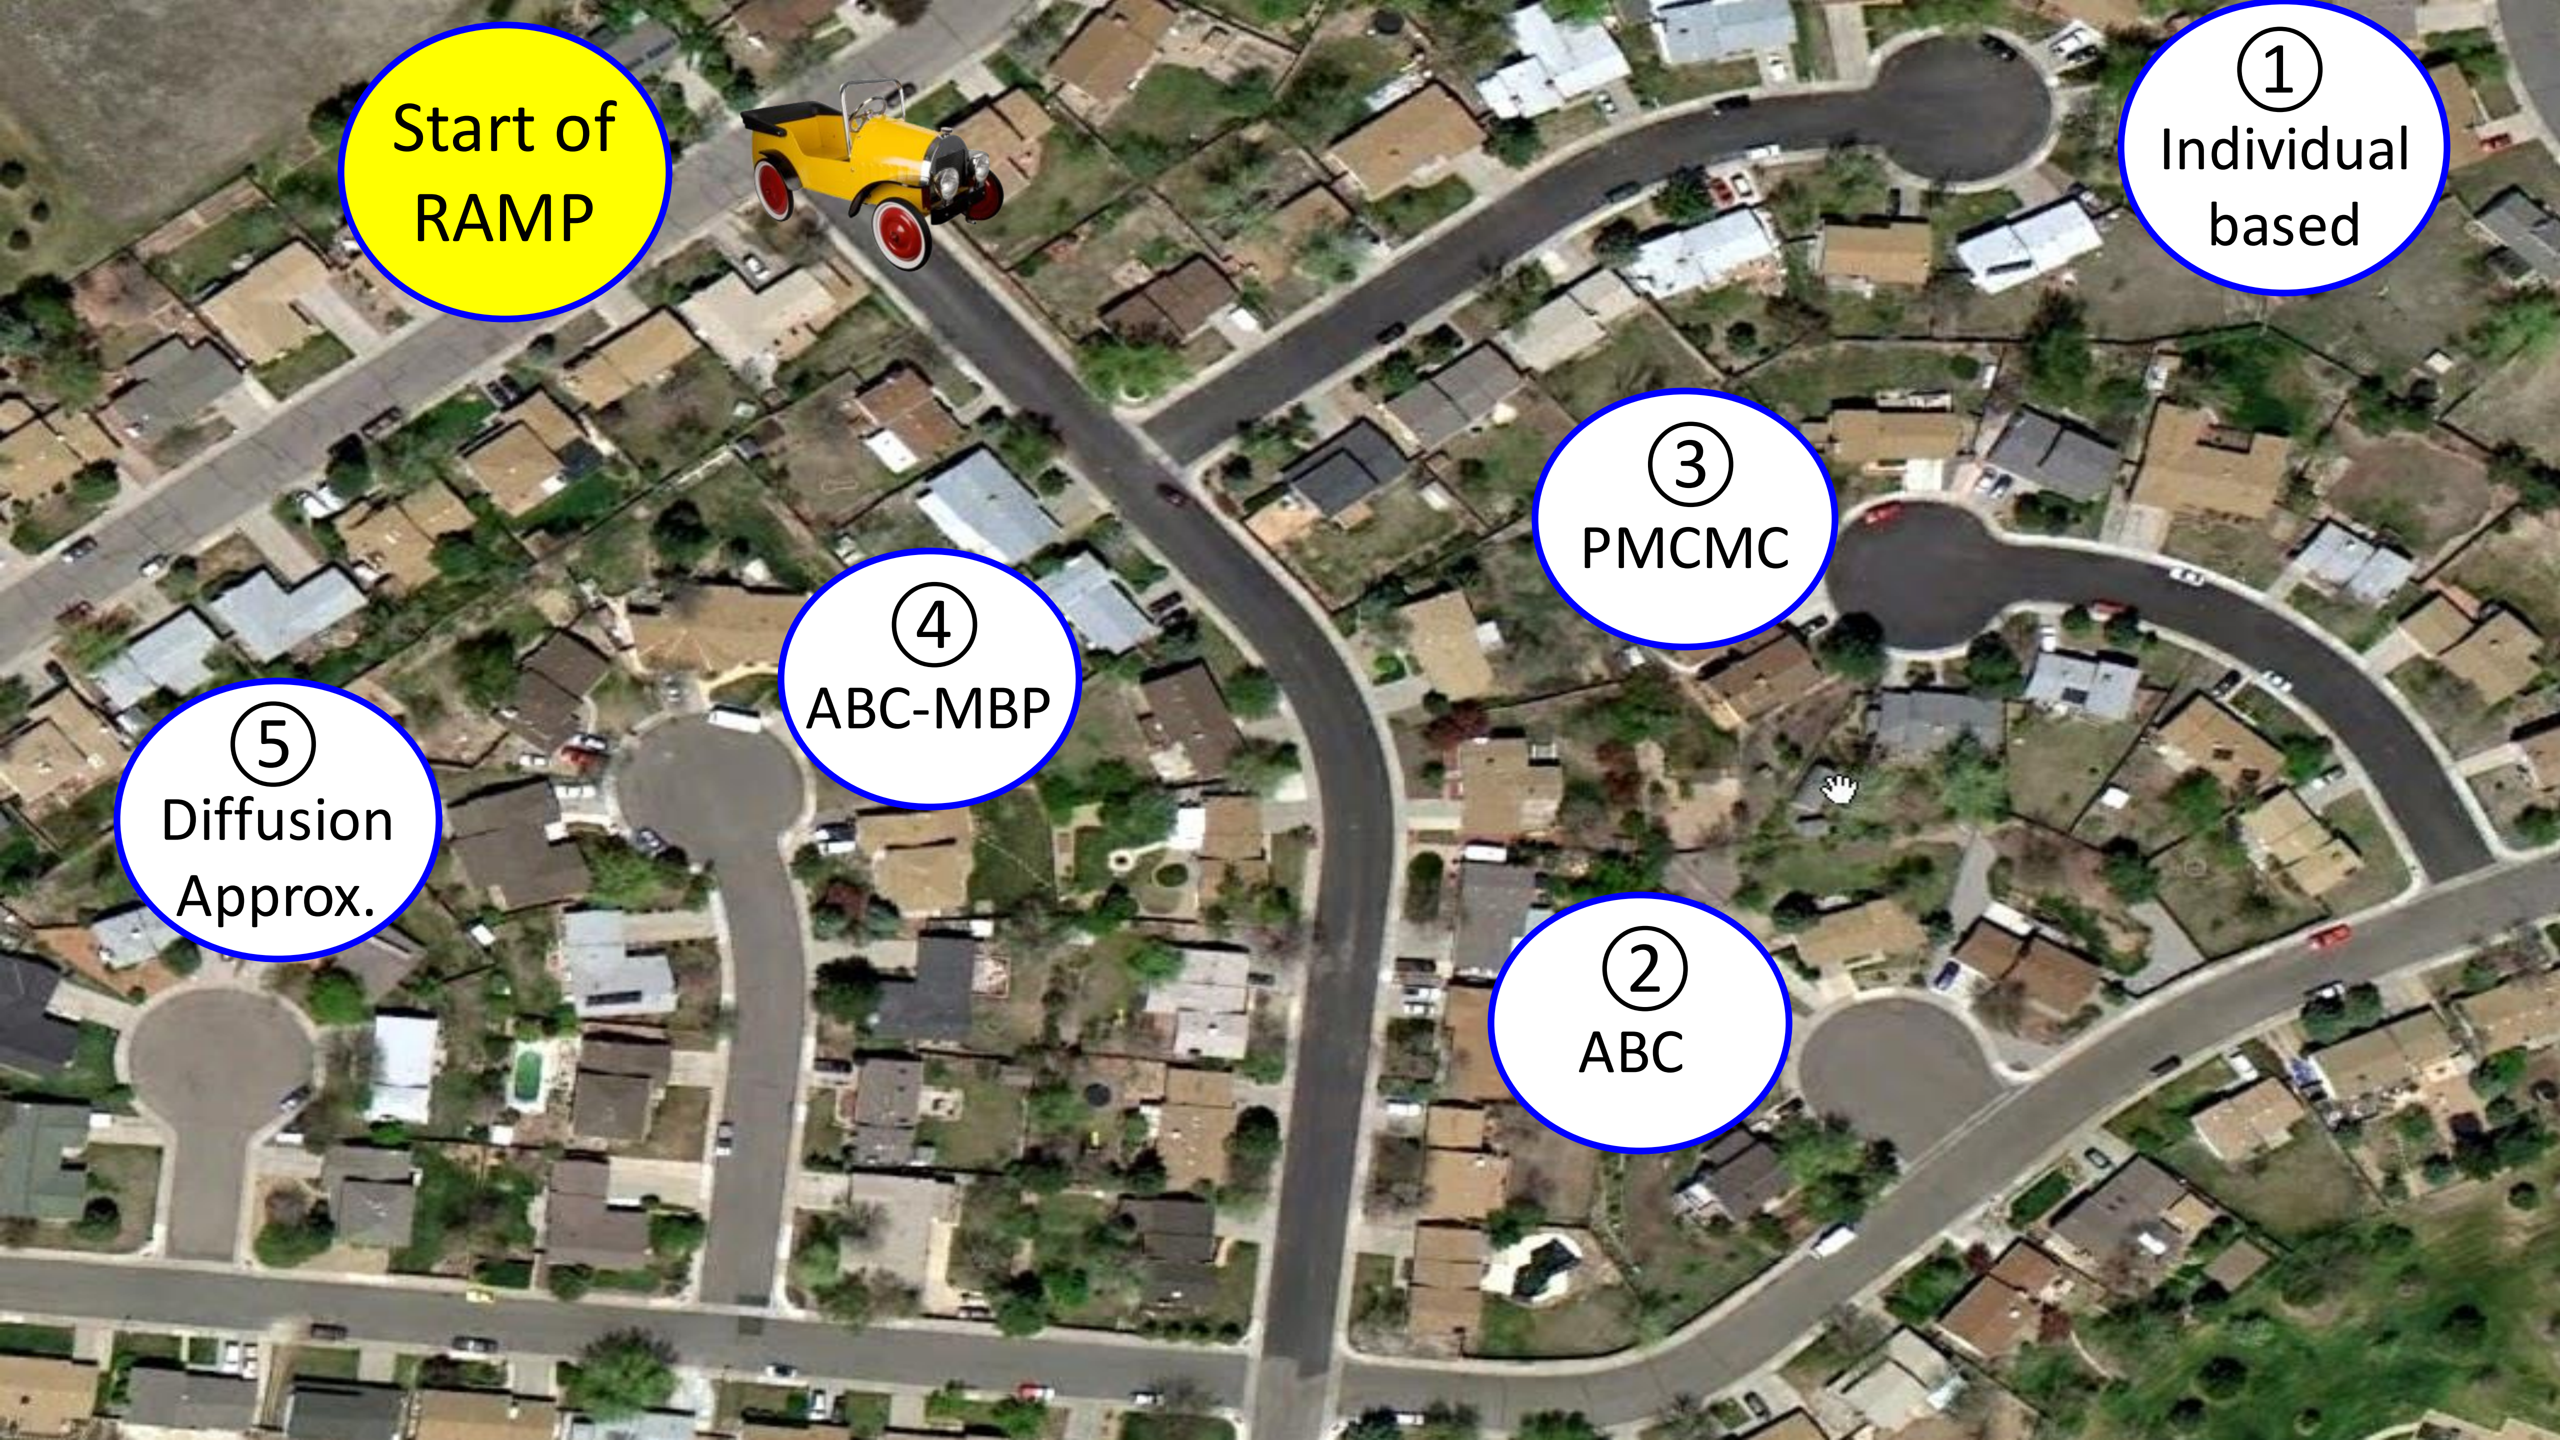
Start of RAMP
① Individual based
③ PMCMC
④ ABC-MBP
⑤ Diffusion Approx.
② ABC

Overview

- Outline existing methods:
  - ABC rejection sampling and ABC-SMC
- Introduce new inference methodology:
  - ABC-MBP
- BEEPmbp software
  - Speed comparison between different approaches
- Application to age-structured model of Covid-19

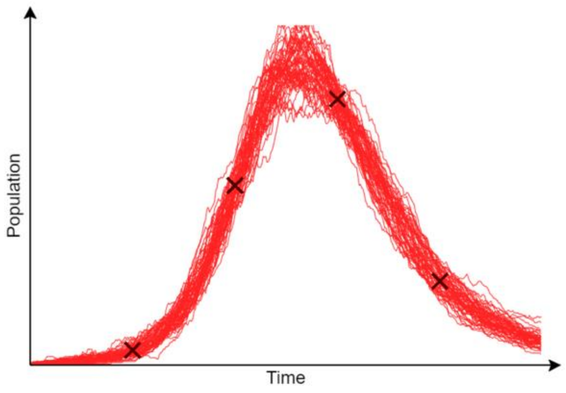

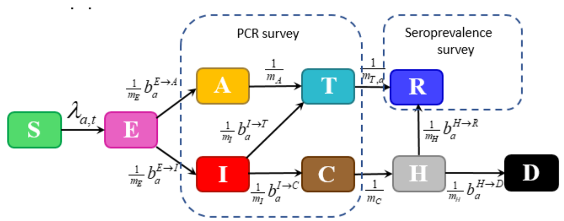

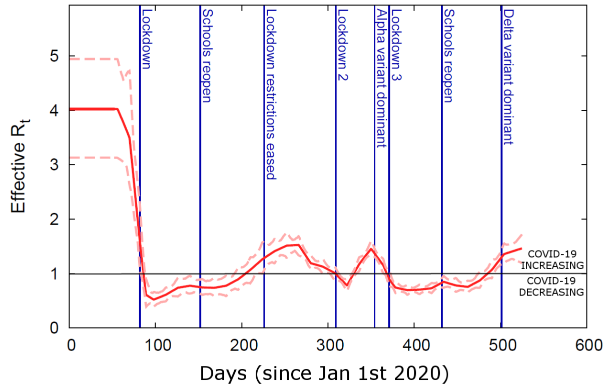

SIR compartmental model

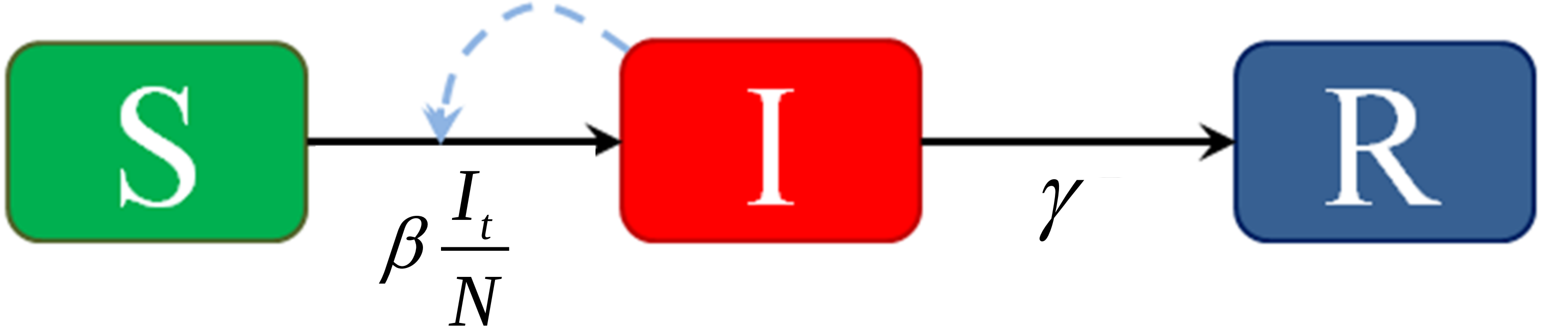

Susceptible population

Infected population

Recovered population

Simulation

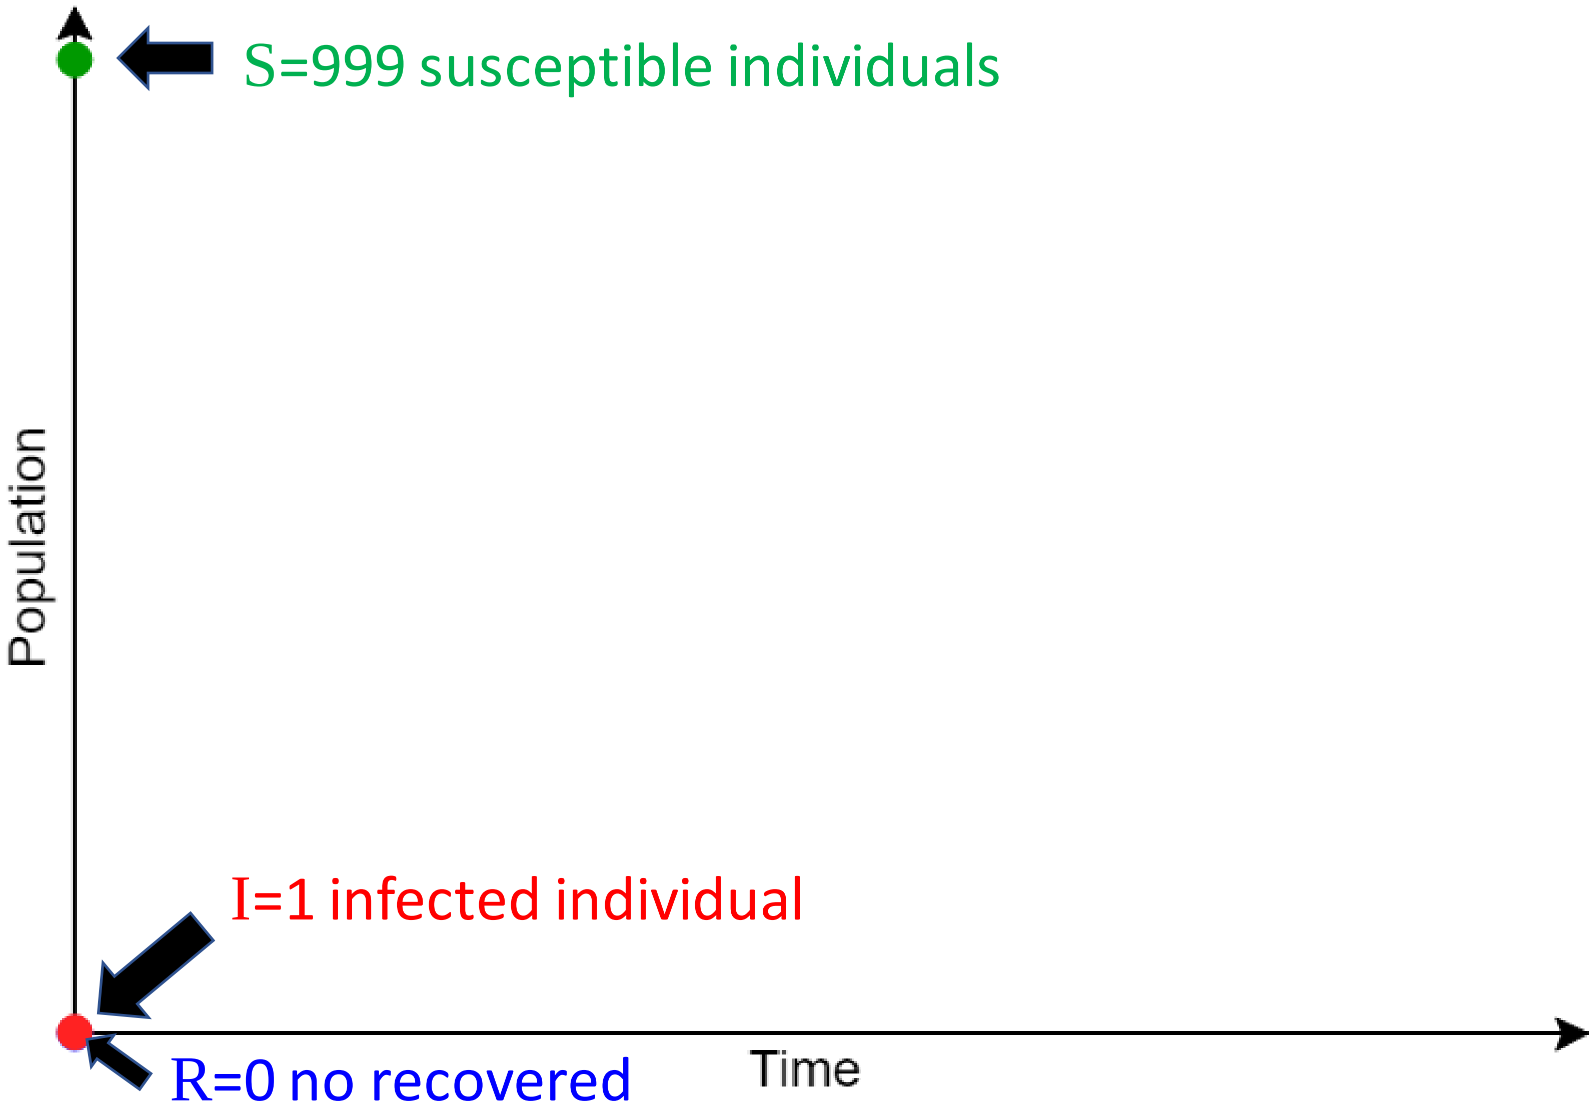
S=999 susceptible individuals
Population
I=1 infected individual
R=0 no recovered
Time

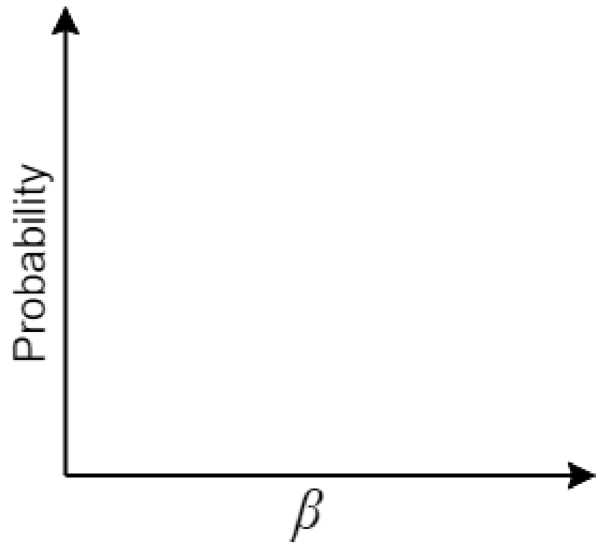
Probability
β

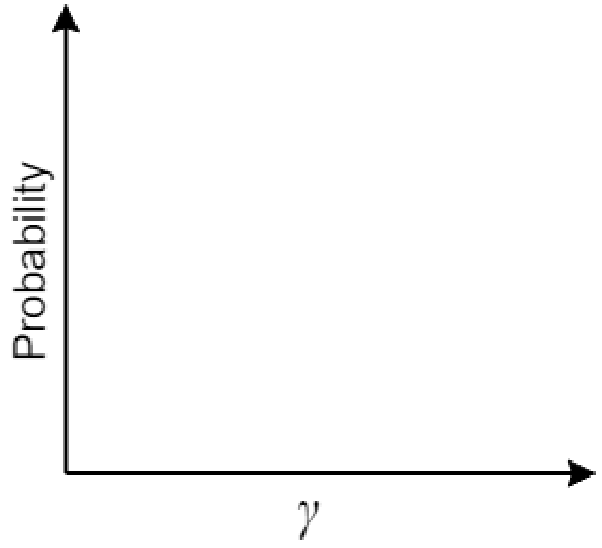
Probability
γ

Simulation

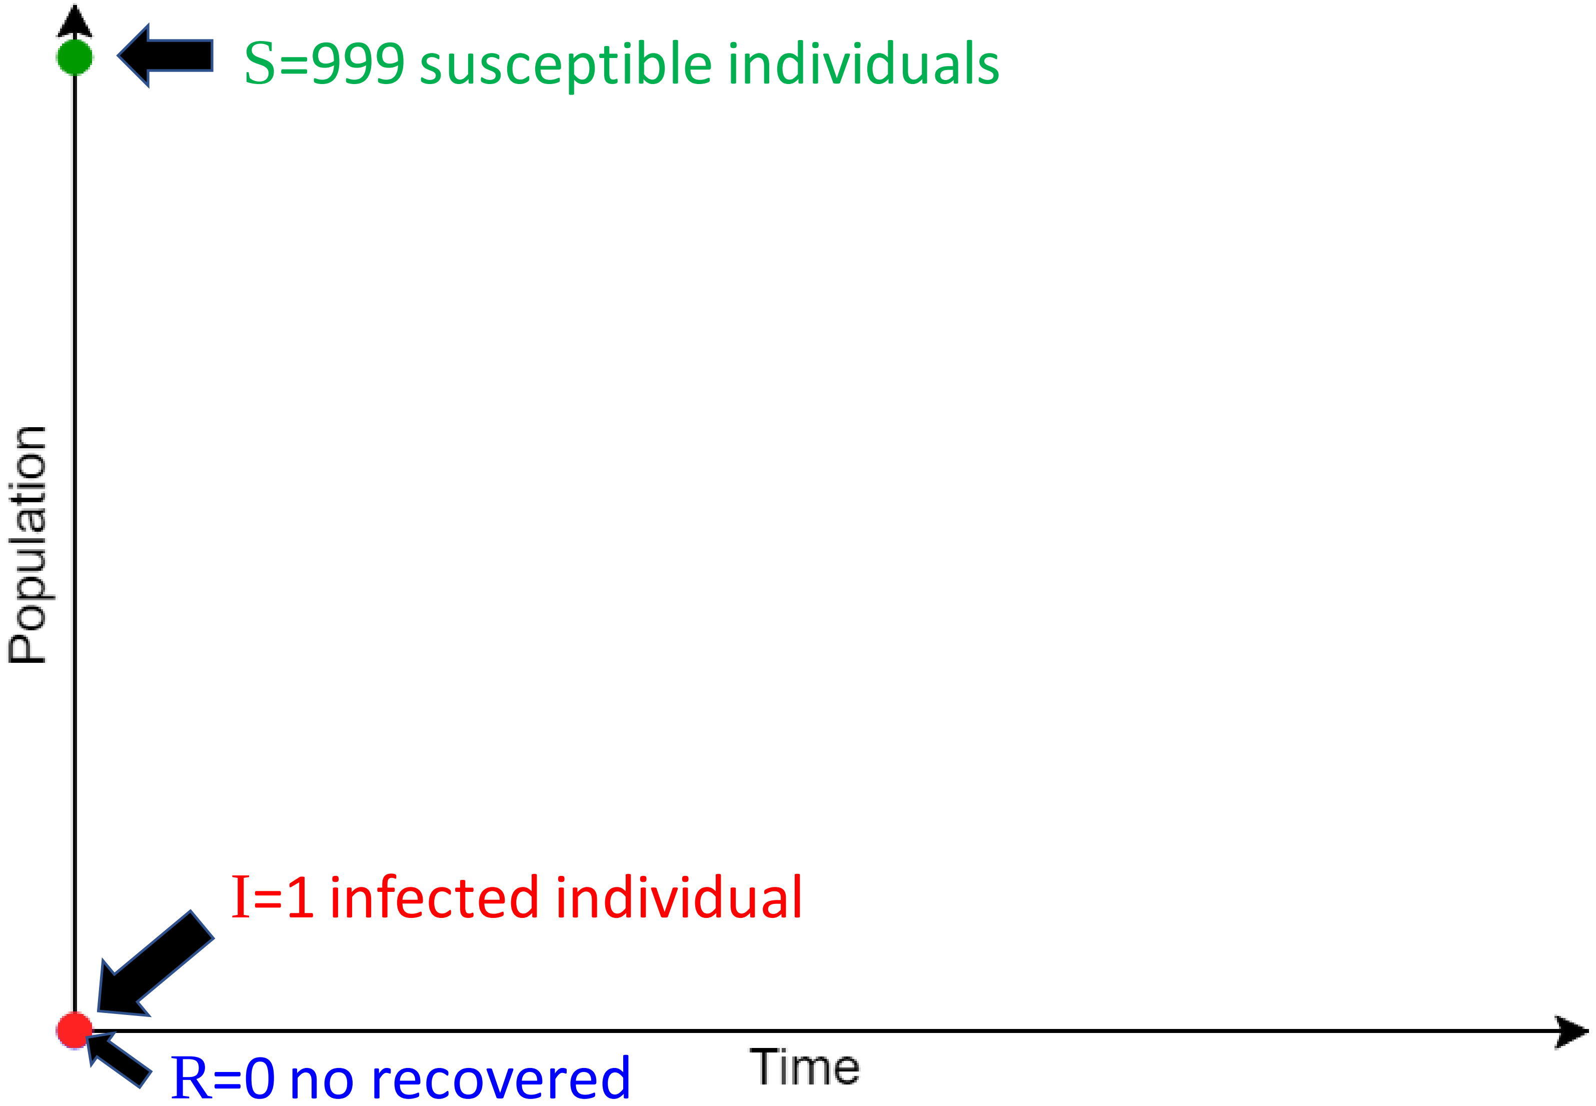

We select a set of suitable model parameters

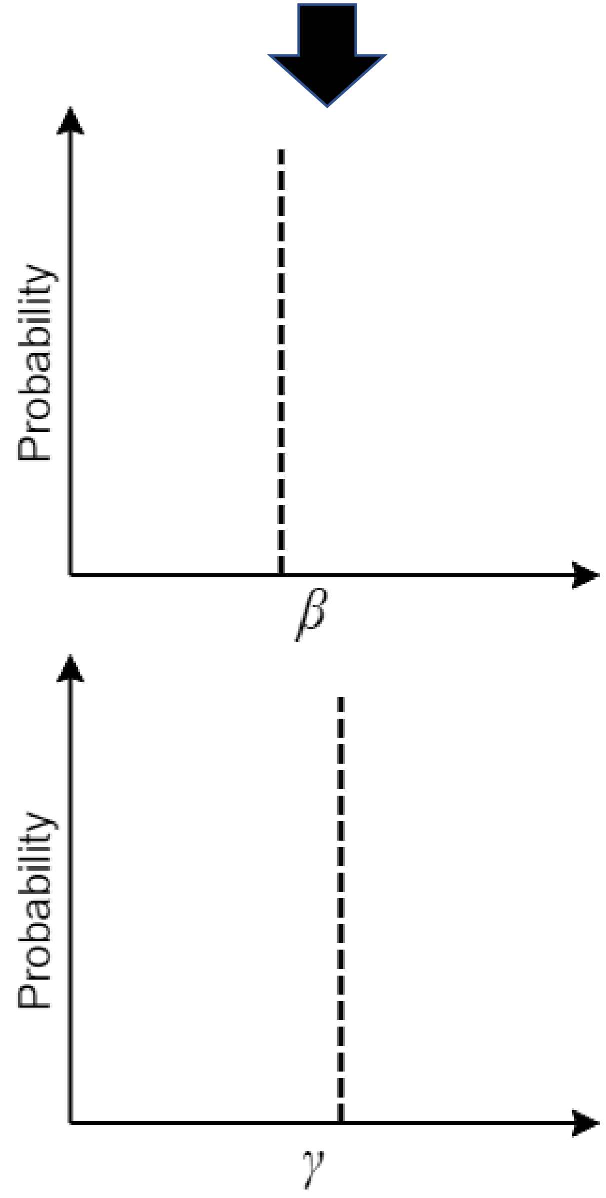

Simulation

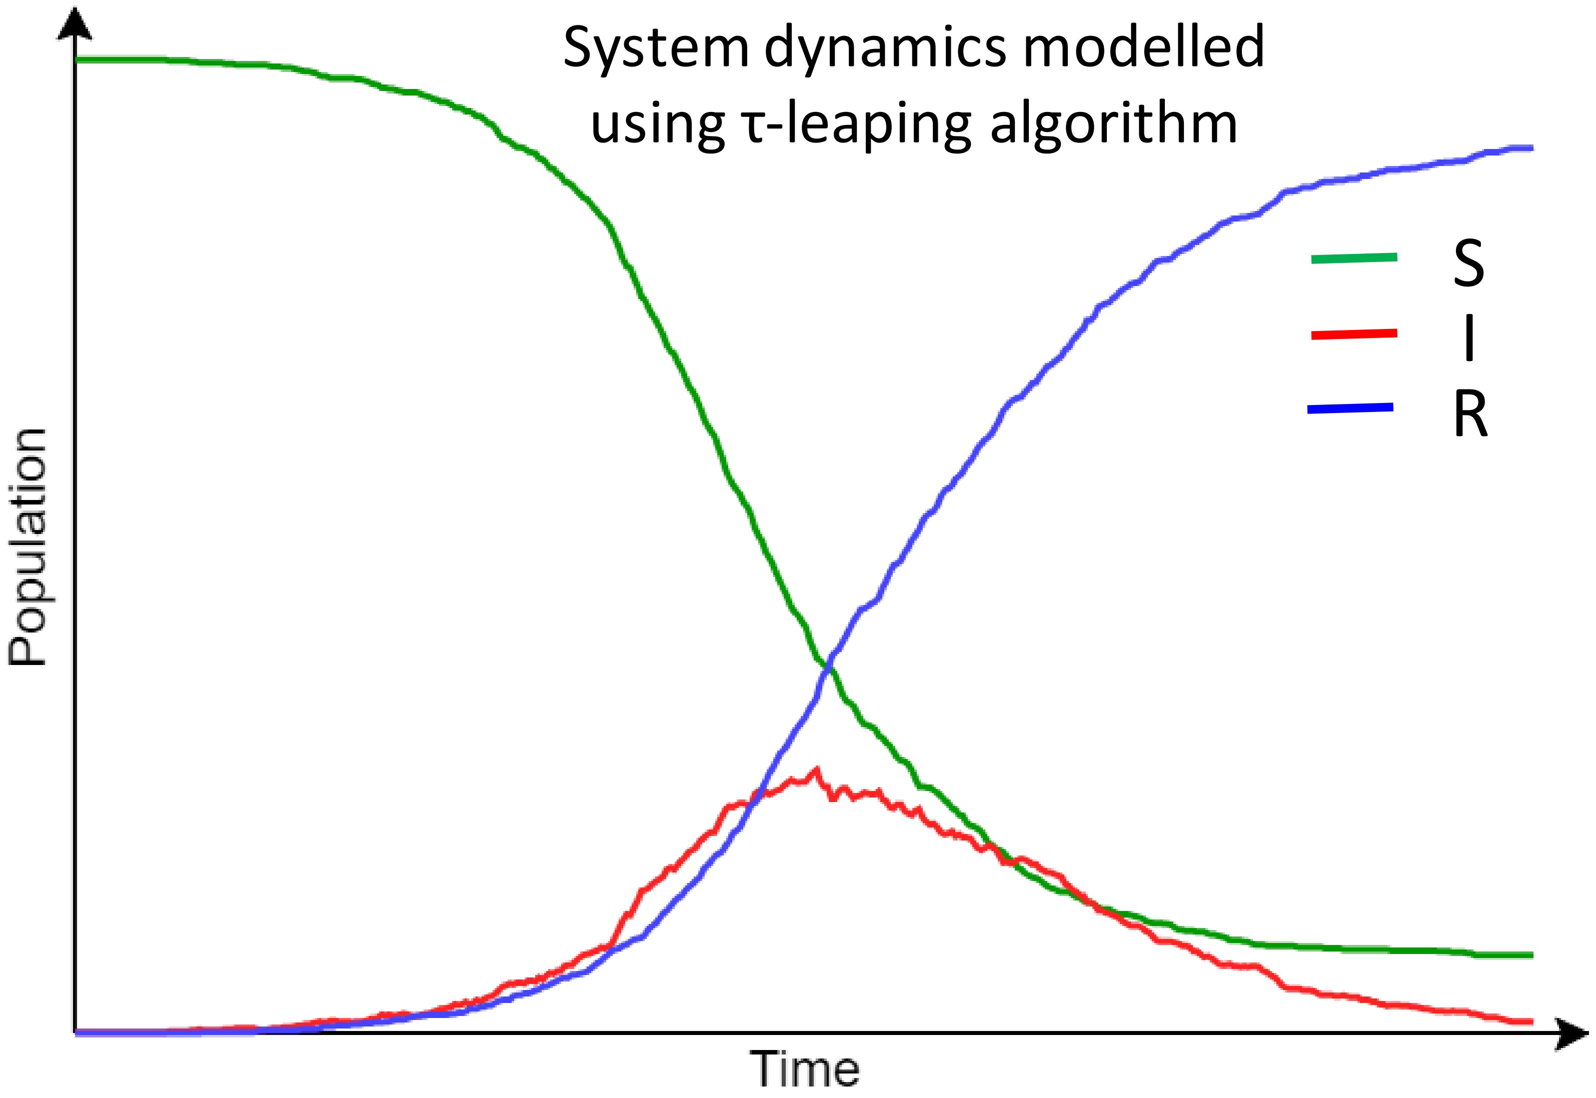
System dynamics modelled
using τ-leaping algorithm
S
I
R
Population
Time

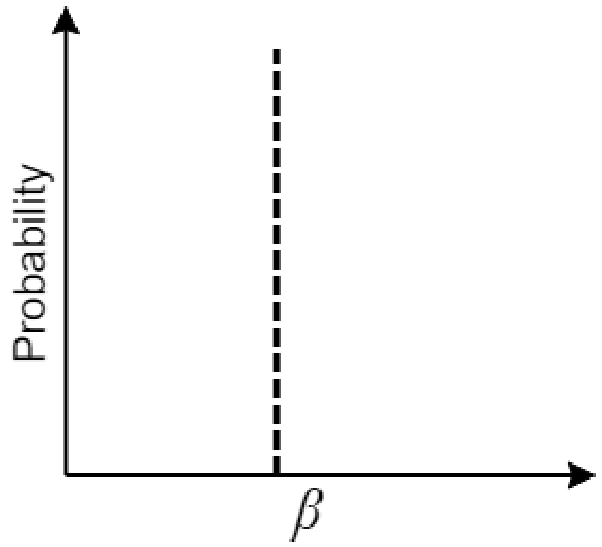
Probability
β

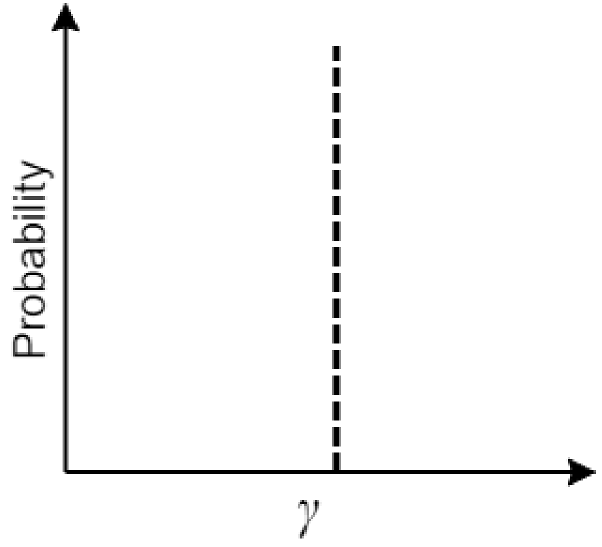
Probability
γ

Simulation

System exhibits stochastic variation

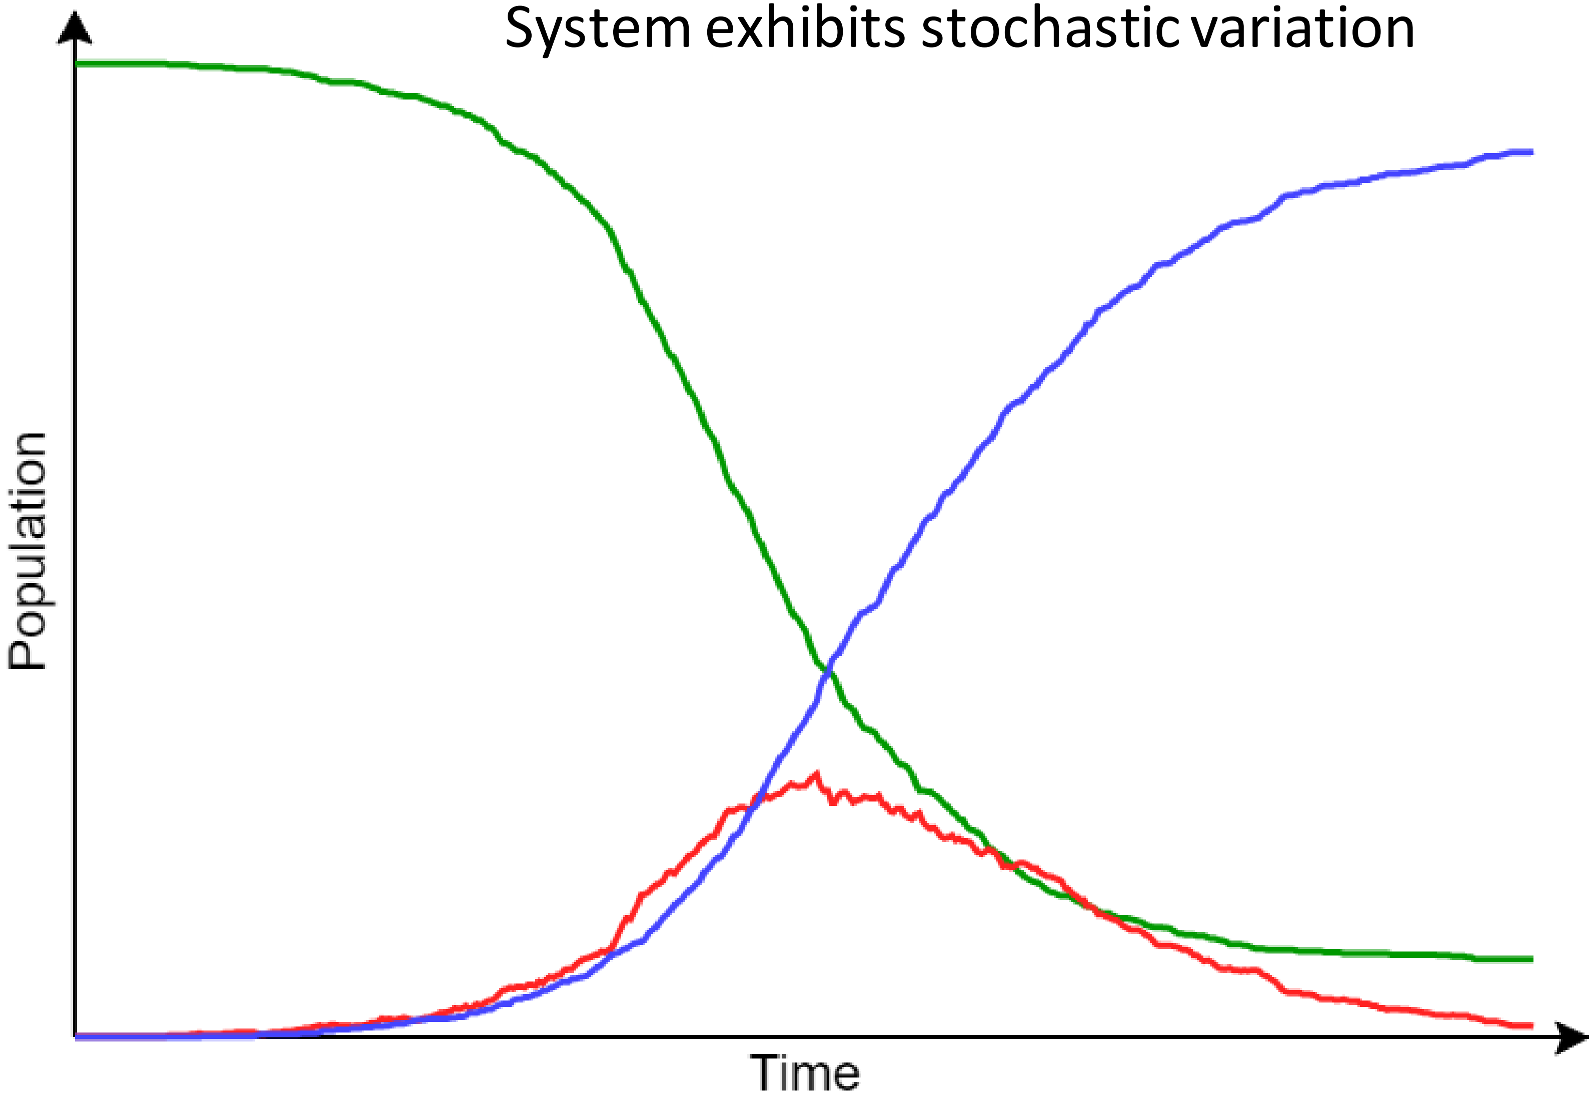

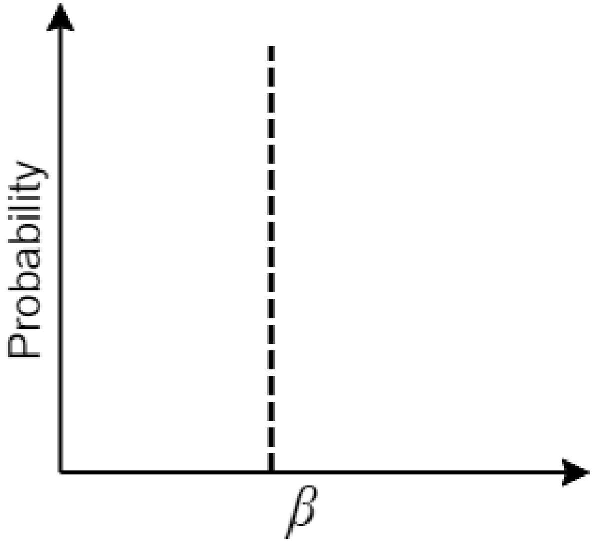

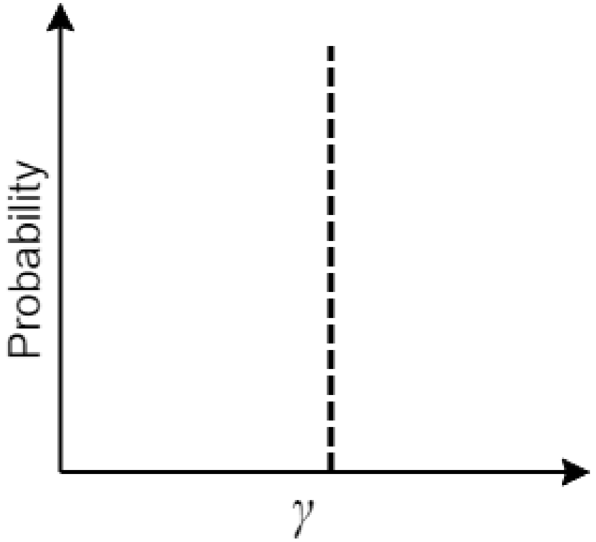

Simulation

System exhibits stochastic variation

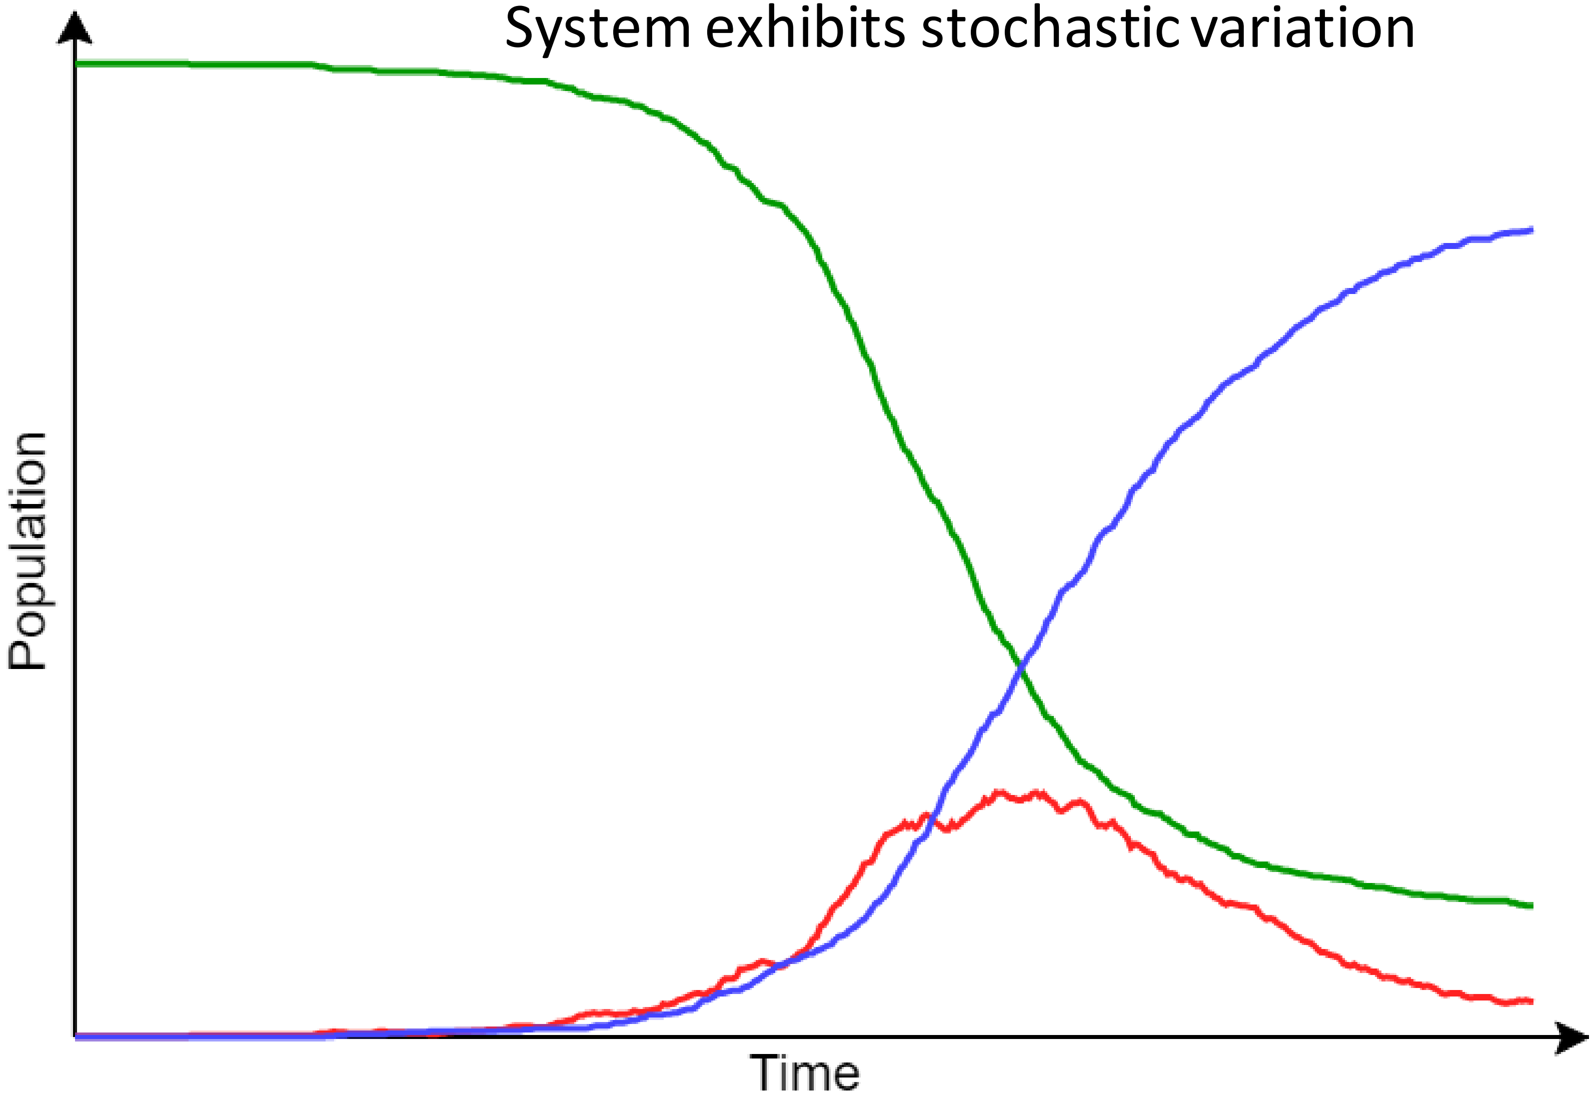

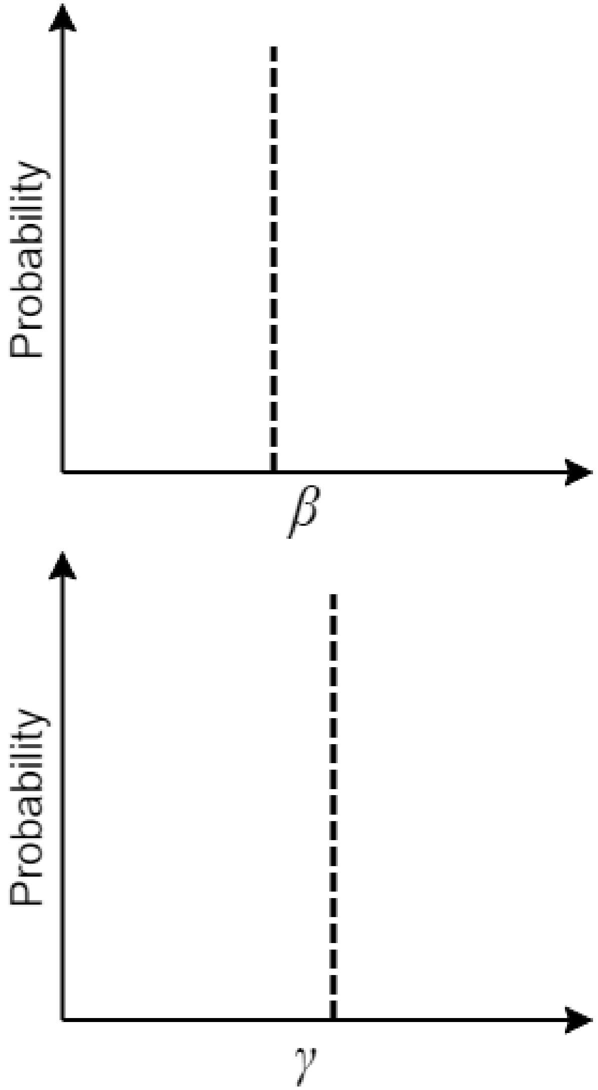

Simulation

System exhibits stochastic variation

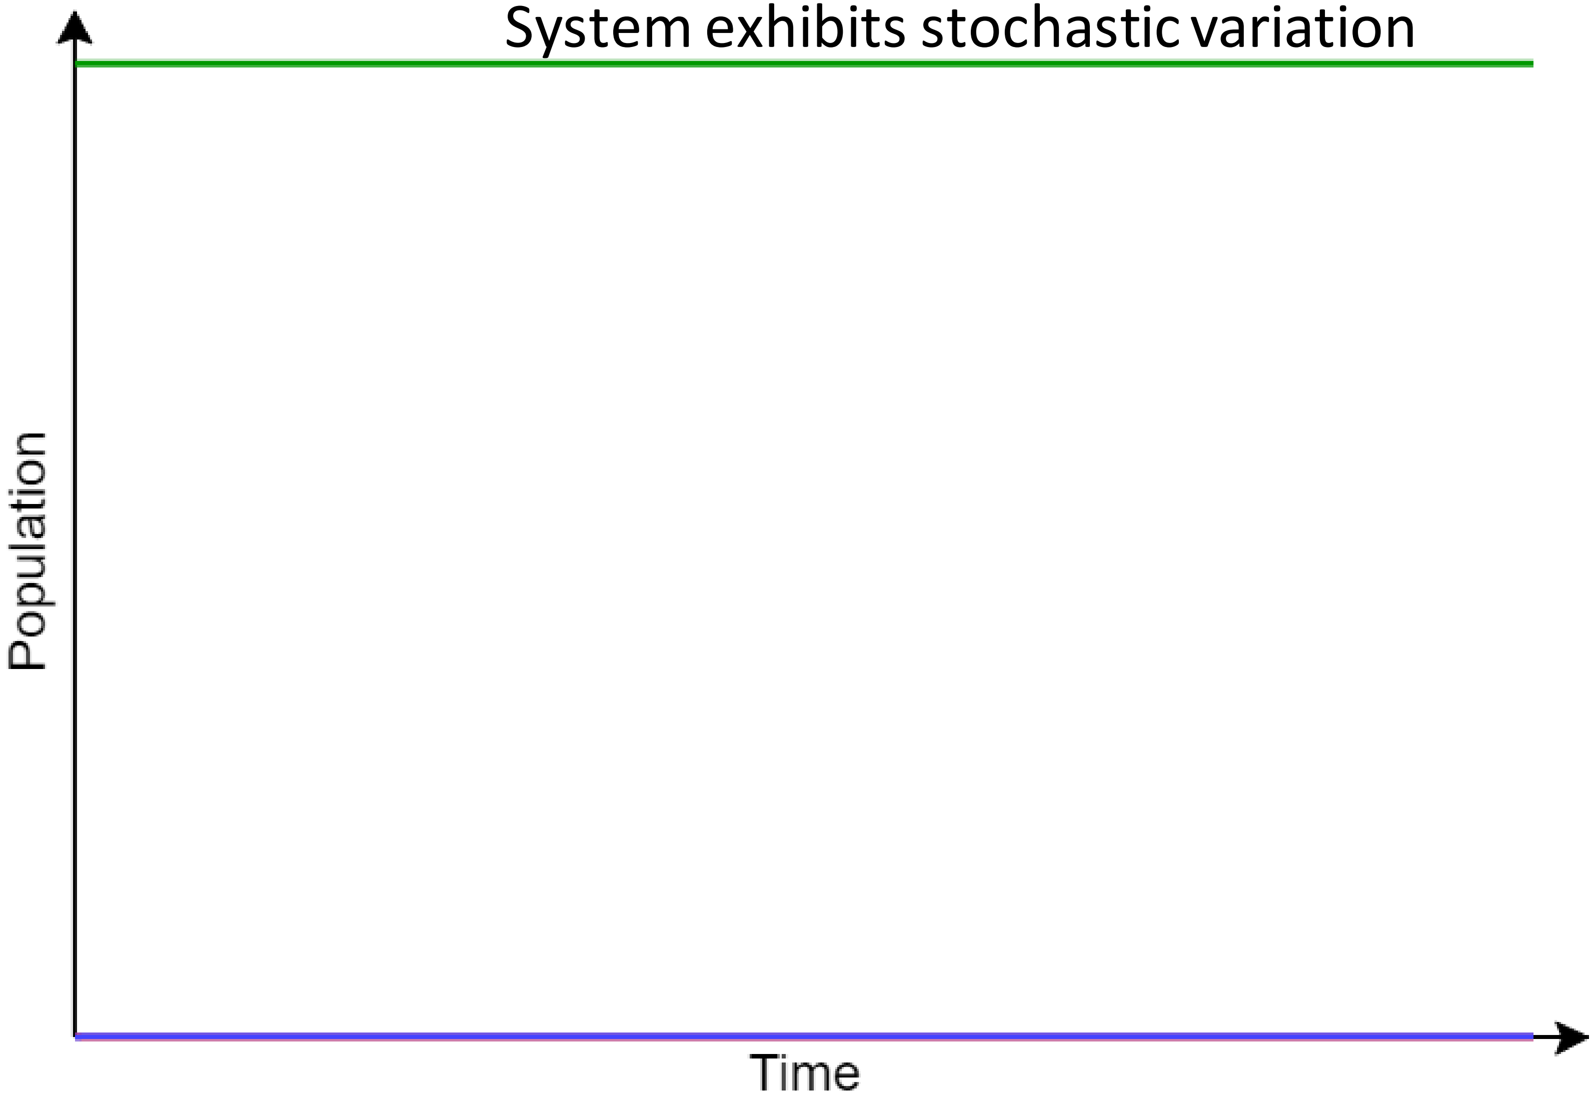

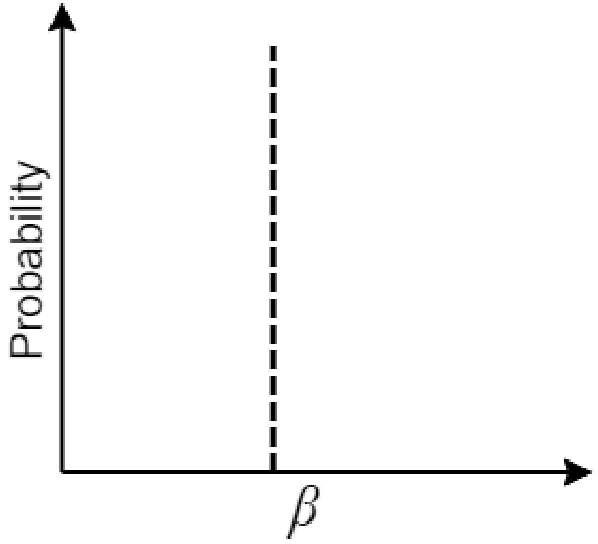

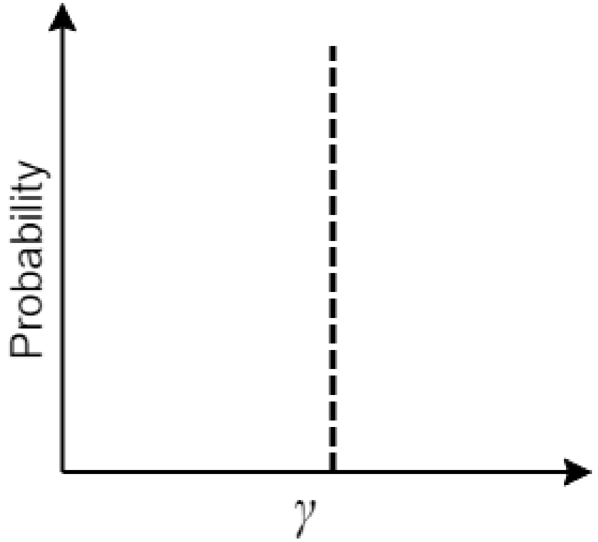

Data

Generate some hypothetical data…

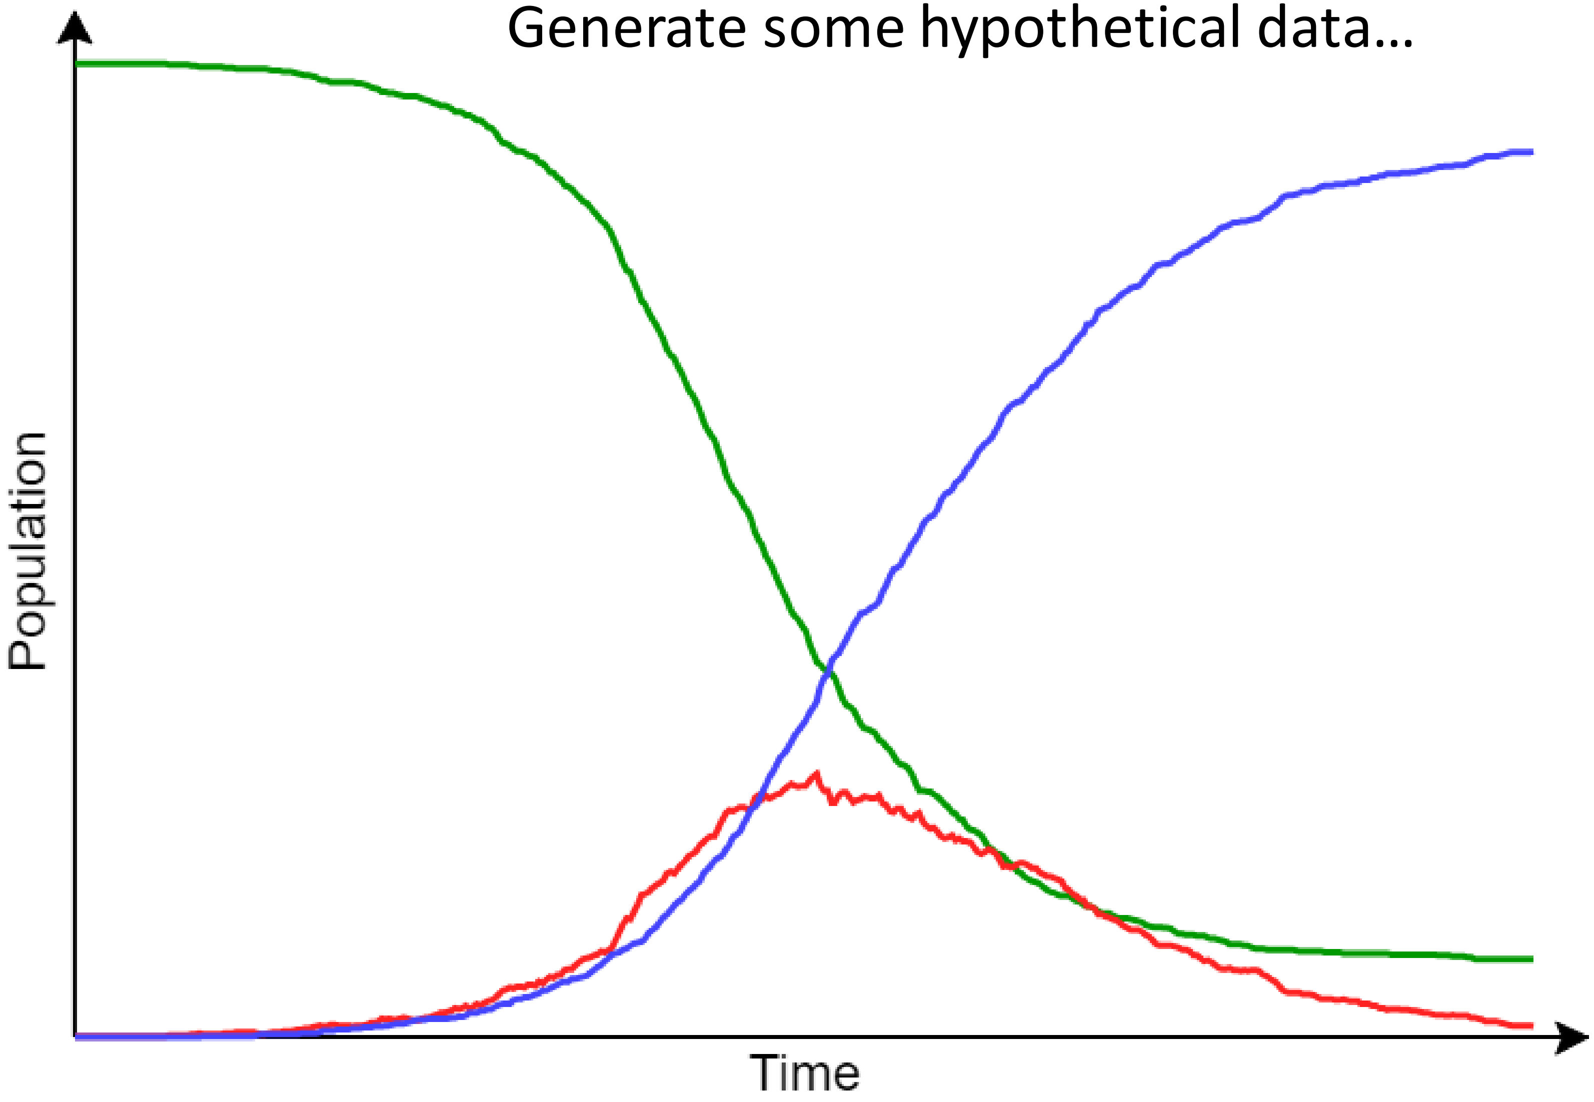

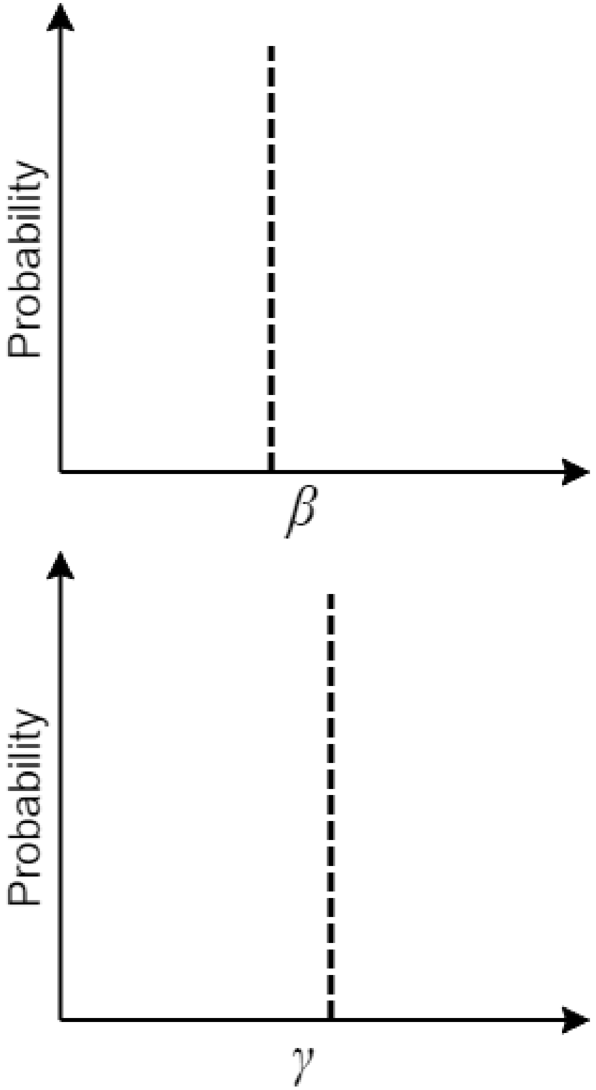

Data

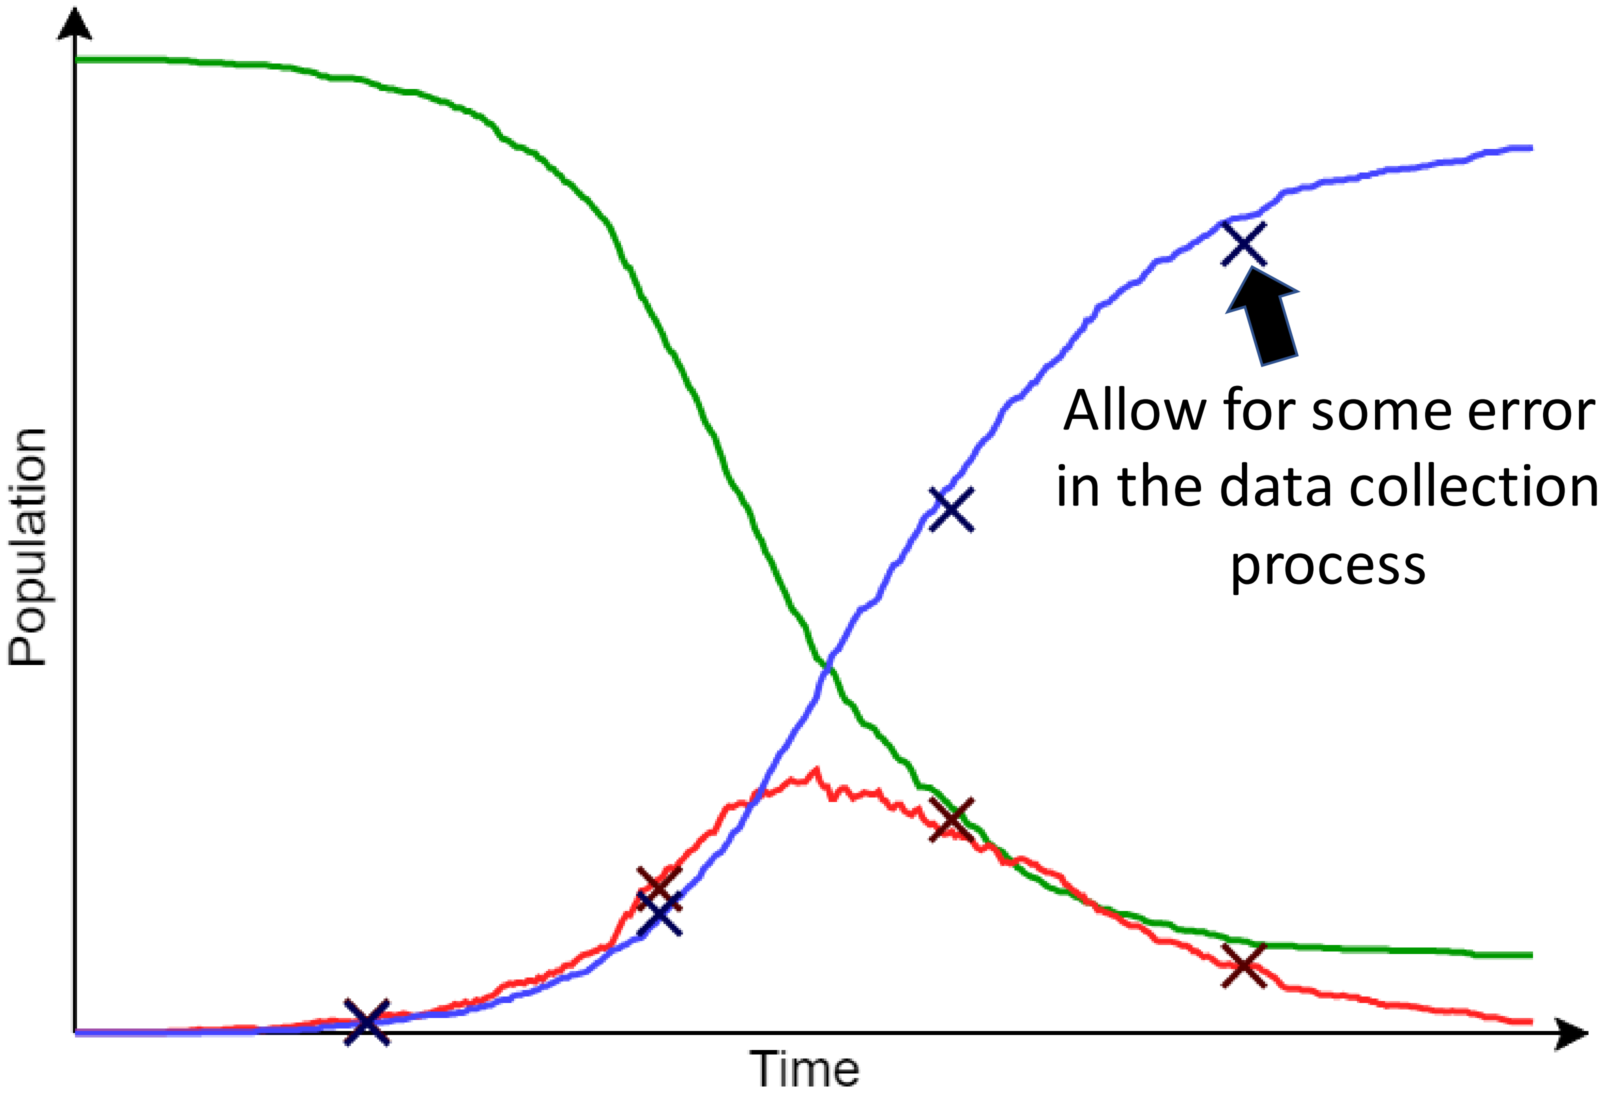
Population
Time
Allow for some error in the data collection process

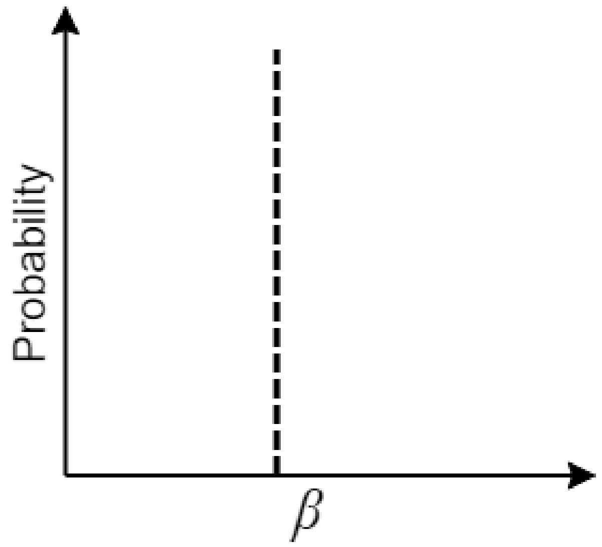
Probability
β

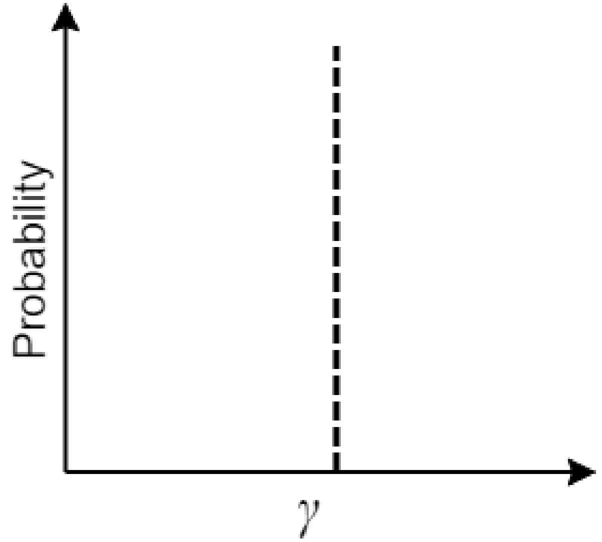
Probability
γ

Terminology

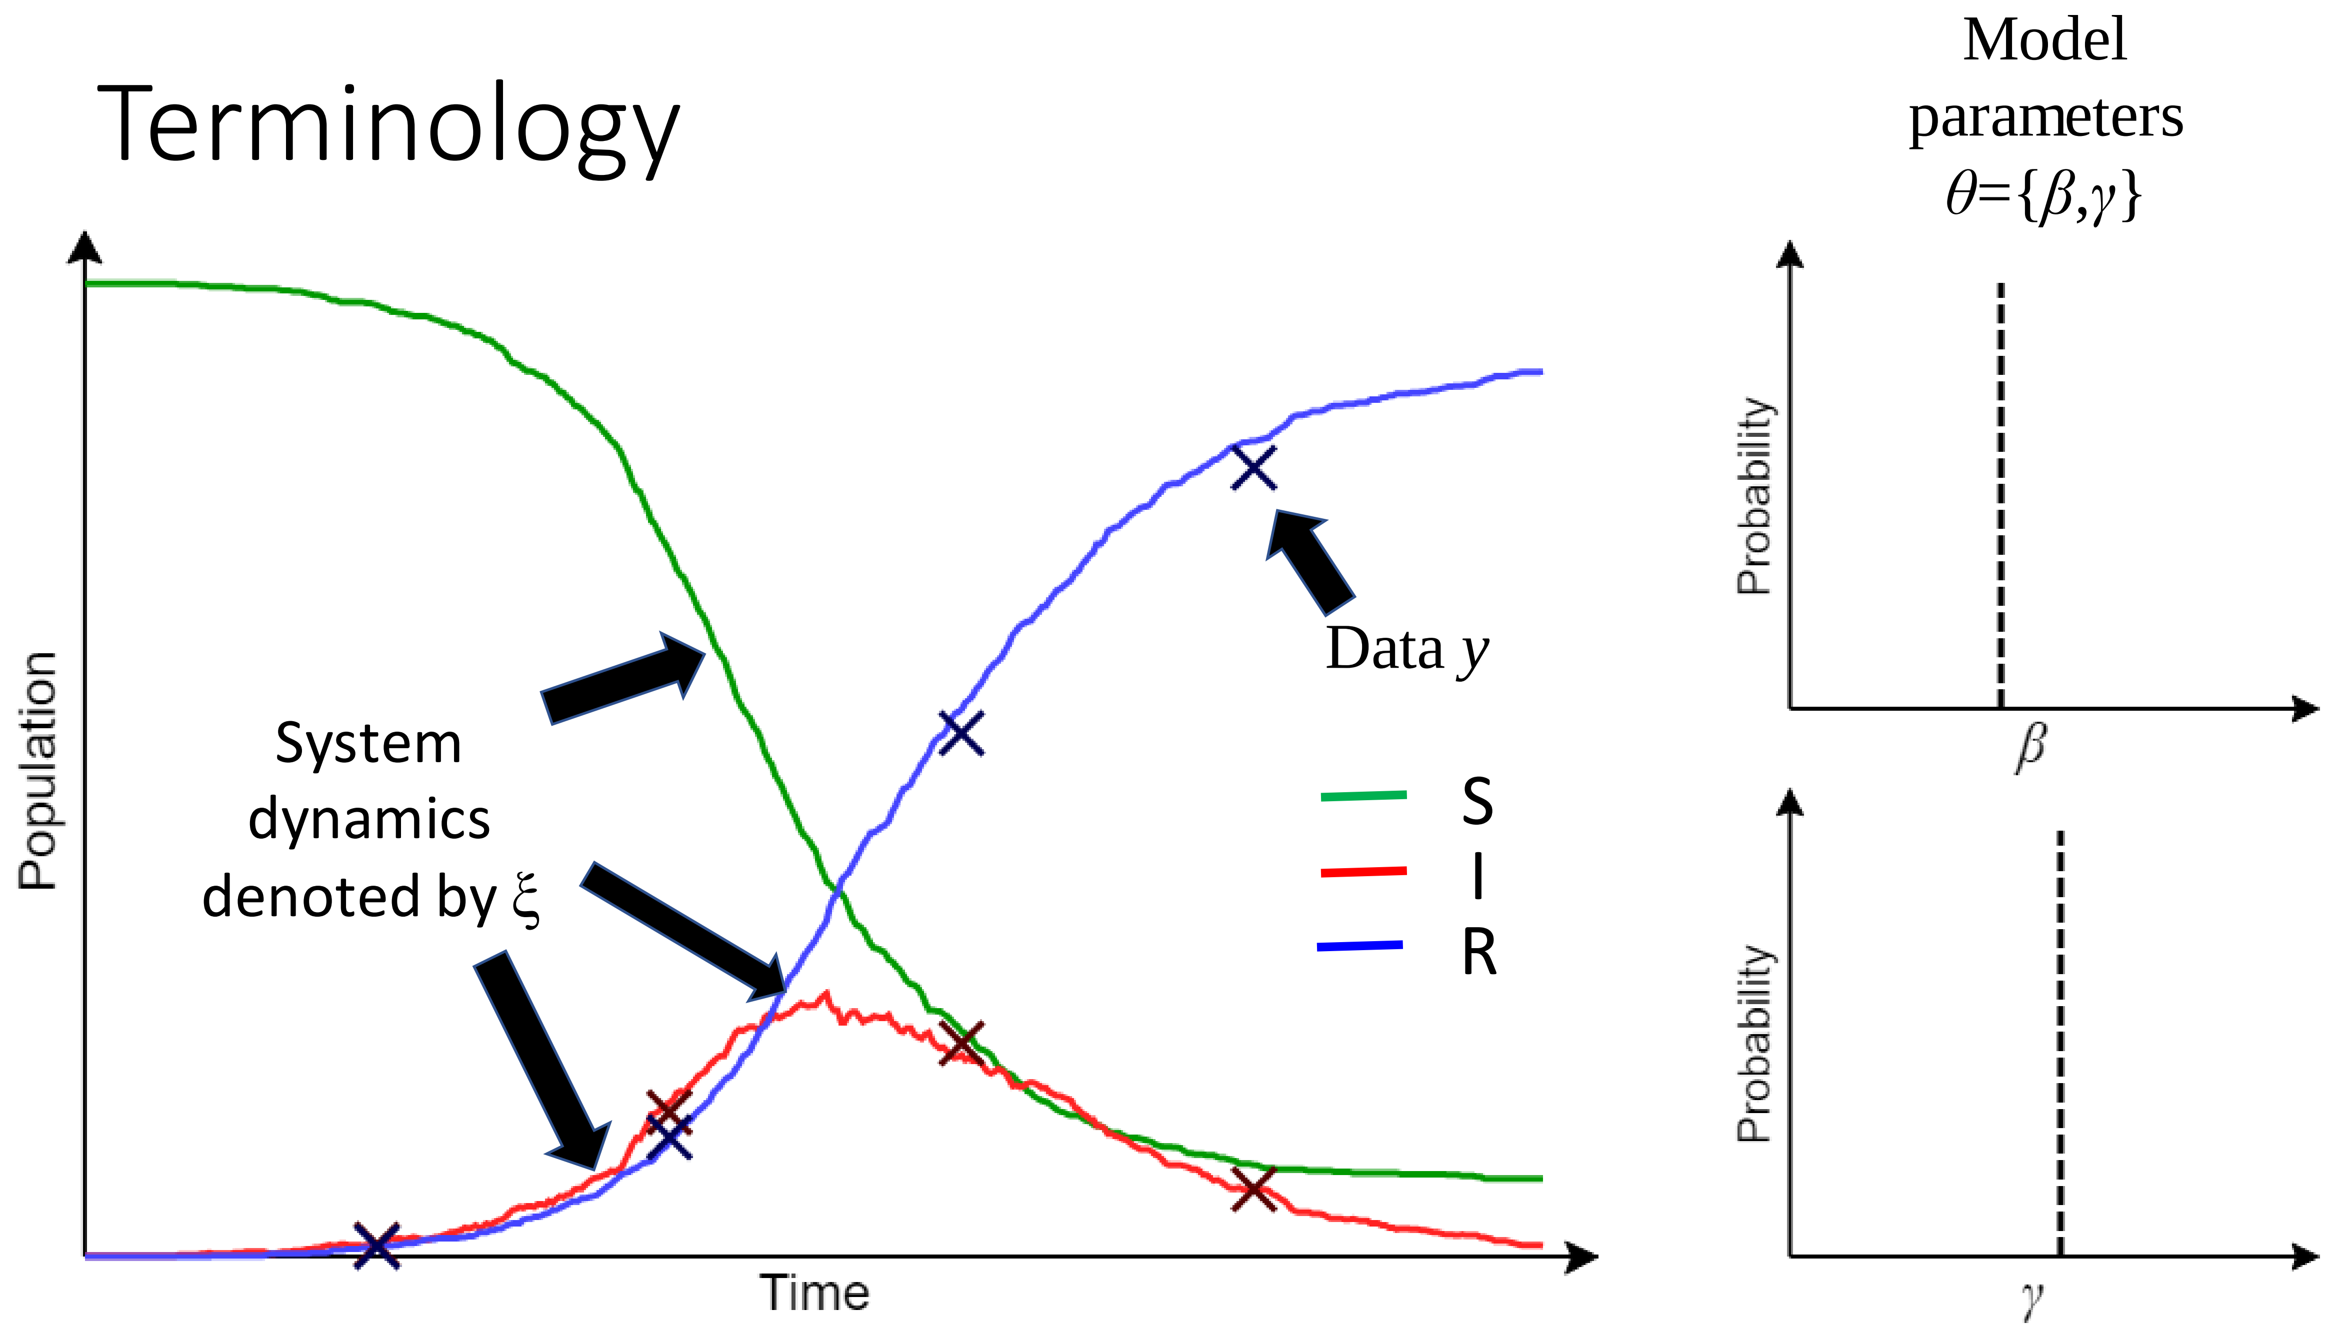

Inference

Can we estimate θ from y?

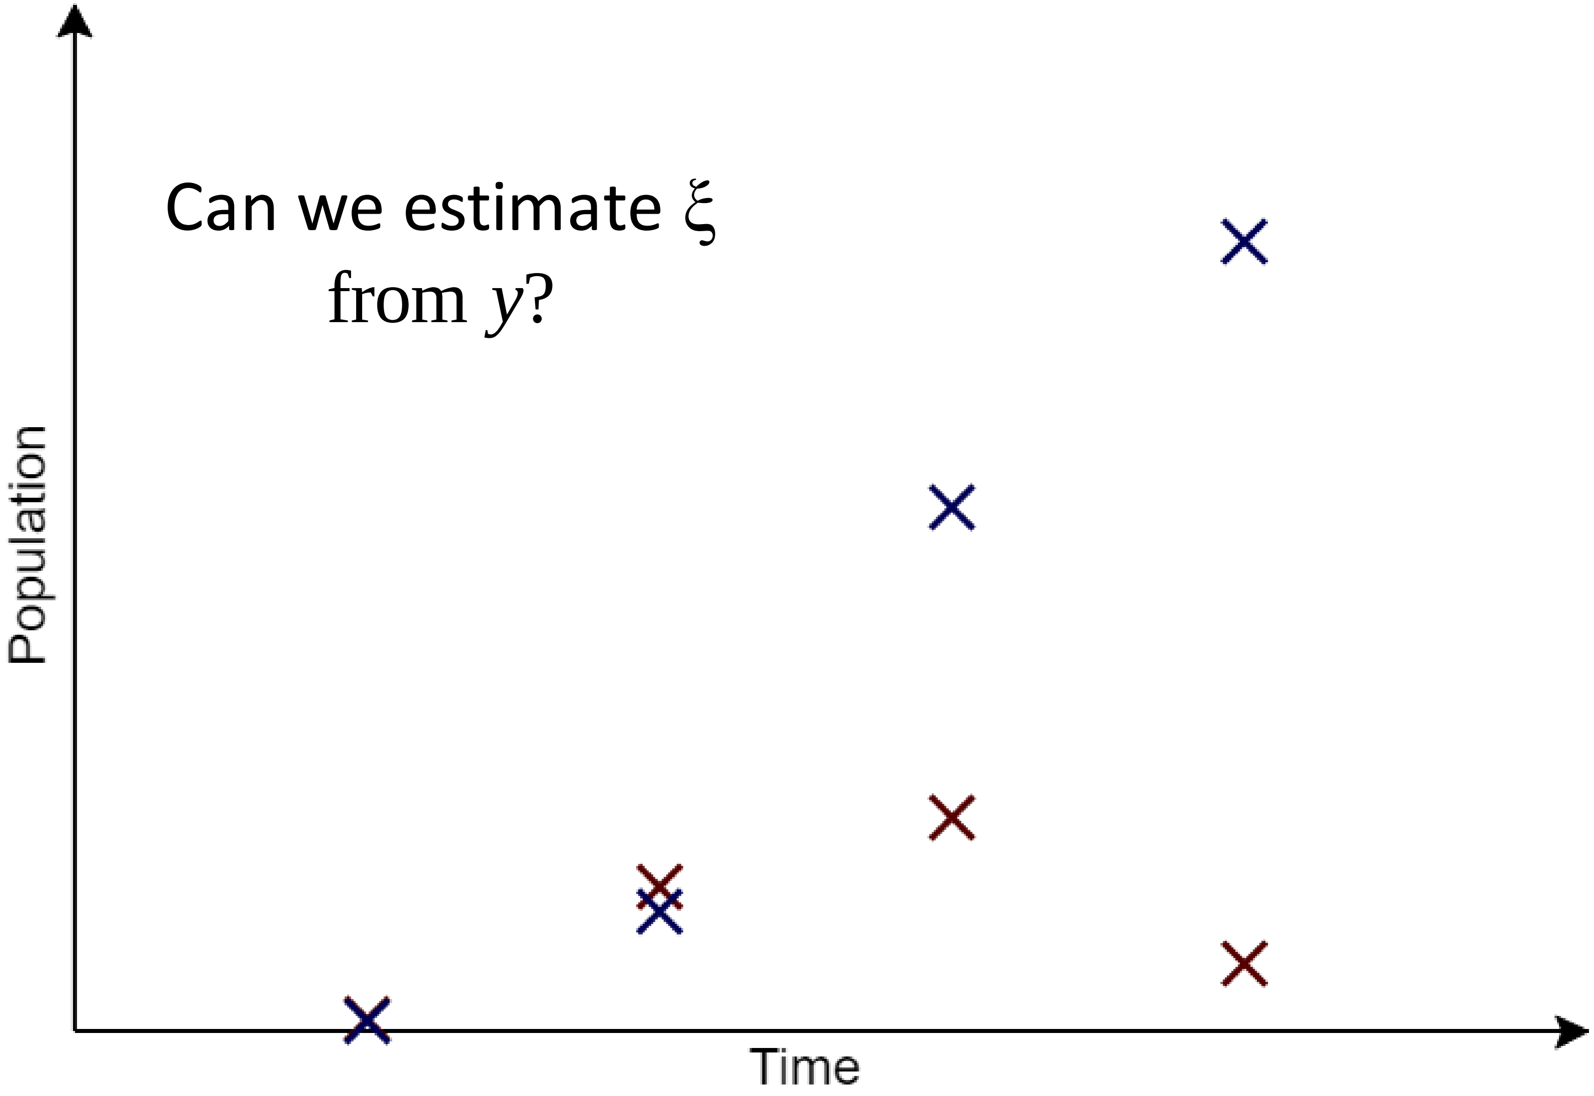

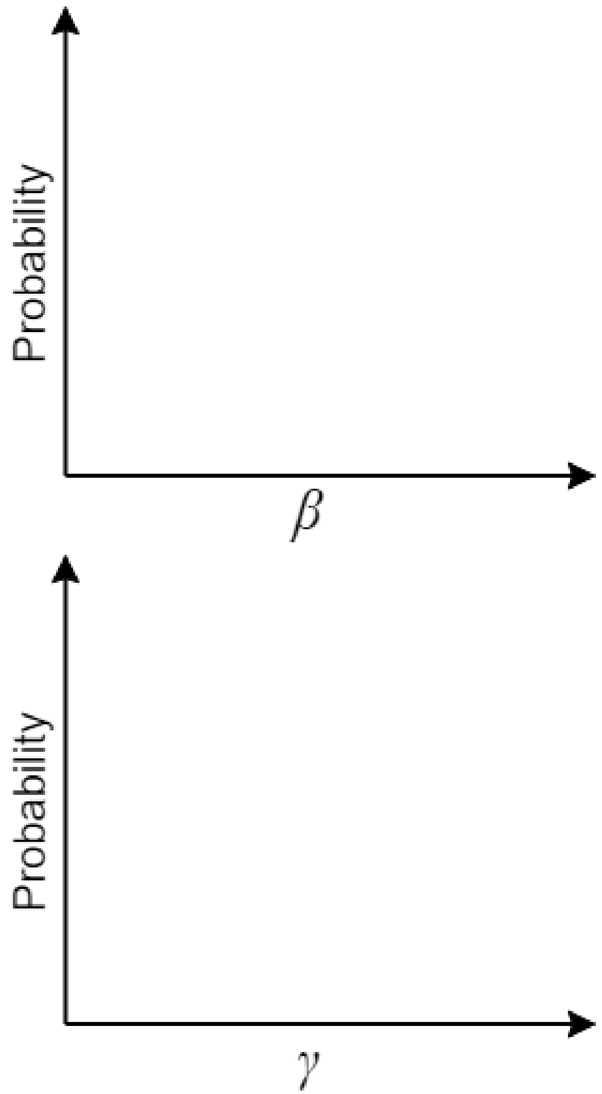

Bayesian inference

Dynamic state

Parameters

Data

$$\pi(\theta, \xi \mid y) \propto \pi(y \mid \xi) \times \pi(\xi \mid \theta) \times \pi(\theta)$$

Posterior — Observation model — Latent process likelihood — Prior

Observation model

We use some θ (not correct value)

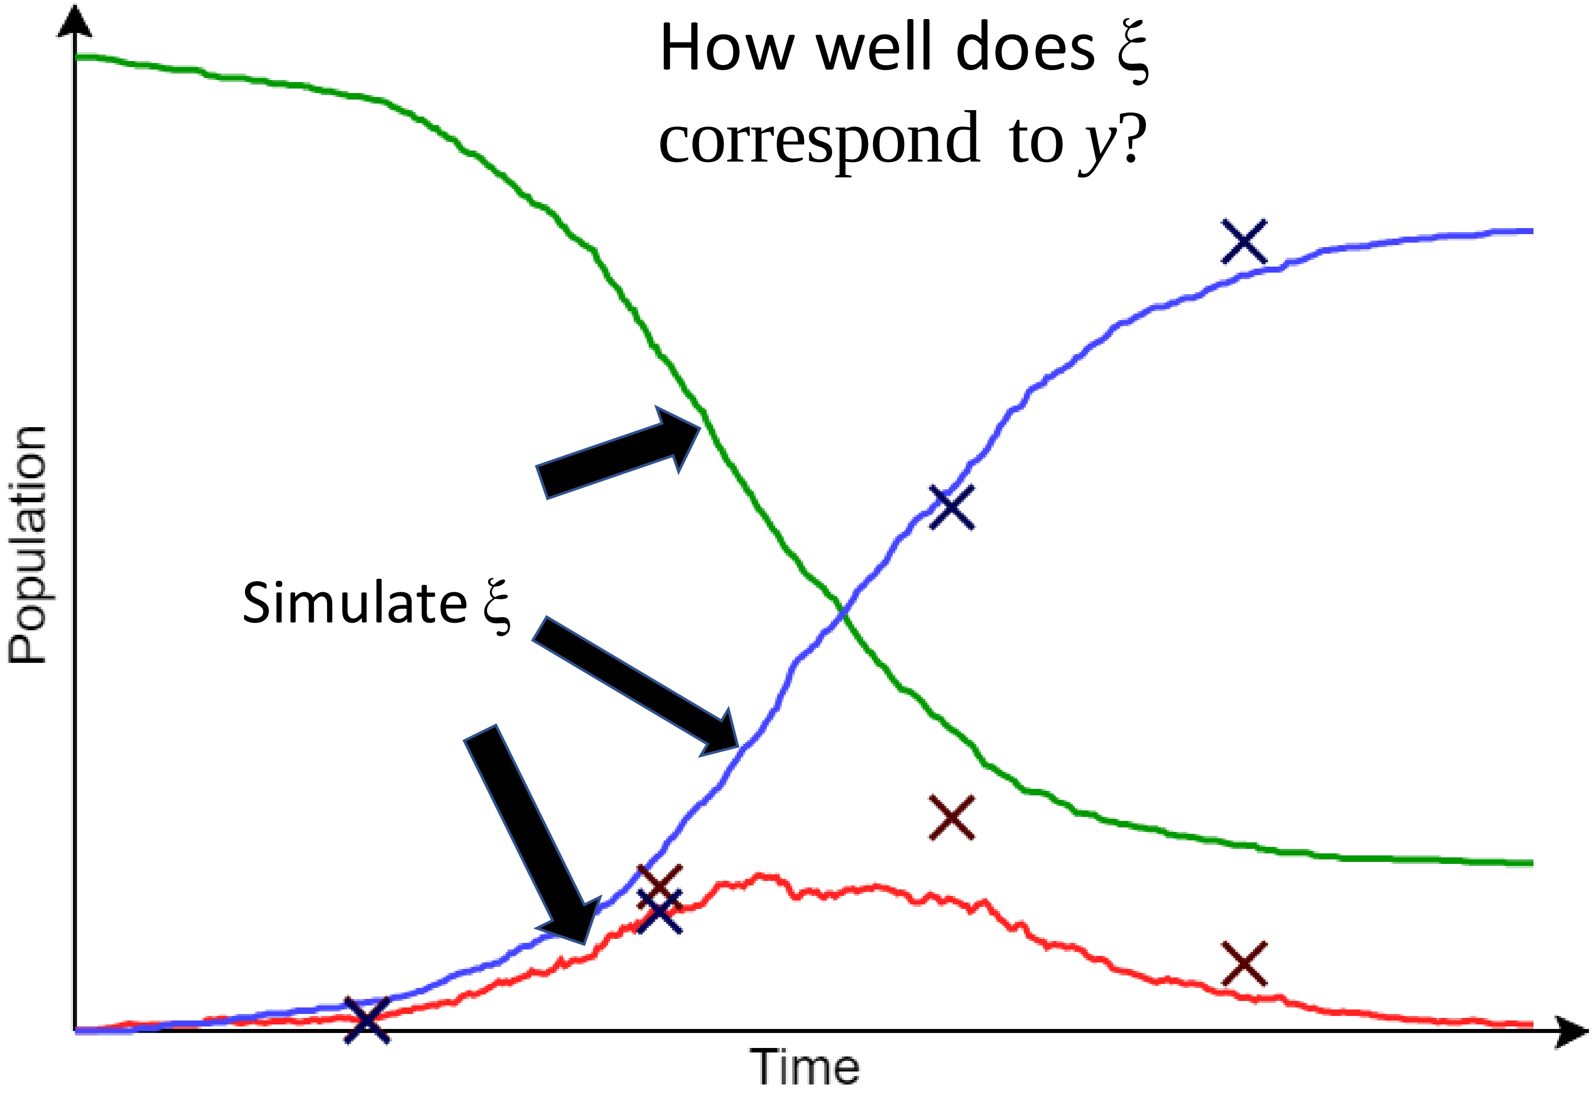

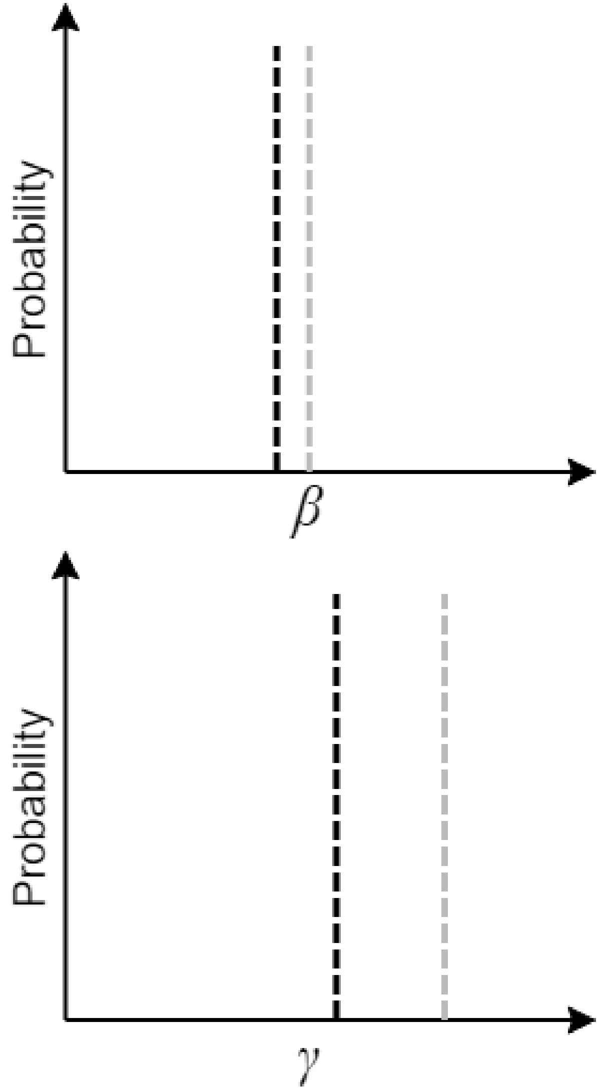

Observation model

We use some θ
(not correct value)

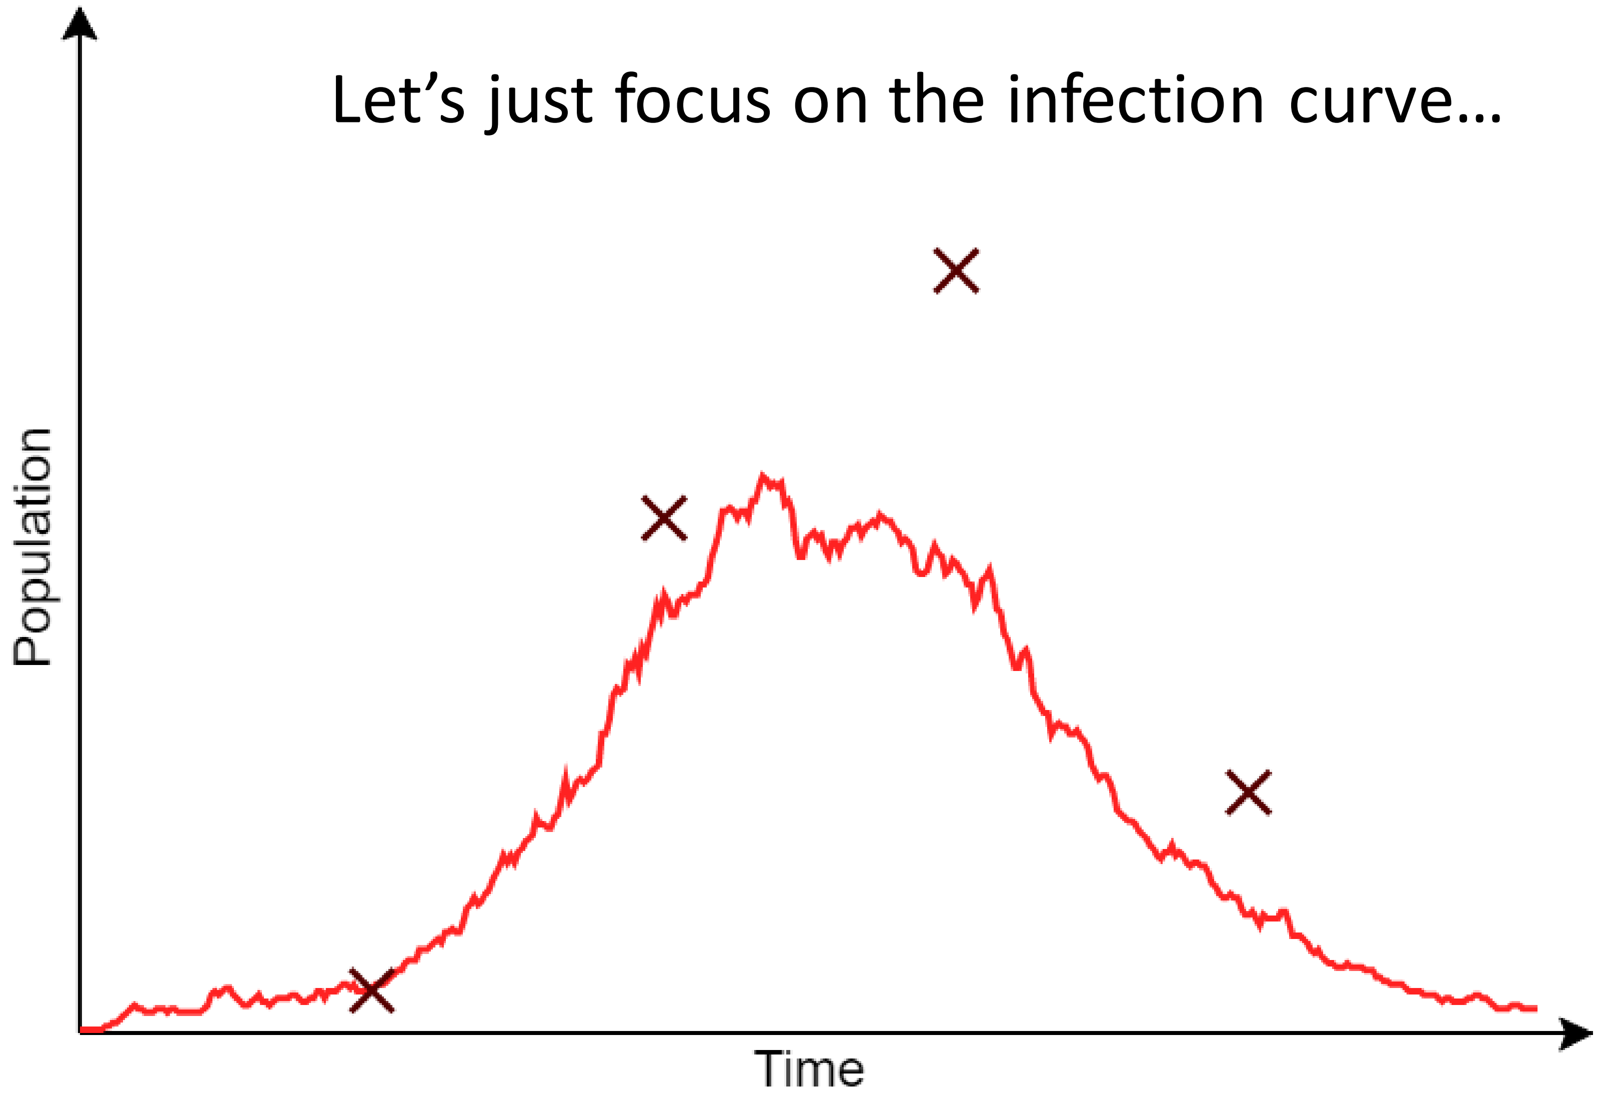

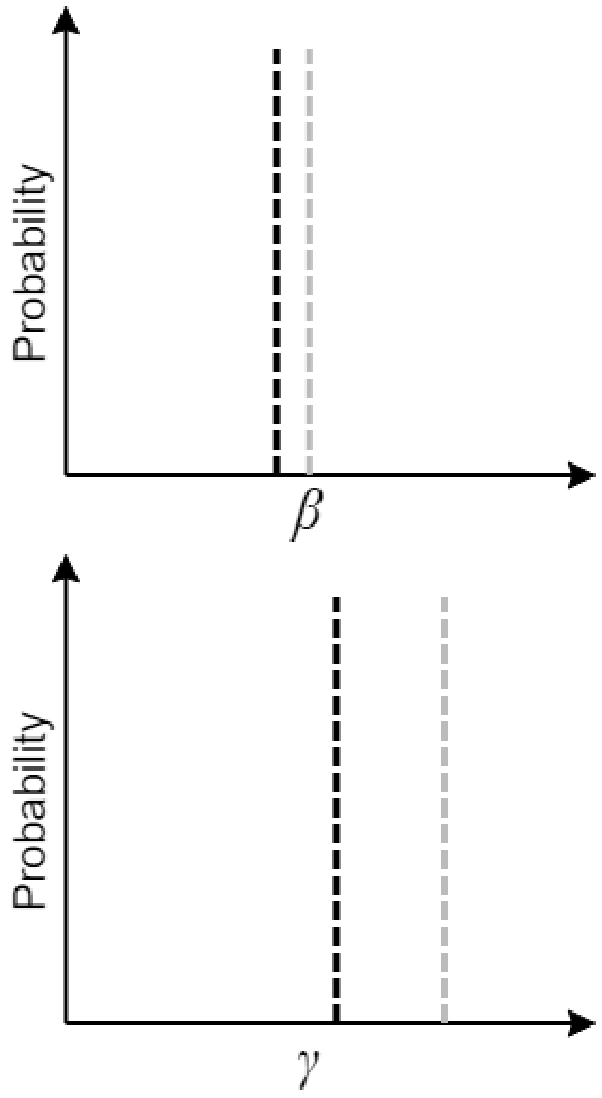

Observation model

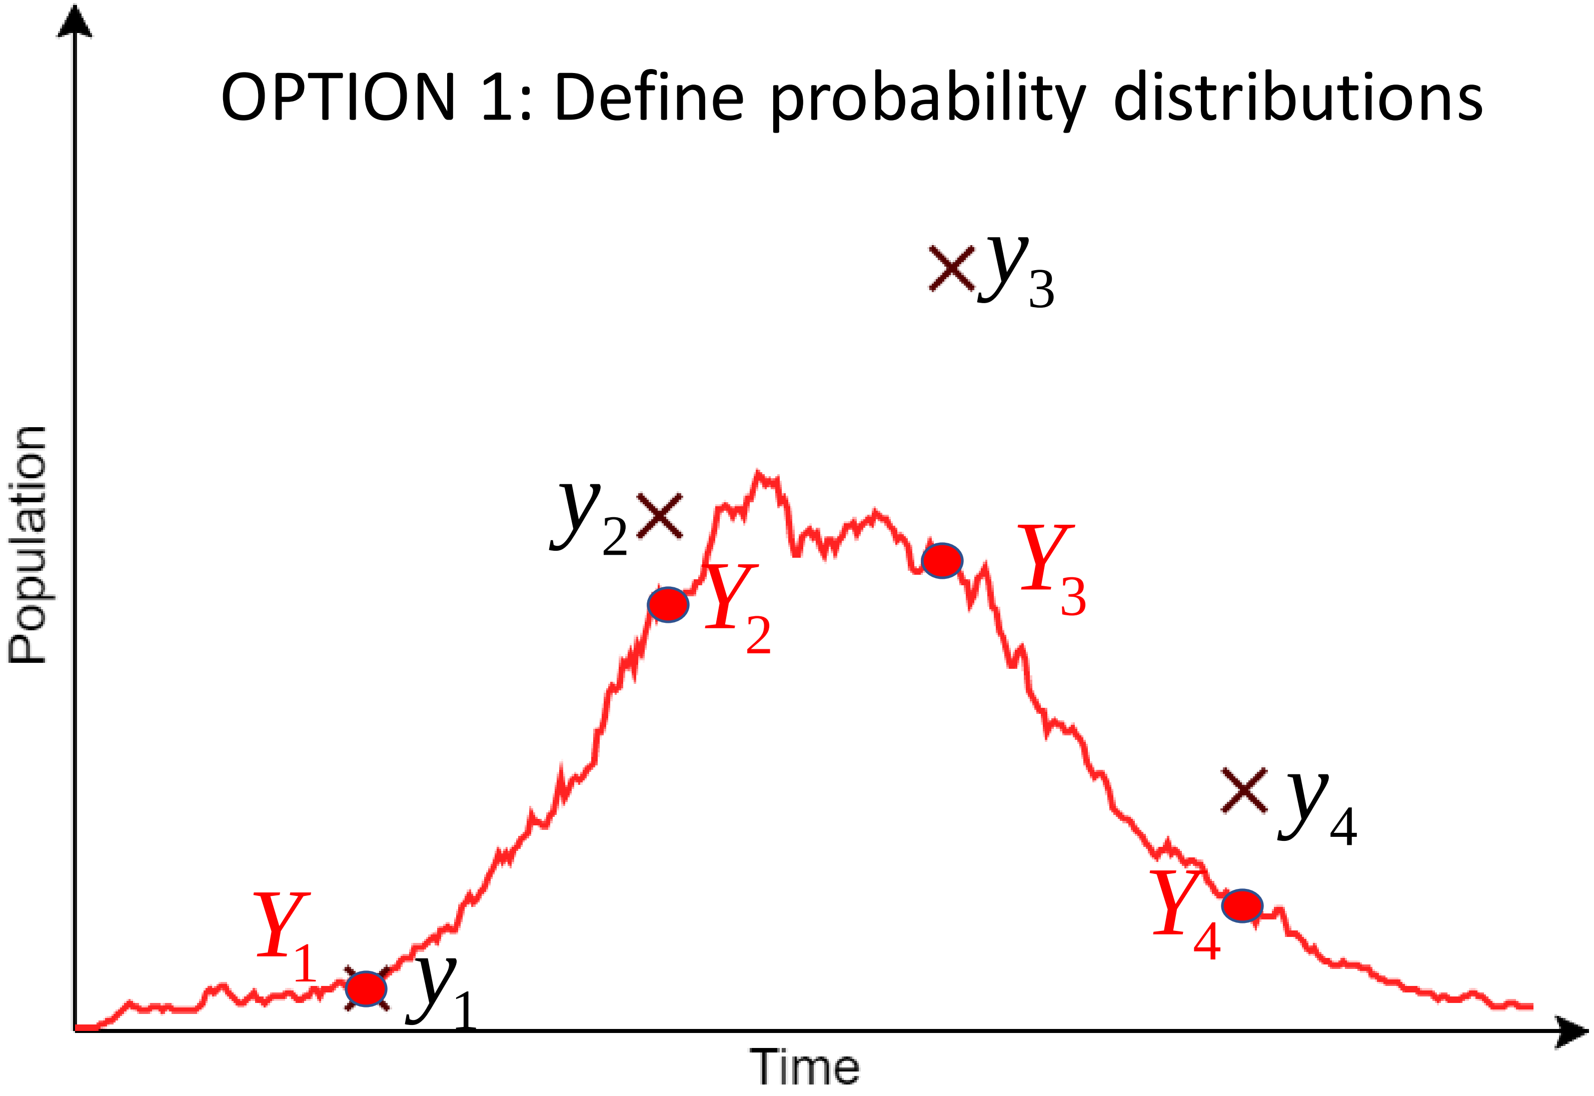

We use some θ
(not correct value)

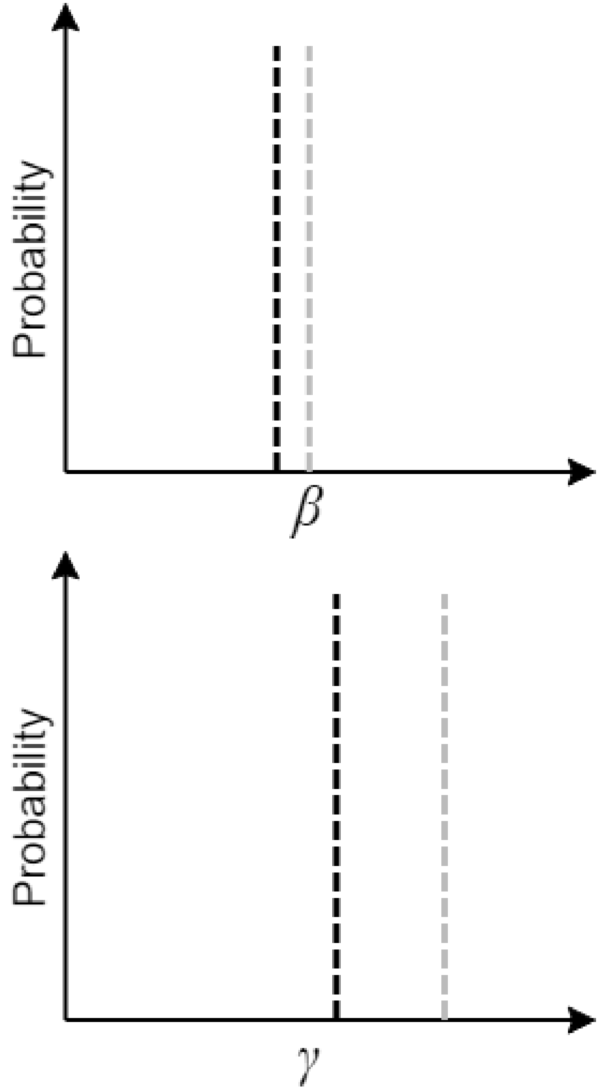

Observation model

Normal distributions: $\pi(y\,|\,\xi) = \prod_i \frac{1}{\sqrt{2\pi\sigma_i}} e^{-\frac{1}{2\sigma_i^2}(y_i - Y_i)^2}$

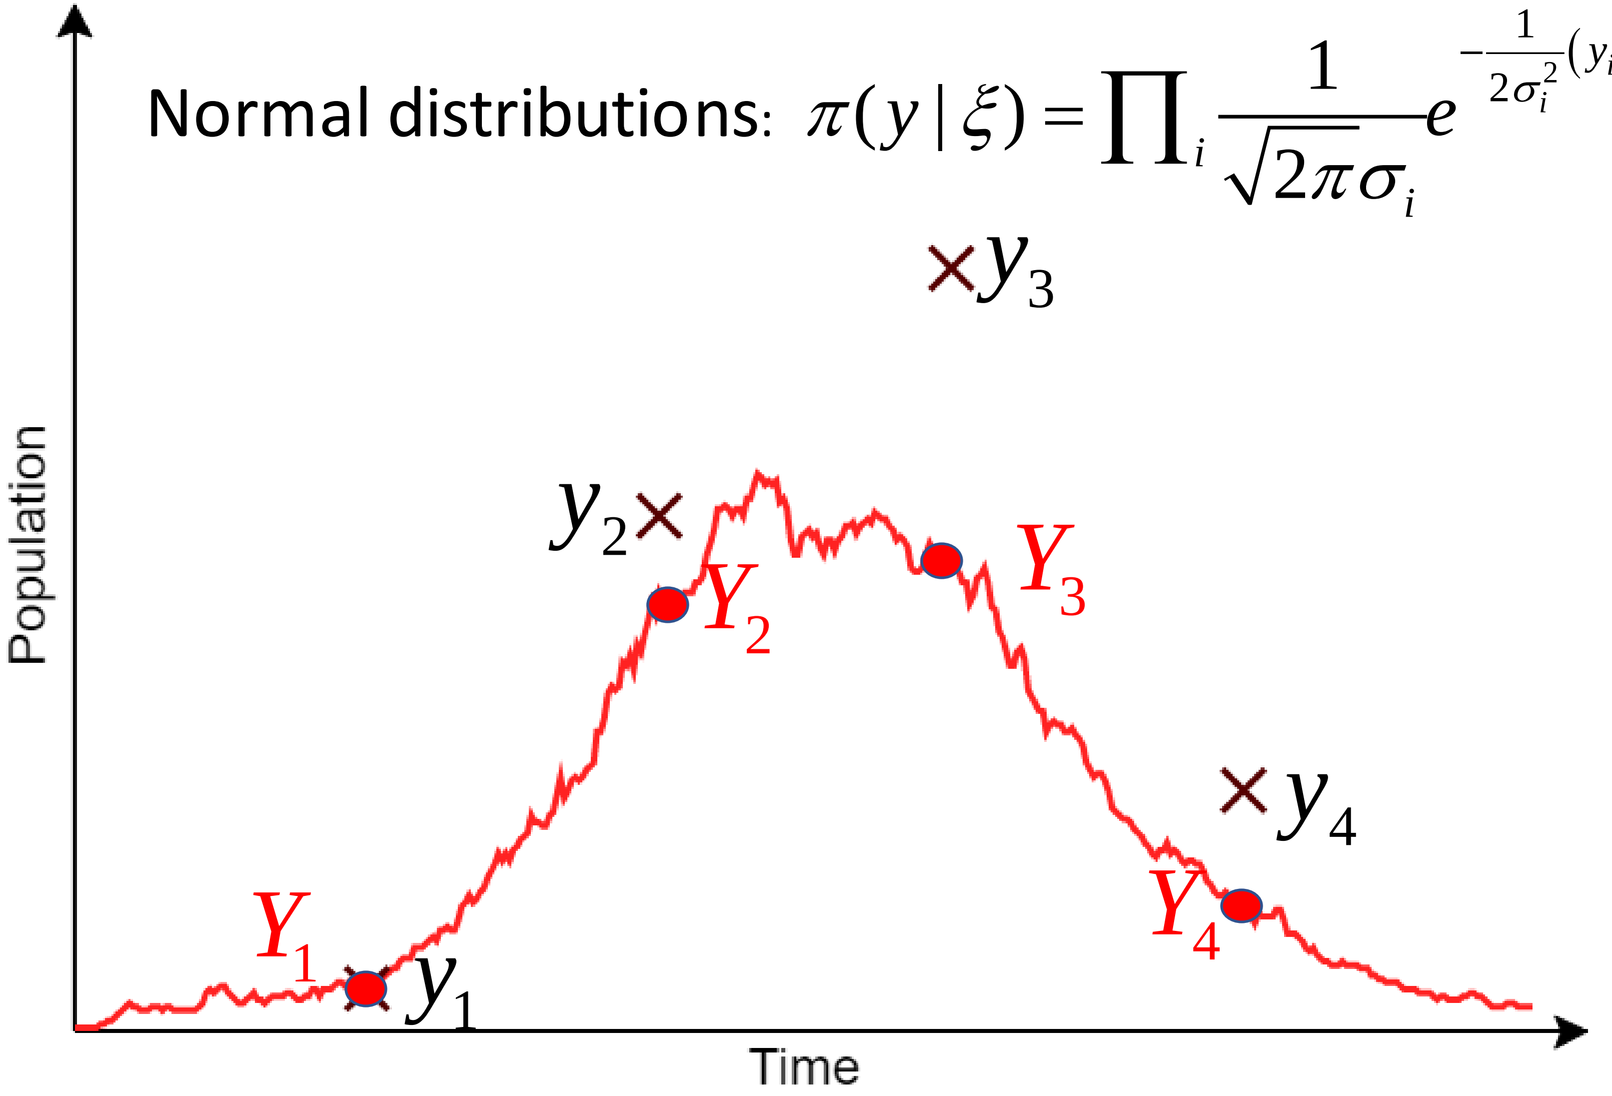

We use some θ
(not correct value)

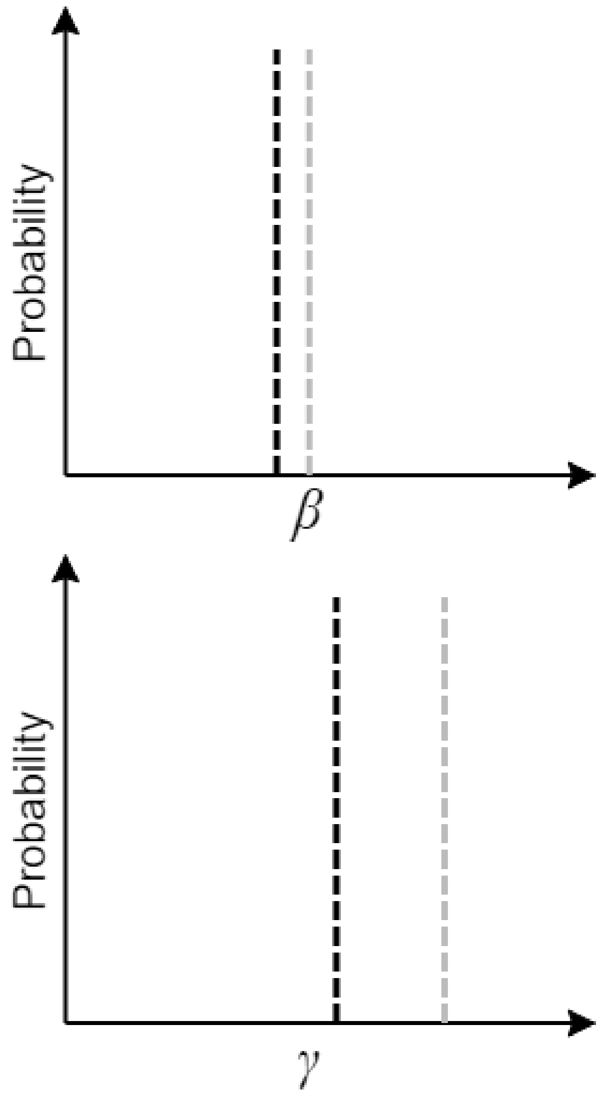

Observation model

We use some θ
(not correct value)

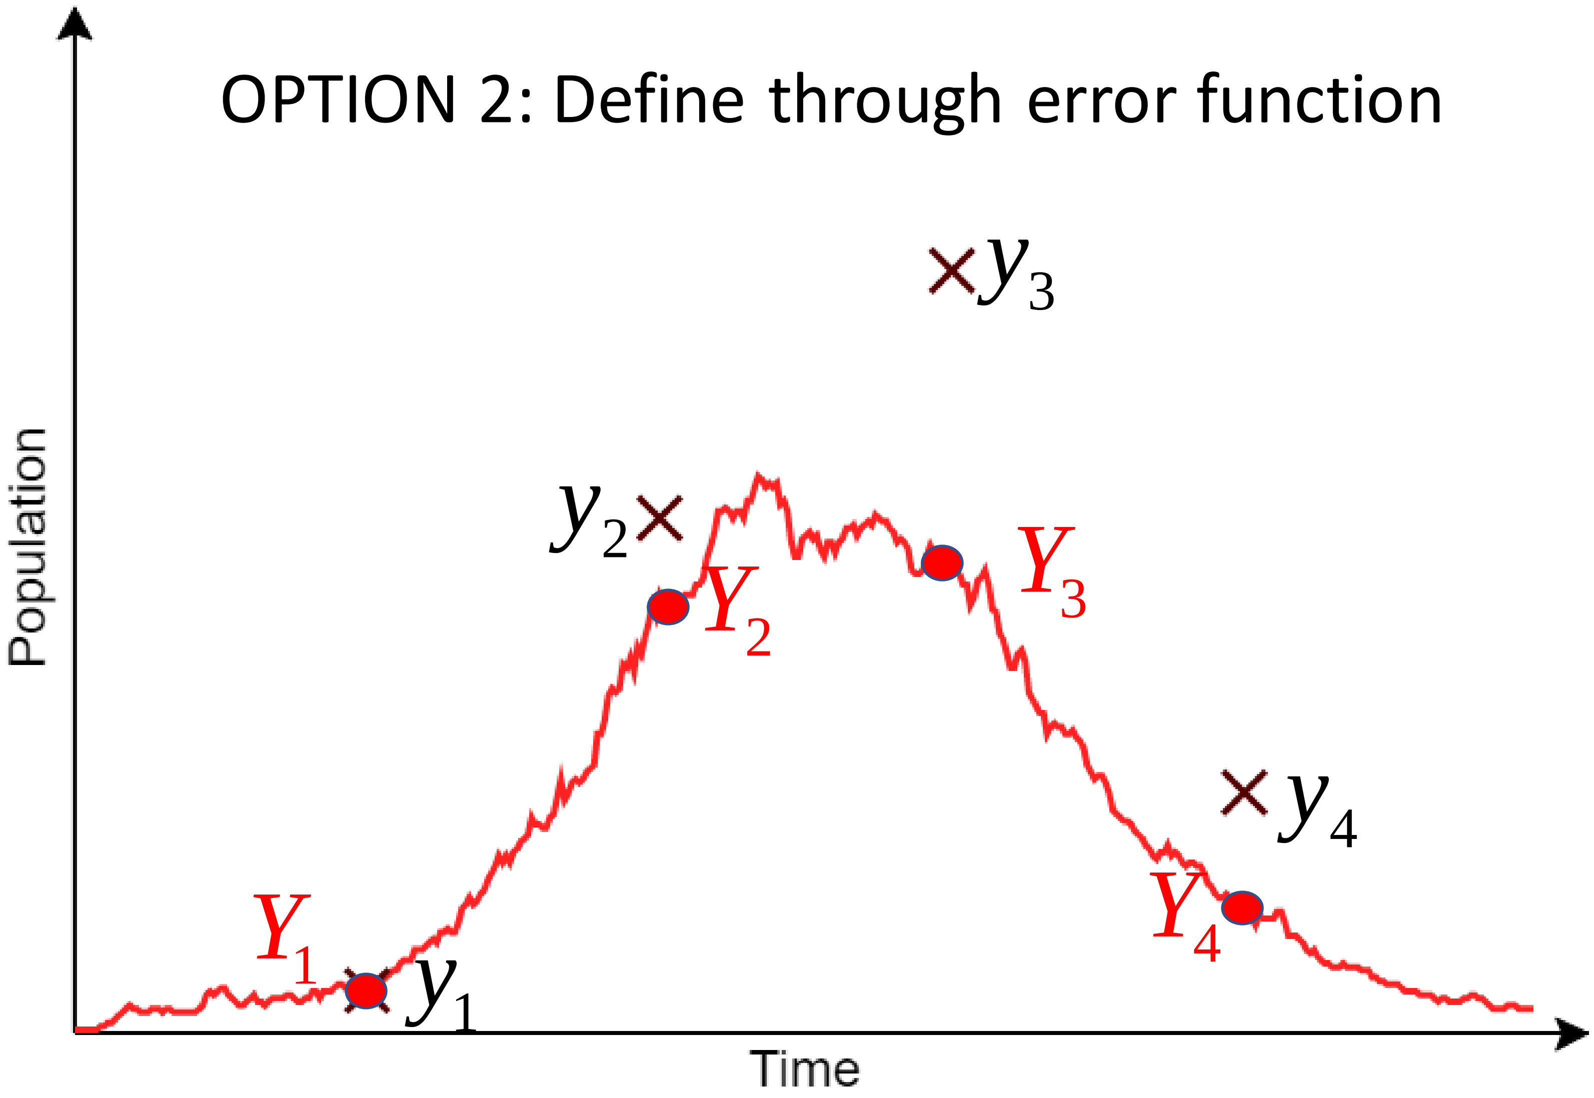

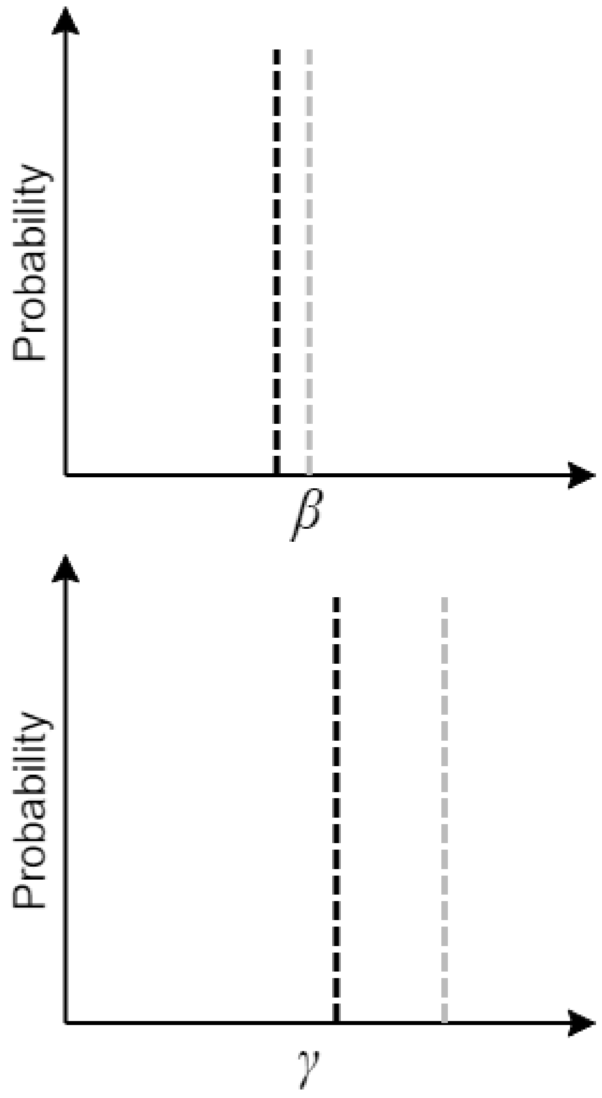

Observation model

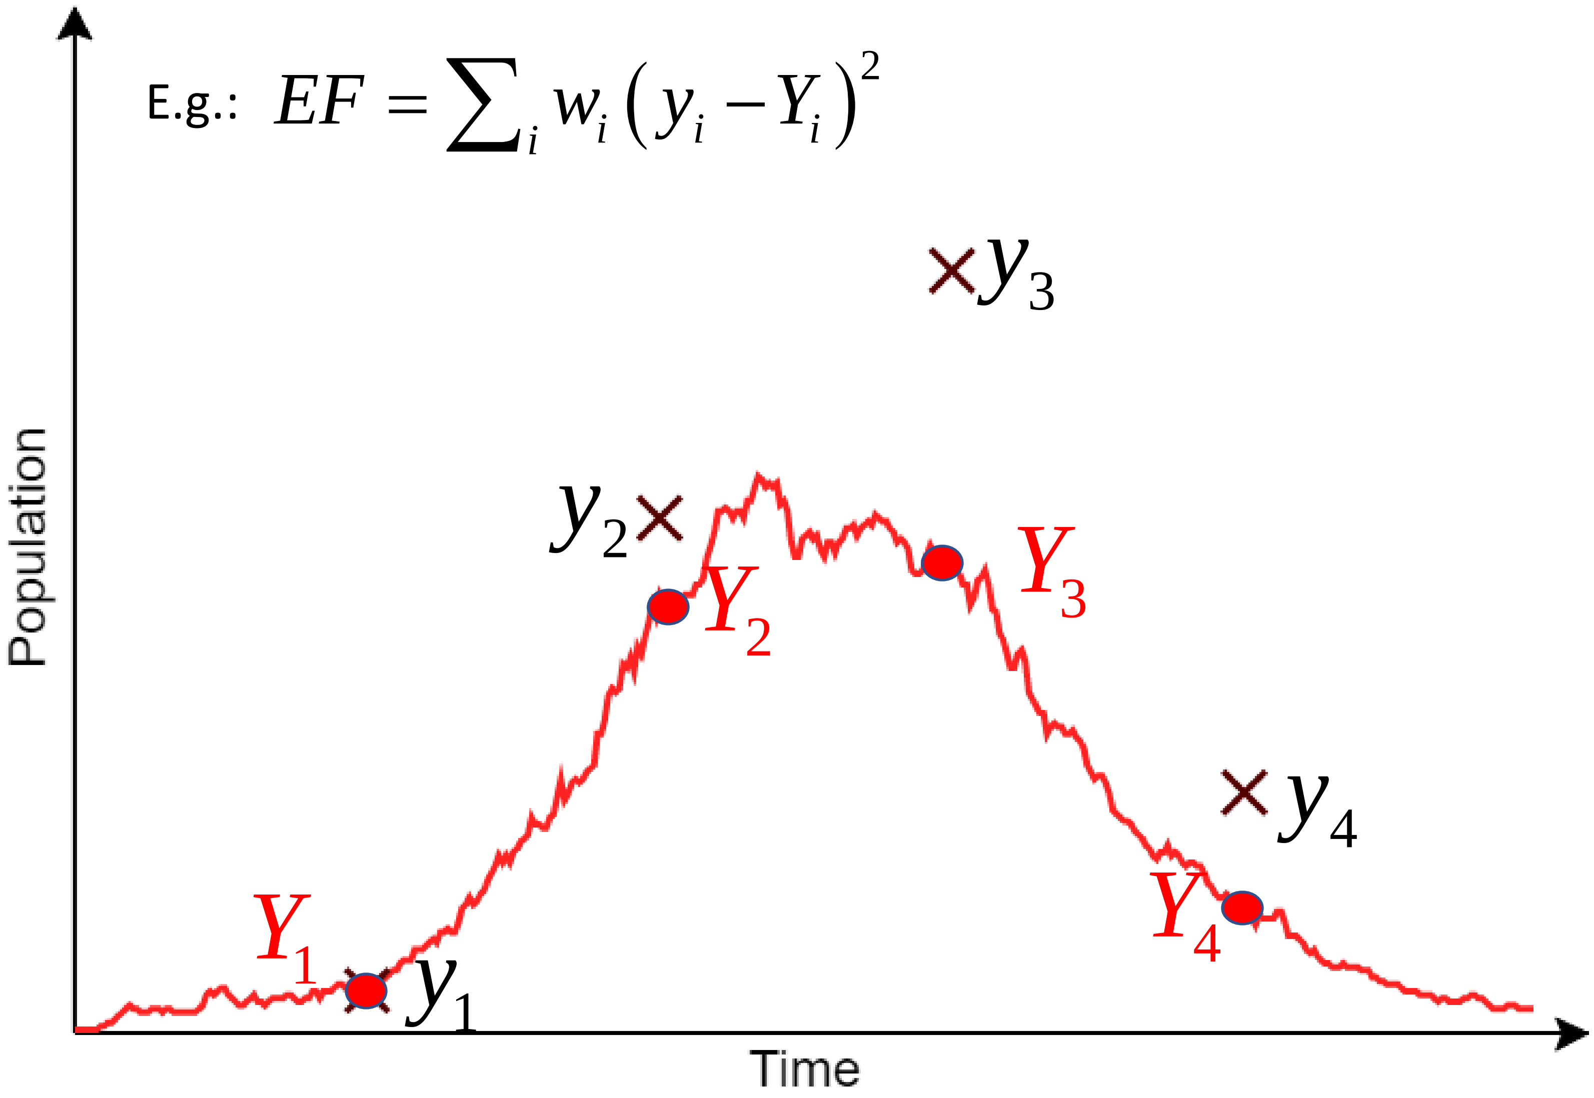

We use some θ (not correct value)

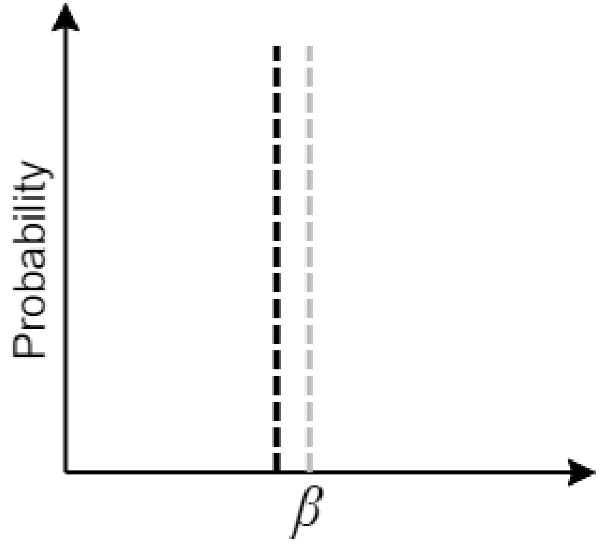

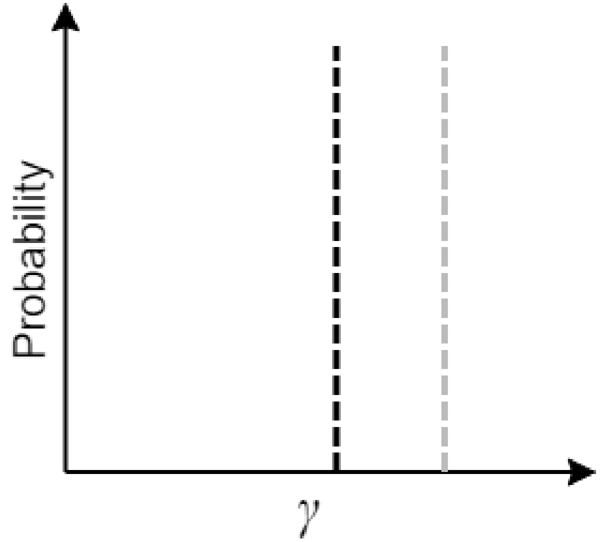

Observation model

We use some θ (not correct value)

ABC approach: $\pi(y|\xi) = \begin{cases} 1 & EF(y|\xi) < EF_{\text{cutoff}} \\ 0 & \text{Otherwise} \end{cases}$

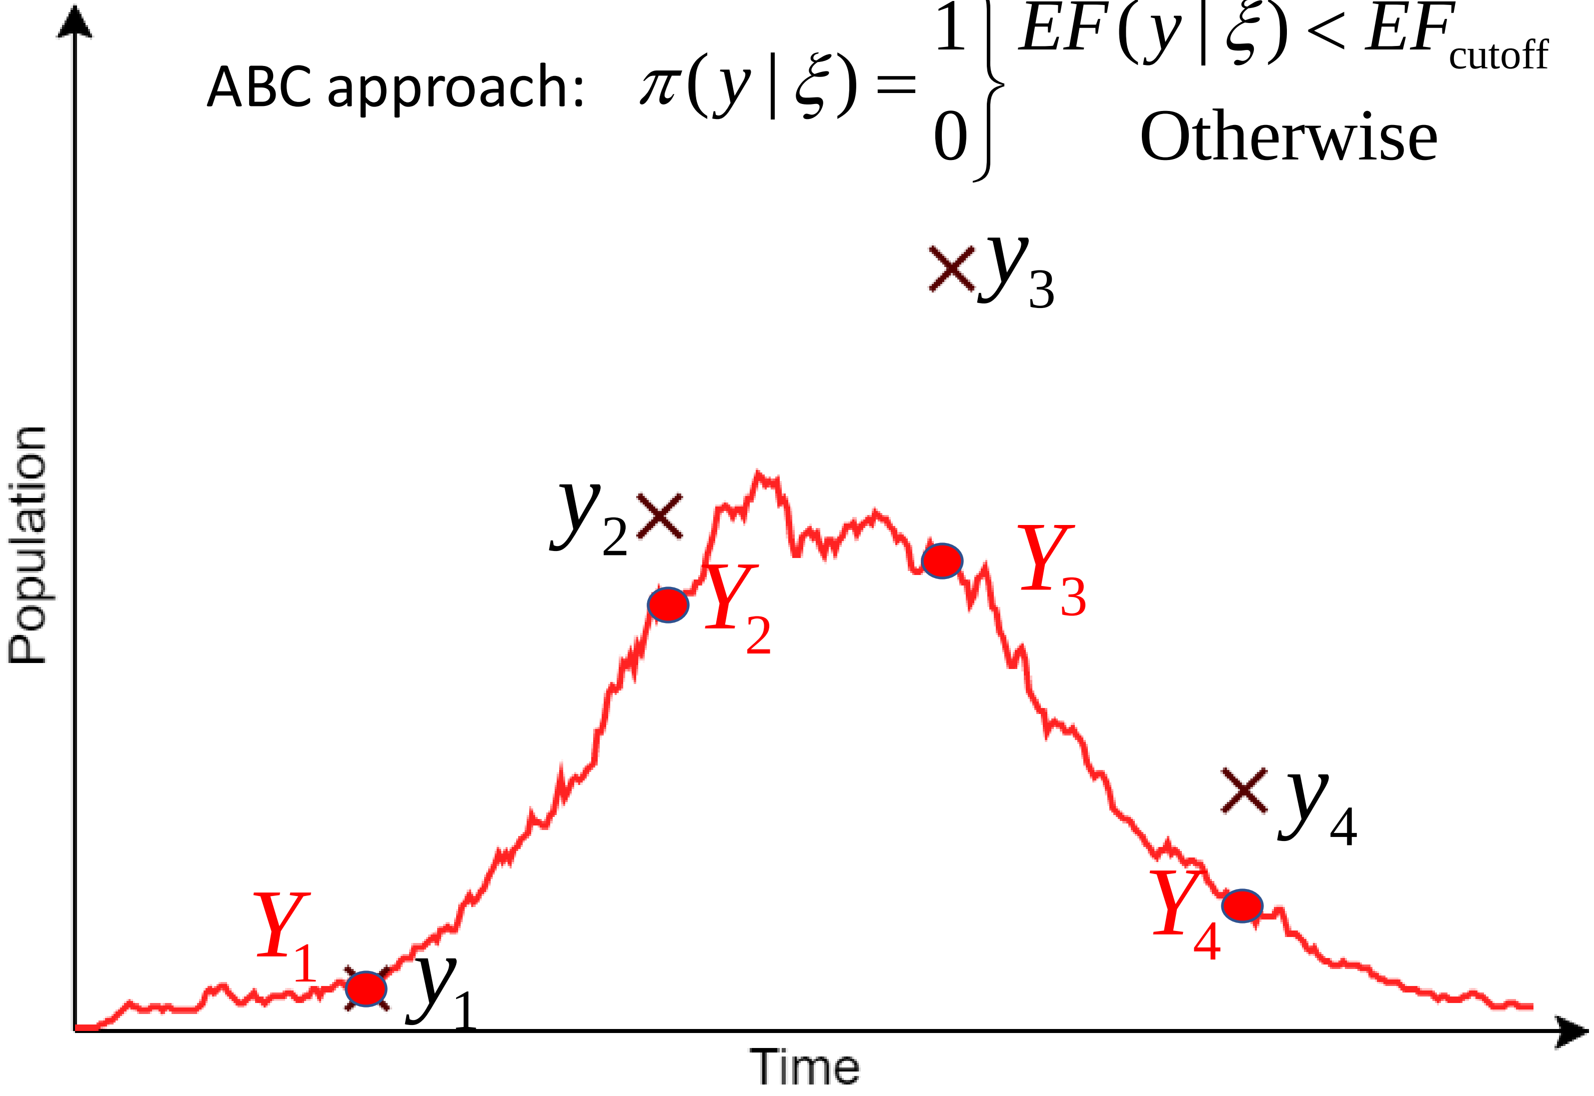

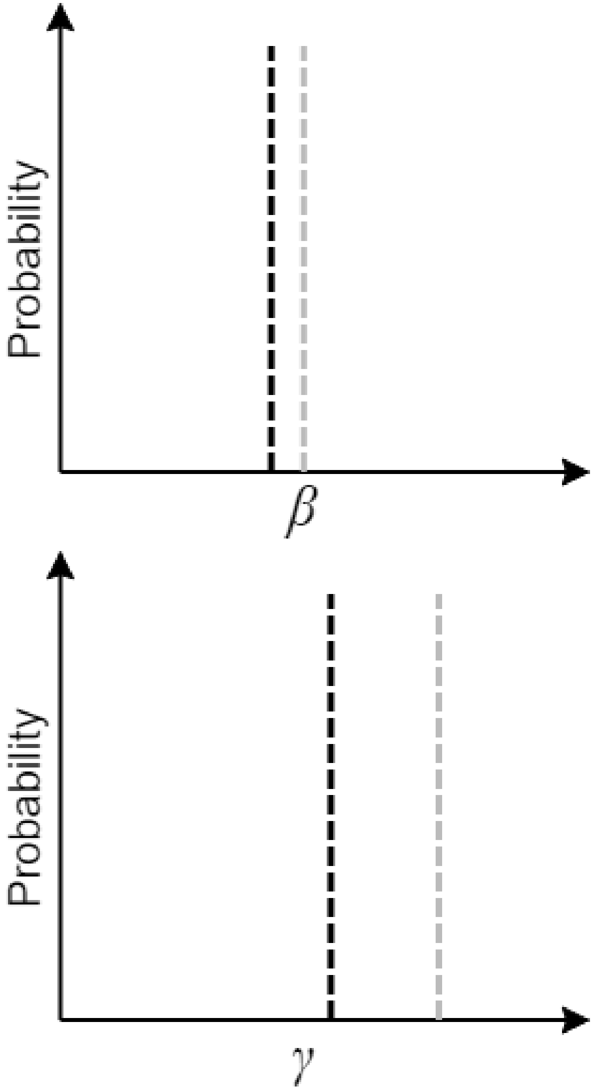

Inference algorithms

OPTION 1: Observation model with probability distributions

- Data augmentation Markov chain Monte Carlo (DA-MCMC)
- Particle MCMC (PMCMC)
- Metropolis coupled MCMC (MC^3)

OPTION 2: Error functions

- Approximate Bayesian computation rejection sampling (ABC)
- ABC sequential Monte Carlo (ABC-SMC)
- ABC with model-based proposals (ABC-MBP)

ABC rejection sampling

Prior $\pi(\theta)$ defines plausible ranges for parameters

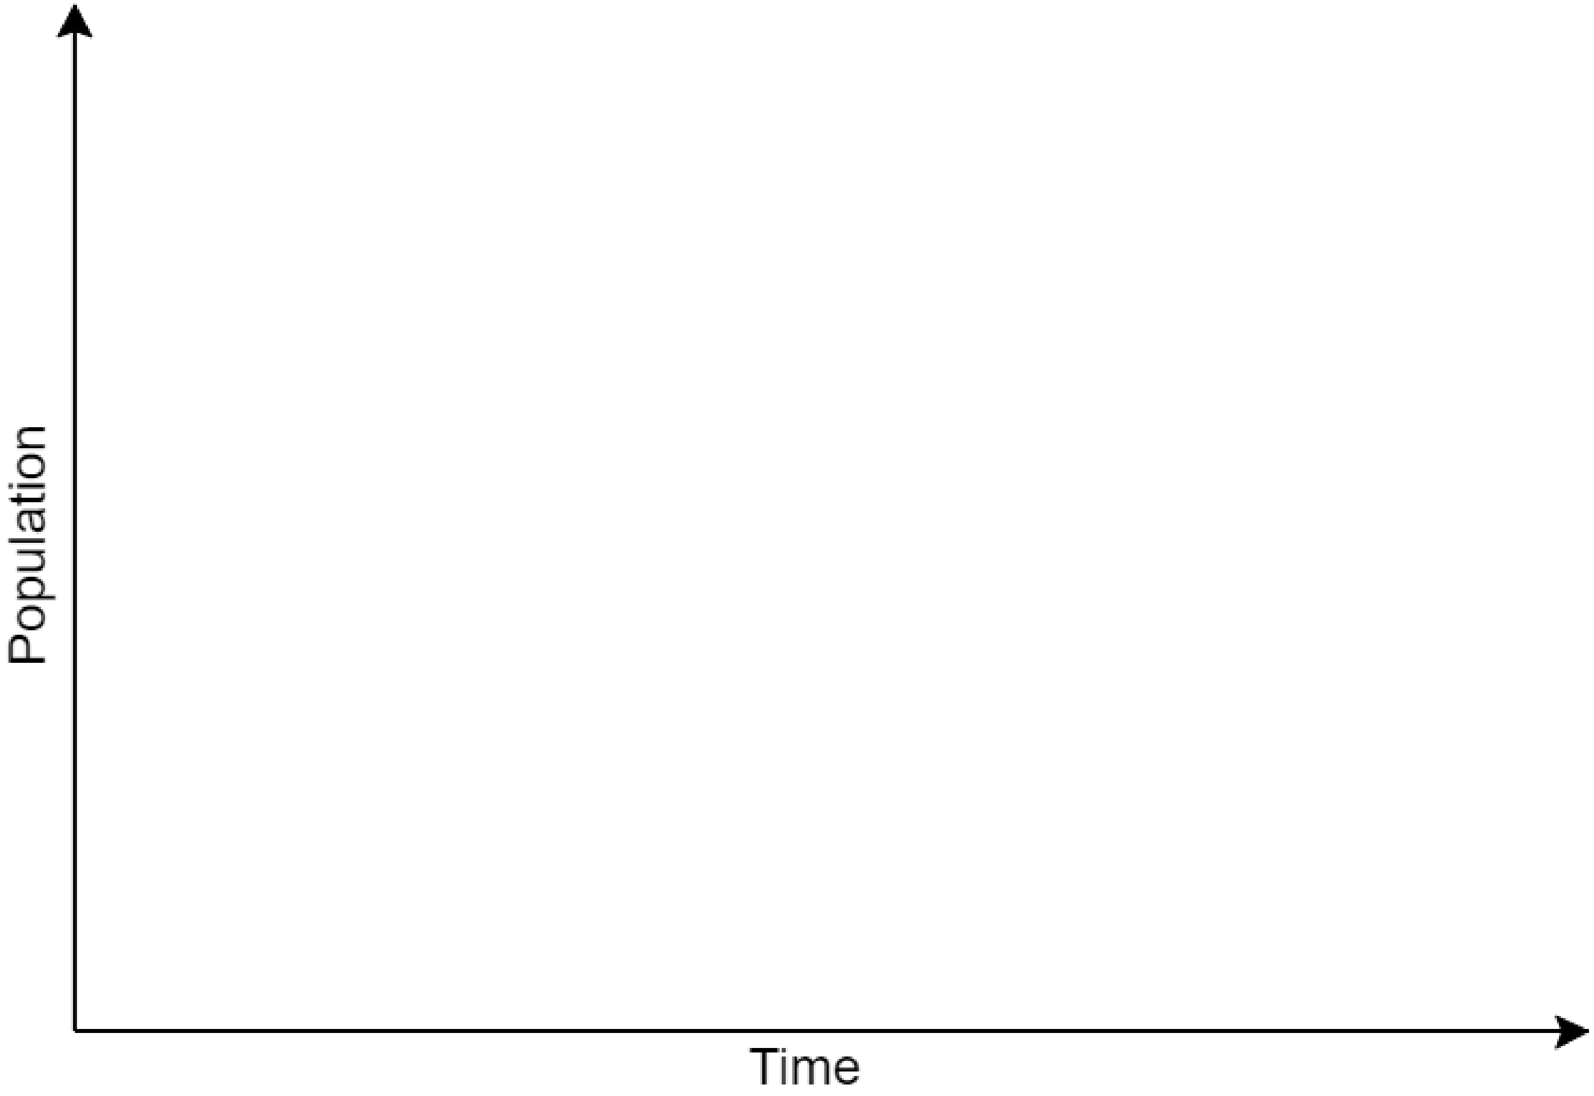

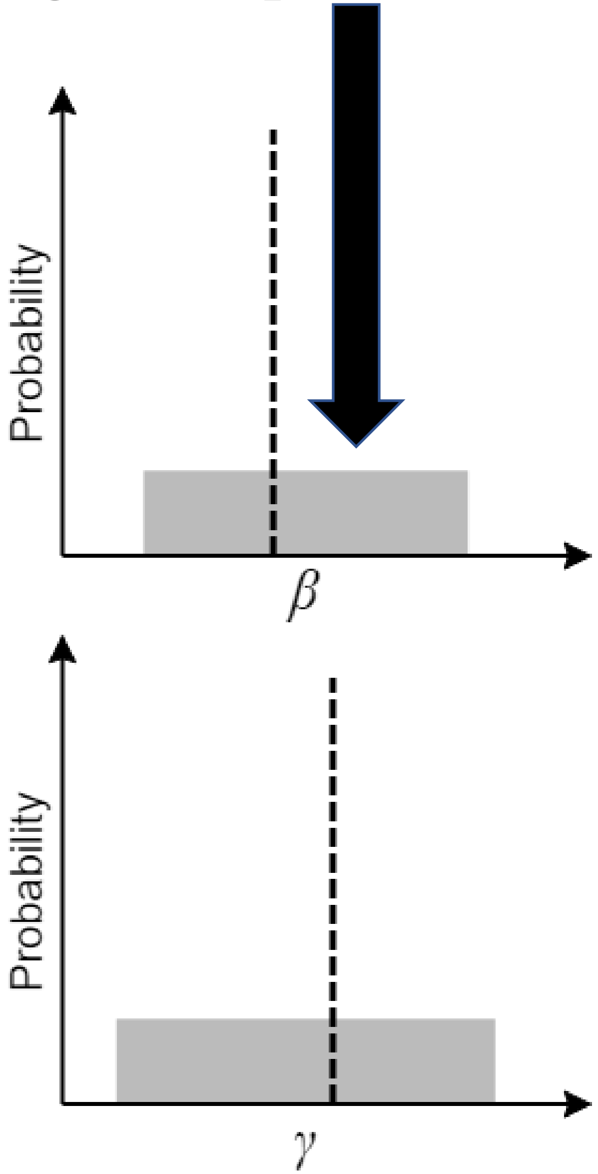

ABC rejection sampling

STEP 1: Sample θ from the prior $\pi(\theta)$

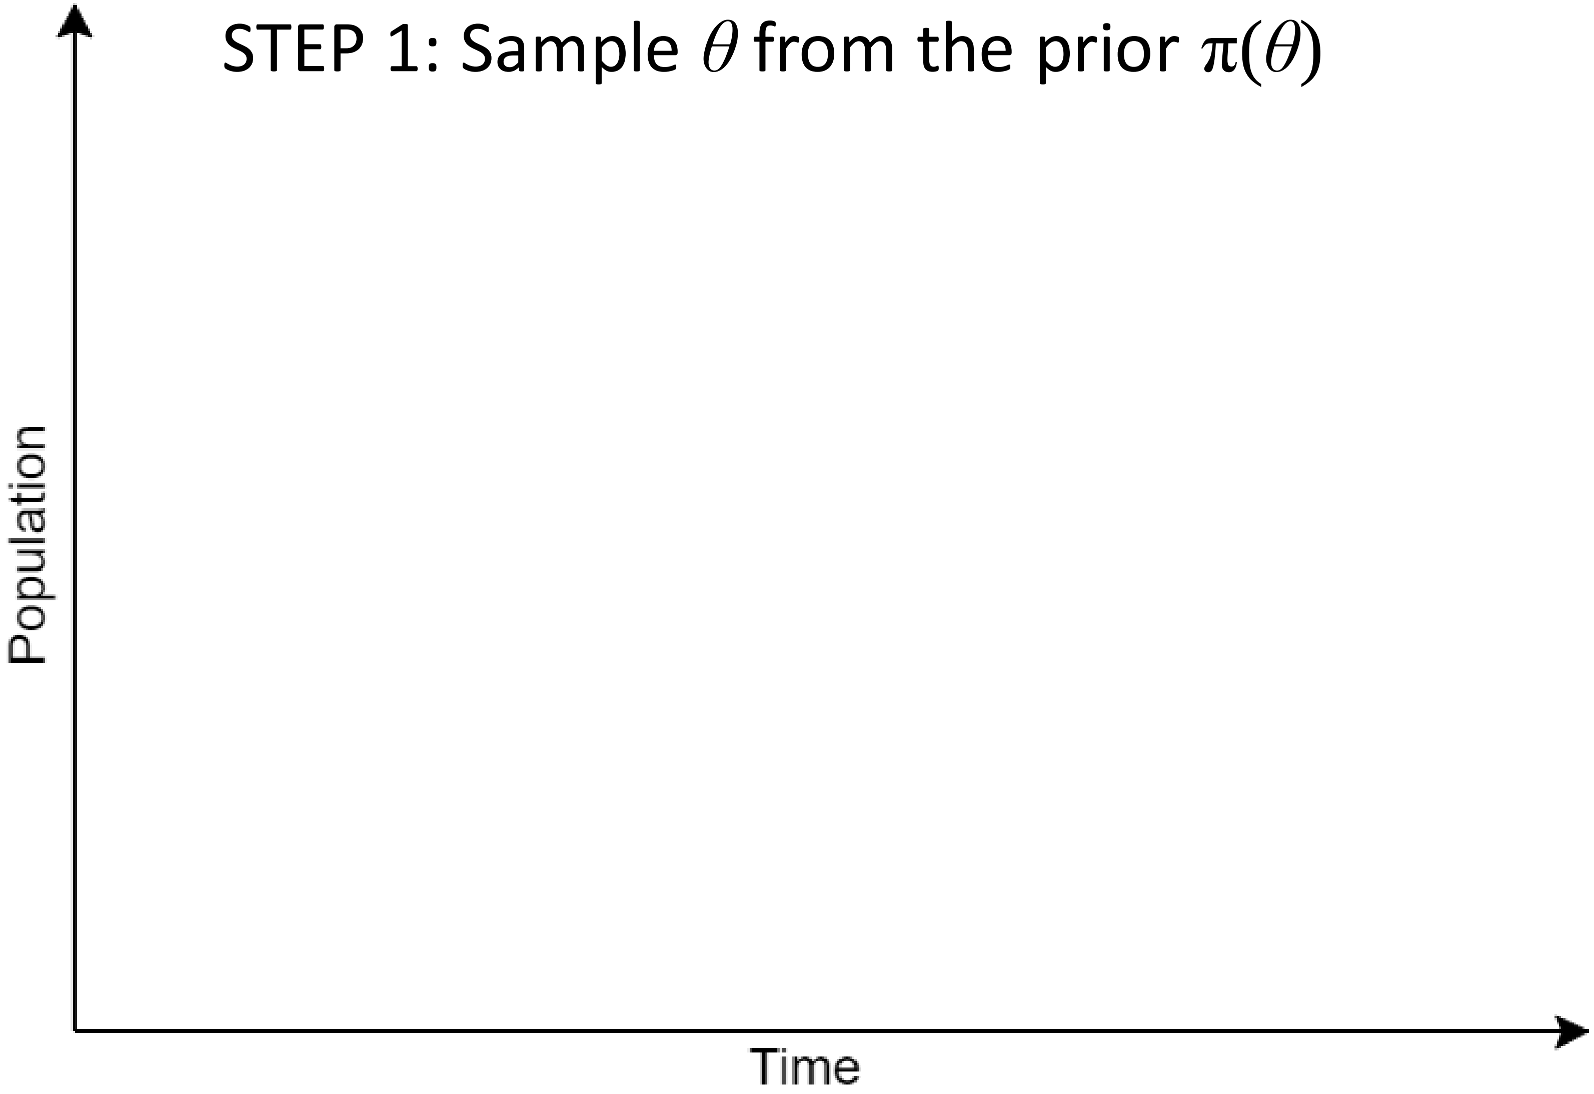

Values sampled (grey)

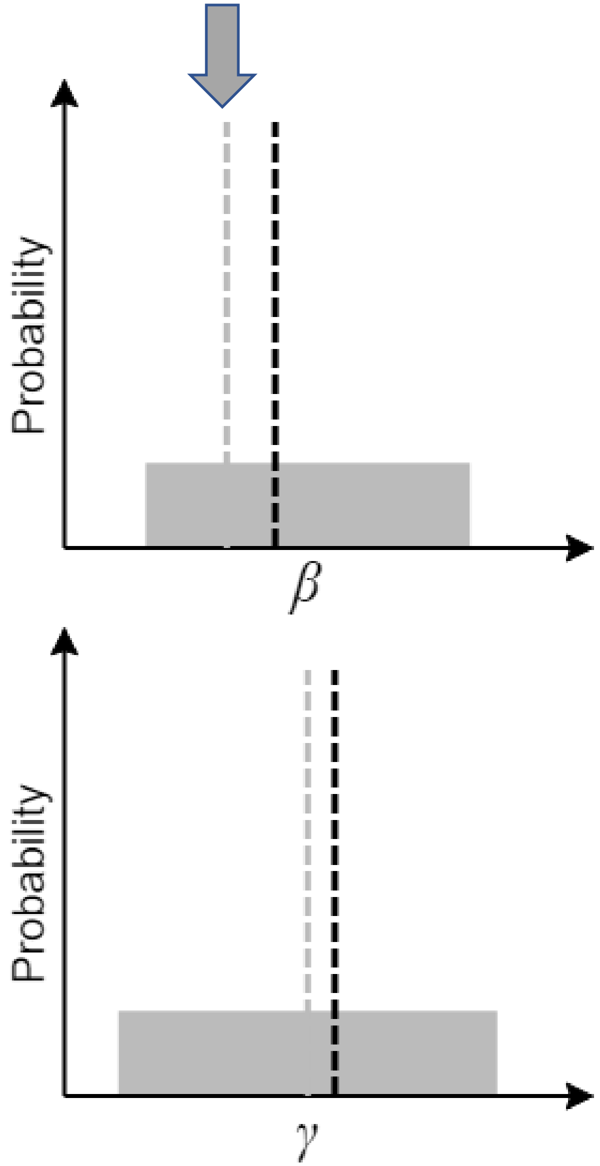

ABC rejection sampling

STEP 2: Simulate ξ from the model

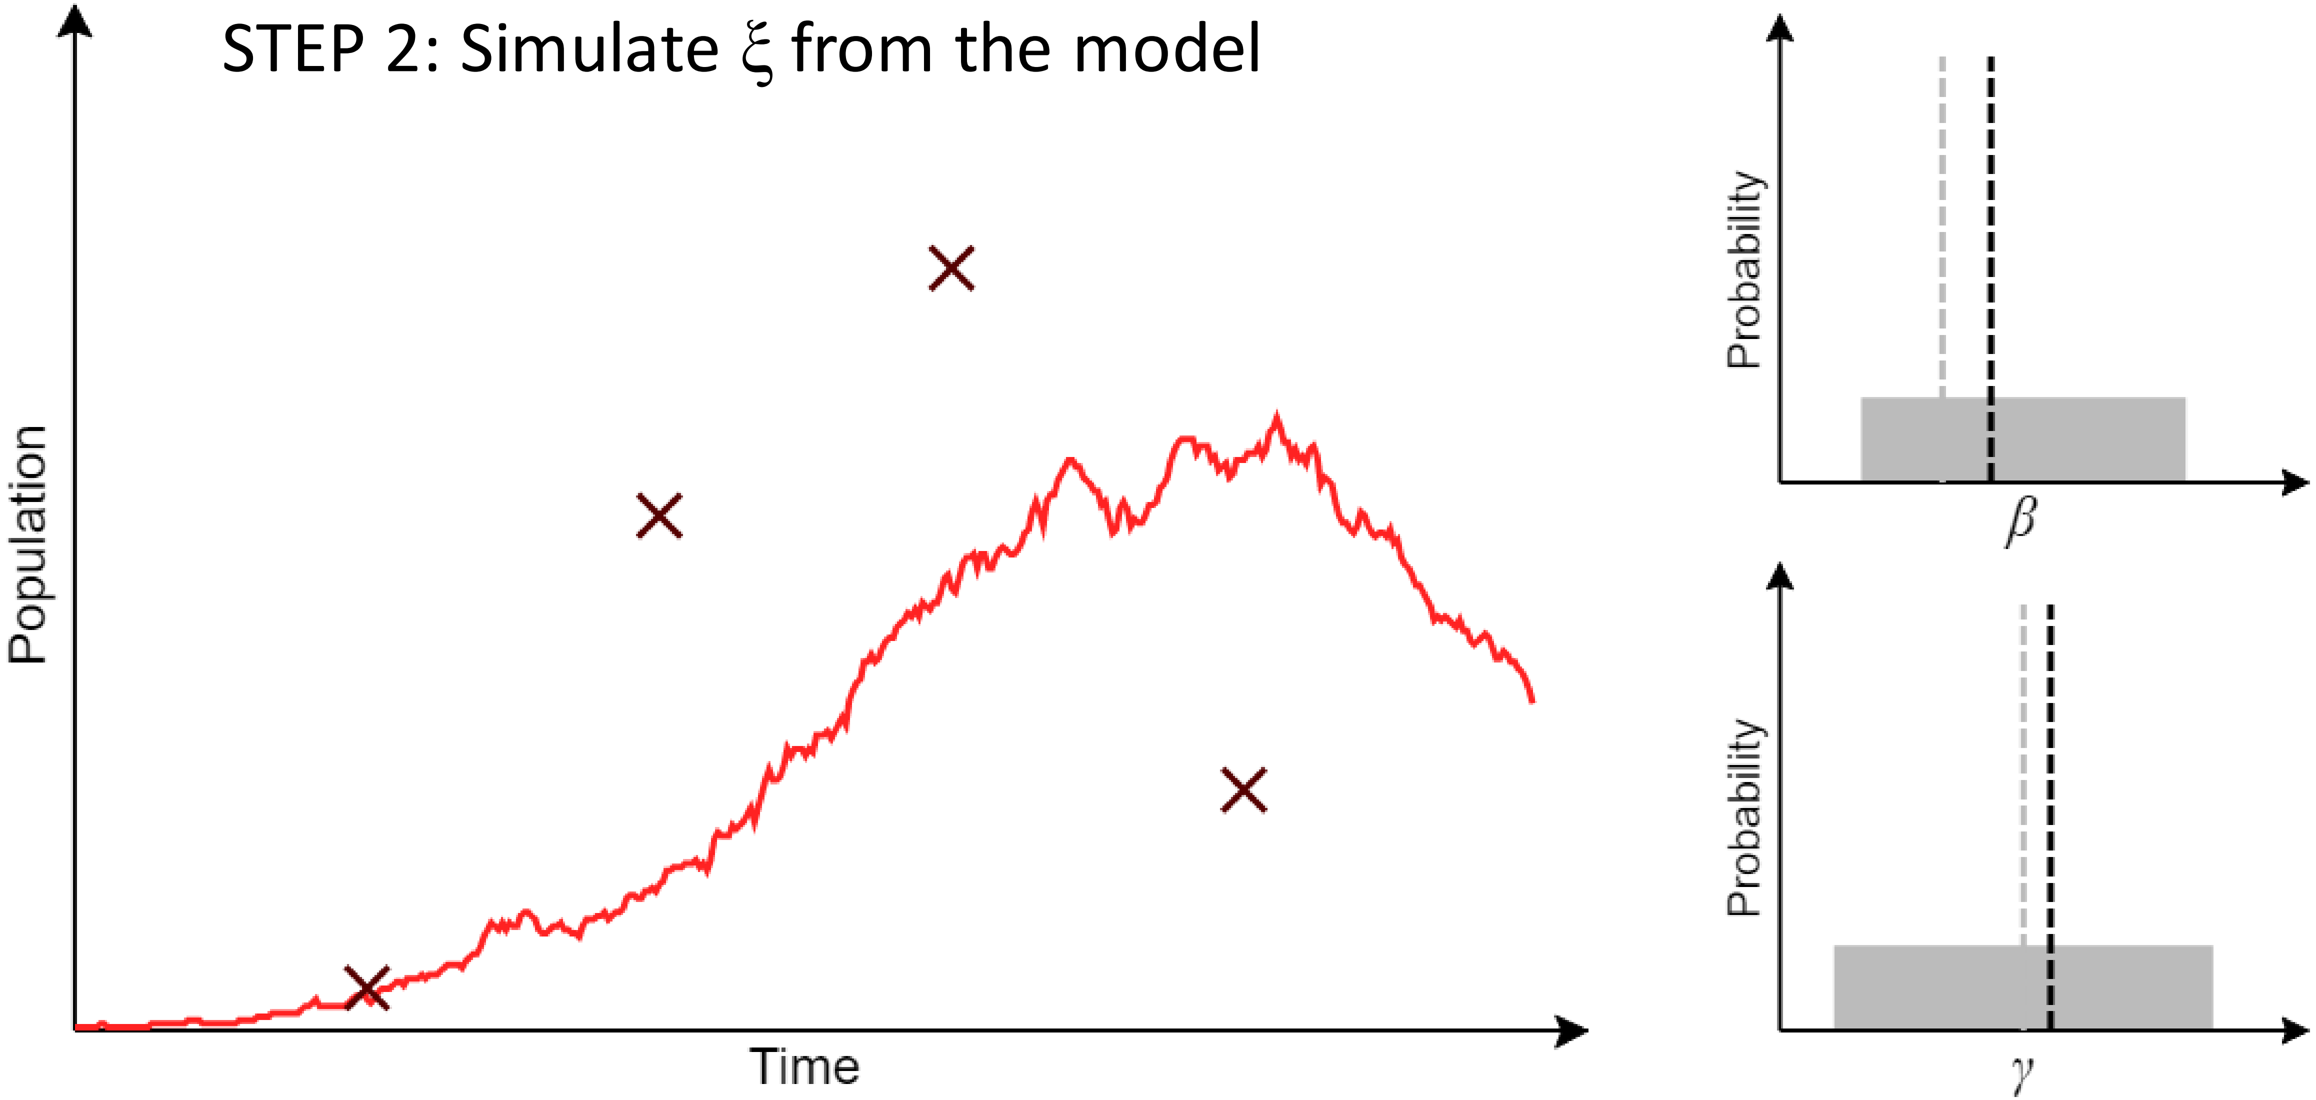

ABC rejection sampling

STEP 3: Calculate $EF = \sum_i (y_i - Y_i)^2$

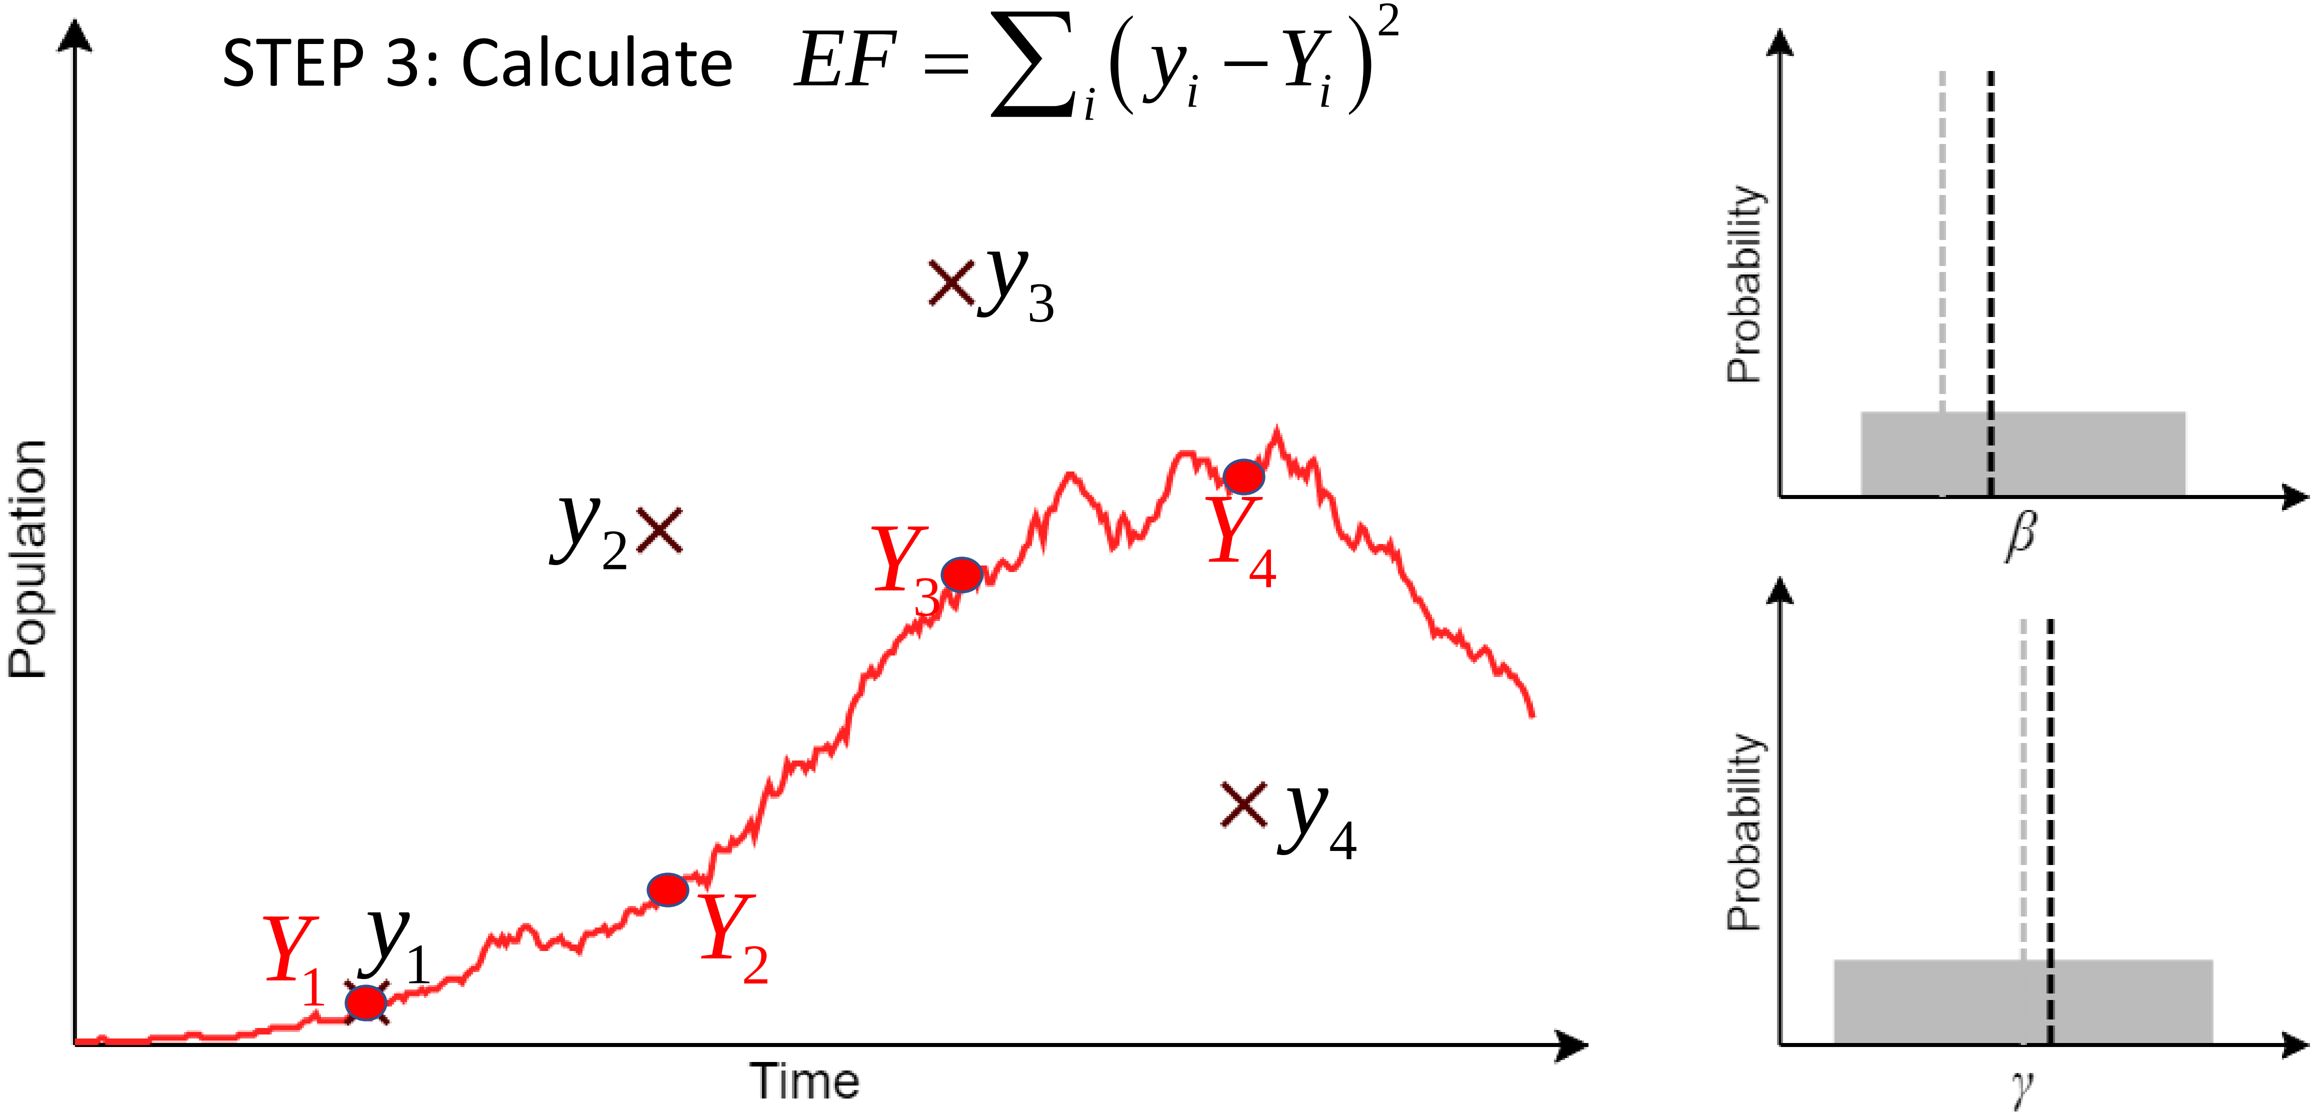

ABC rejection sampling

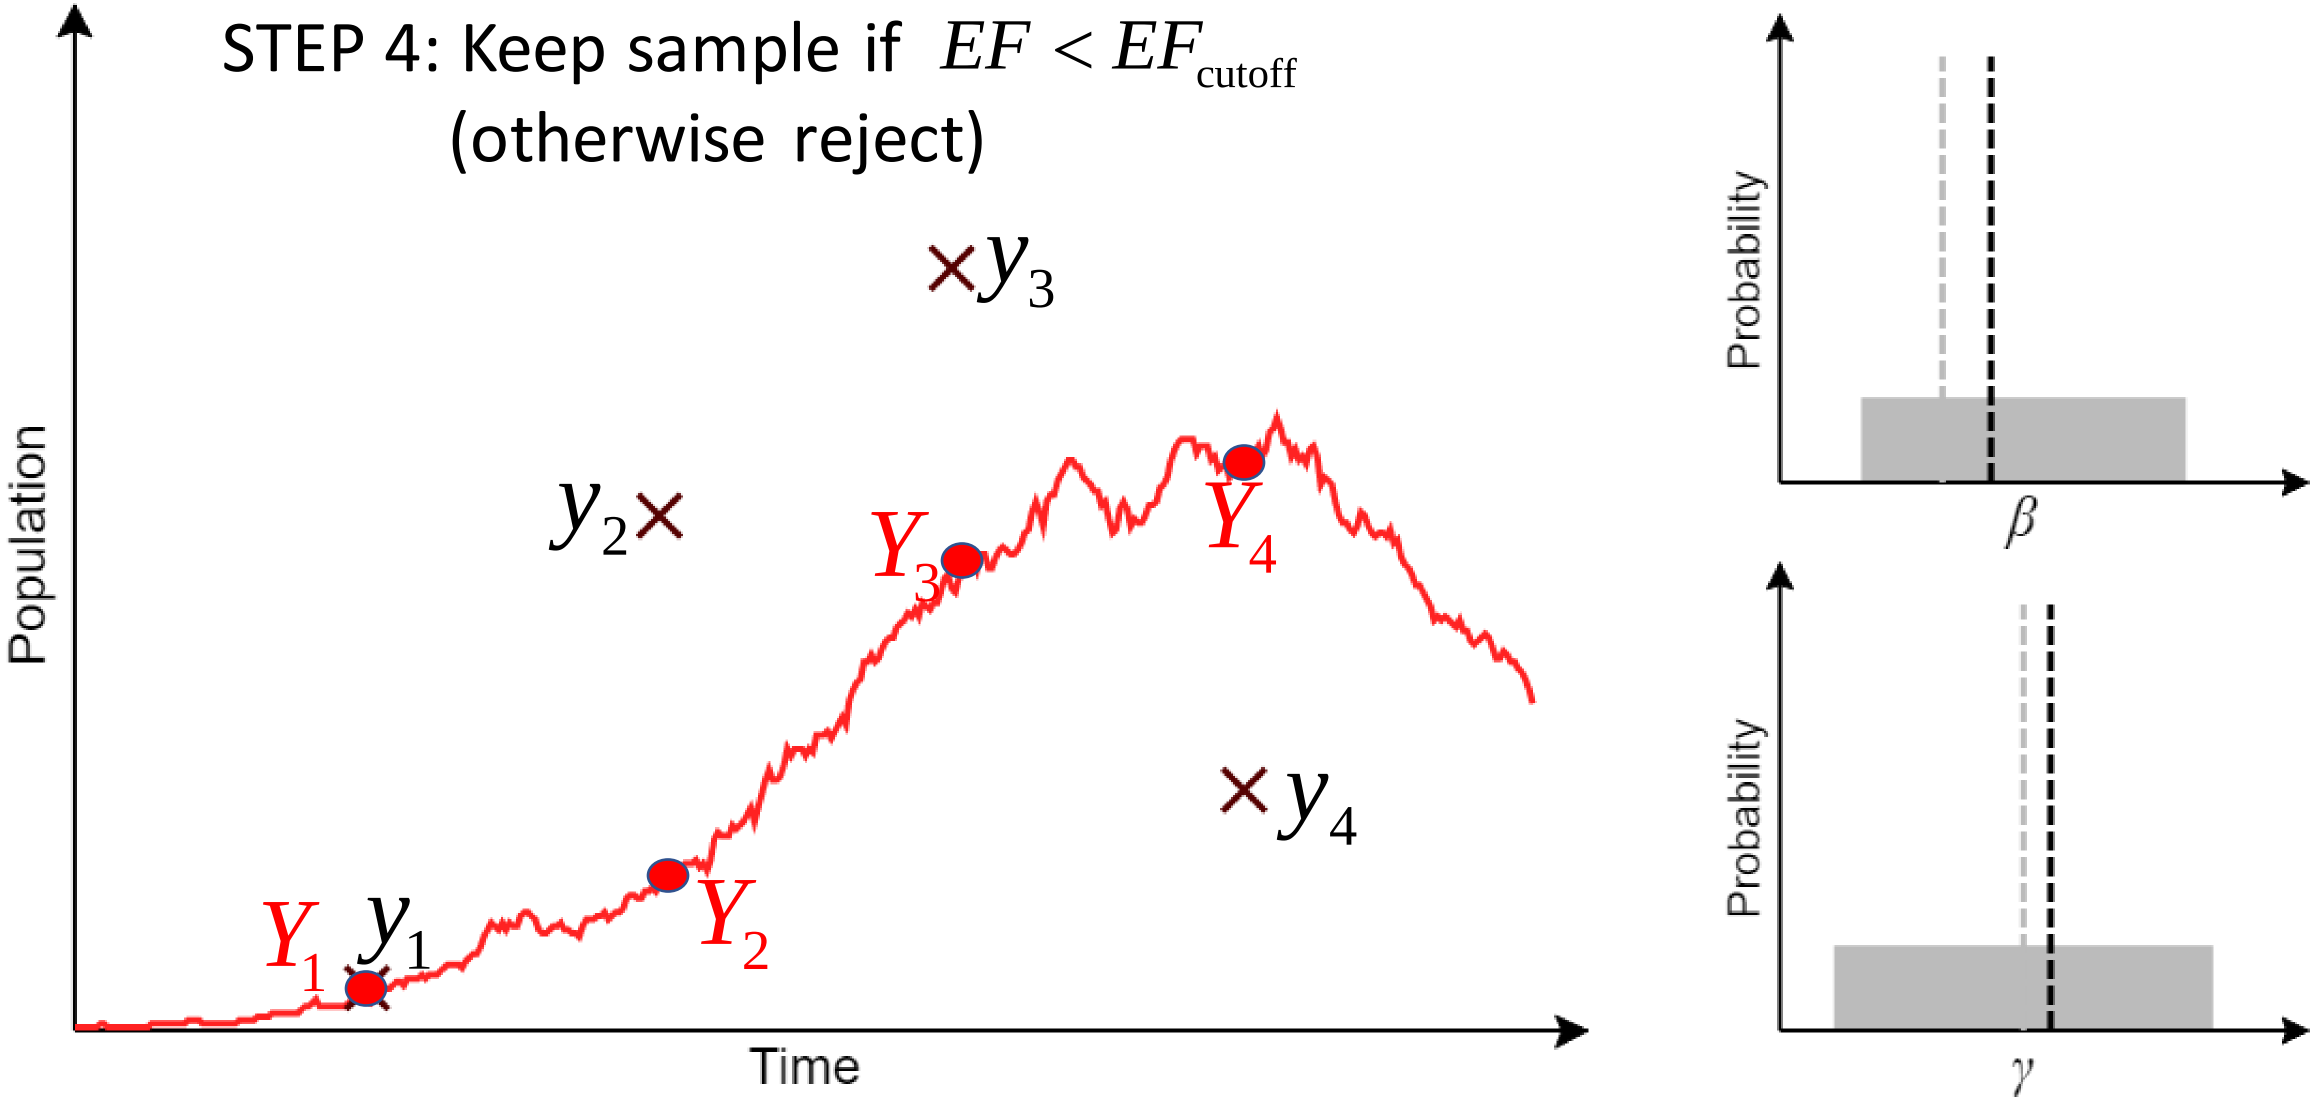

ABC rejection sampling

STEP 5: Repeat many times!

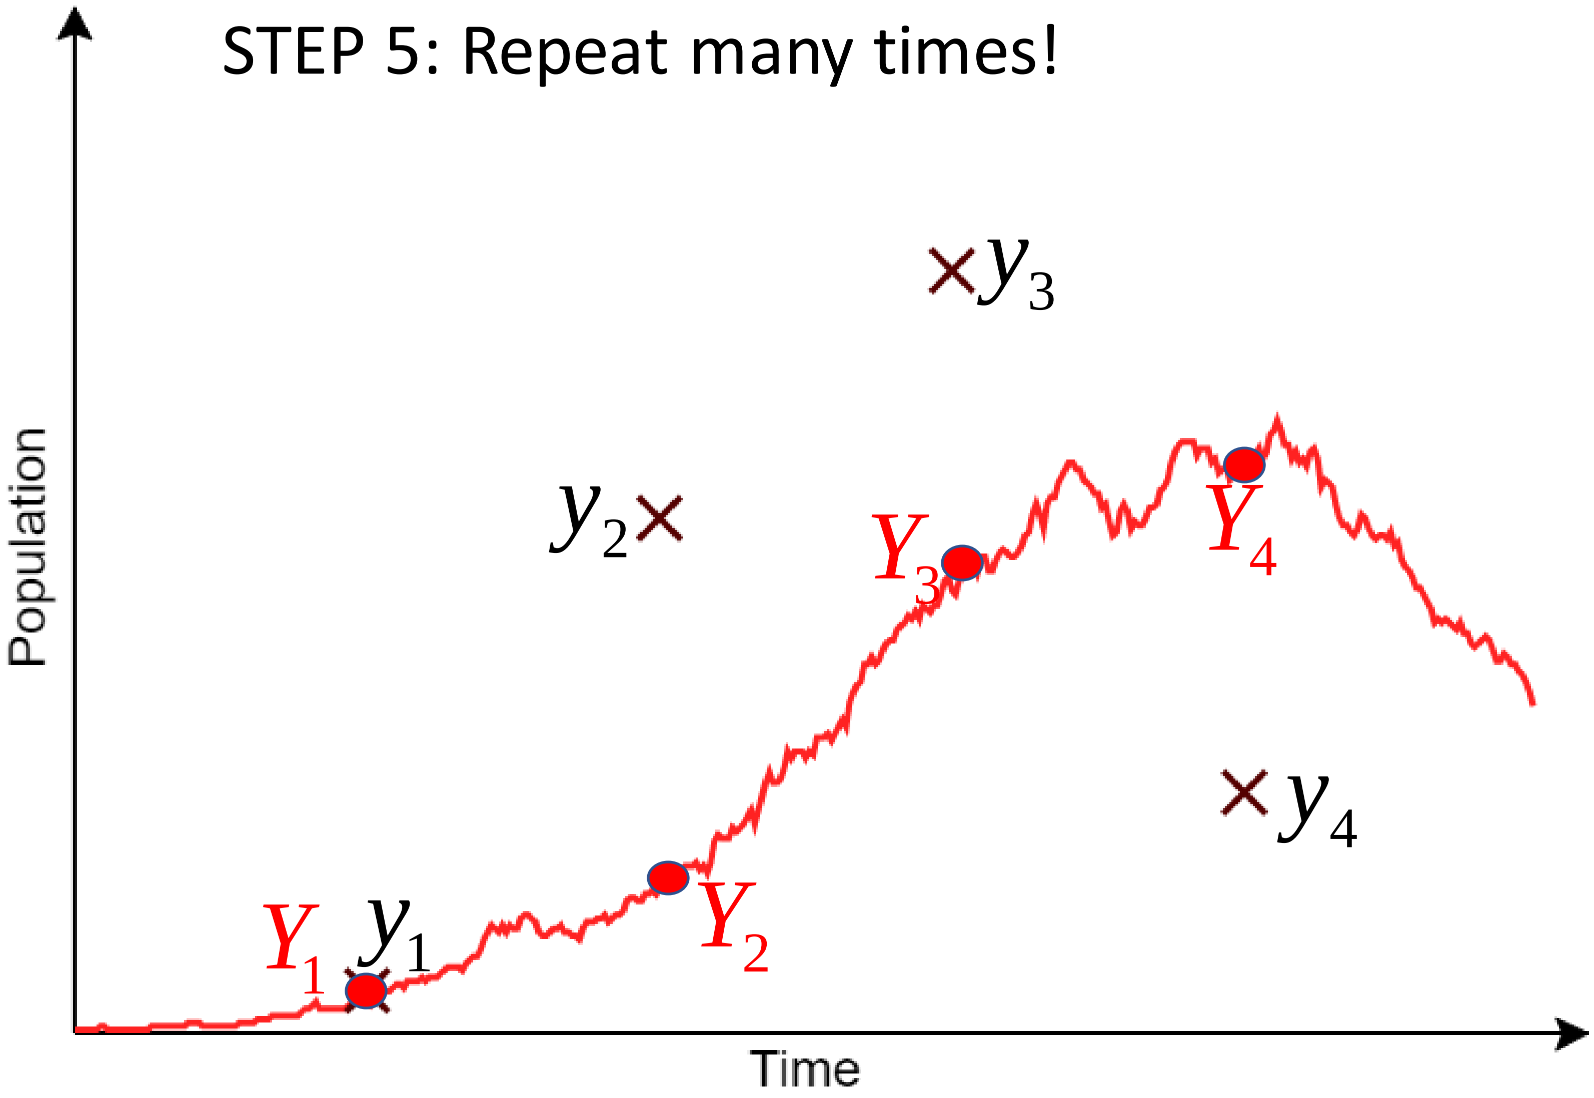

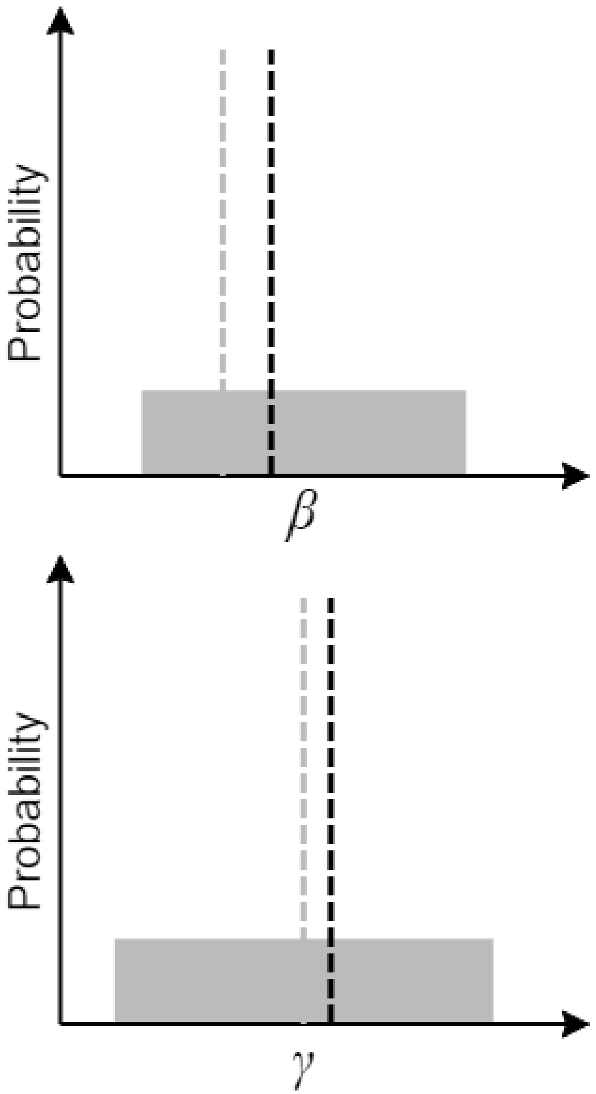

ABC rejection sampling

EXAMPLE: $EF_{\text{cutoff}} = 10^6$

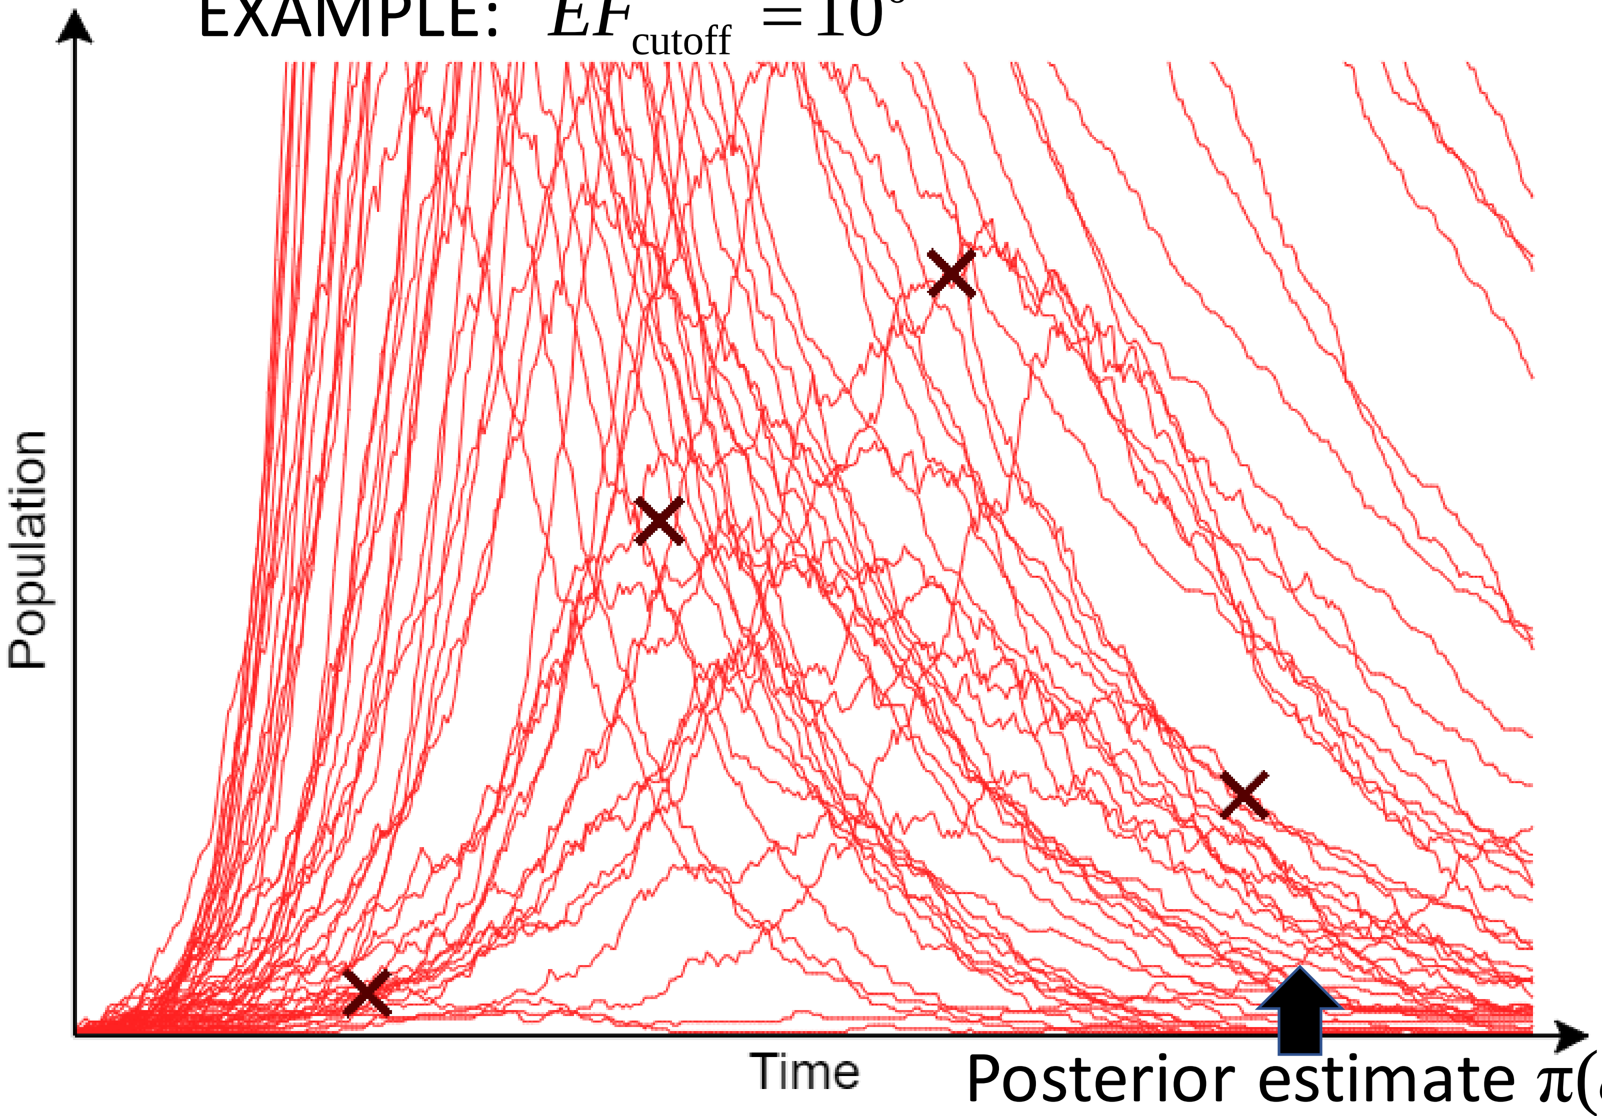

Posterior estimate $\pi(\xi|y)$

Posterior estimate $\pi(\theta|y)$

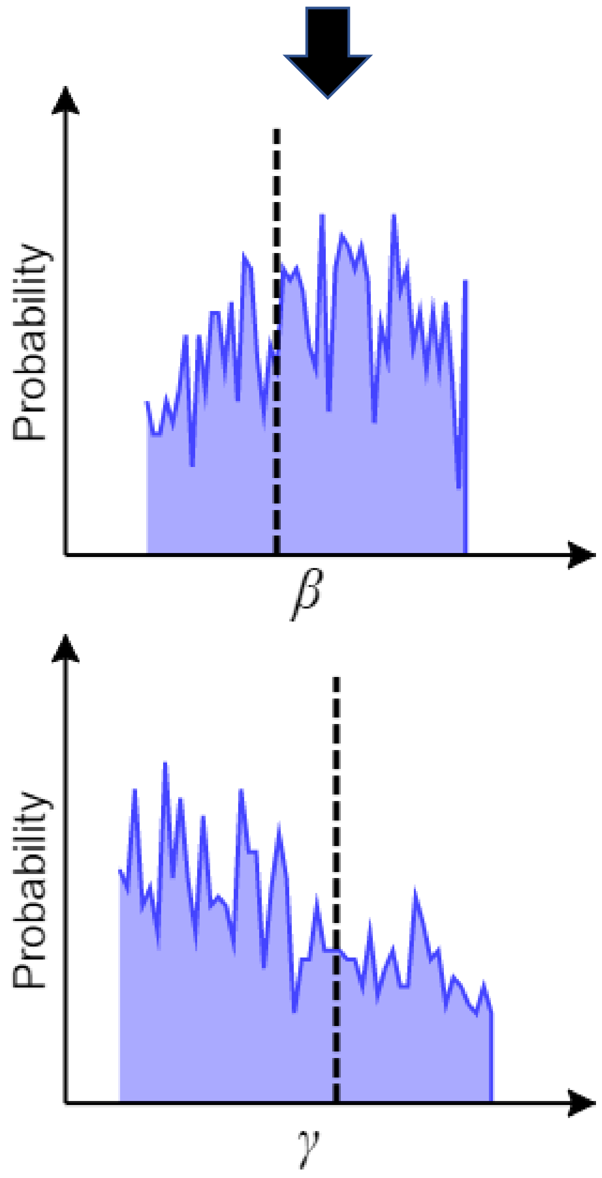

ABC rejection sampling

EXAMPLE: $EF_{\text{cutoff}} = 10^5$

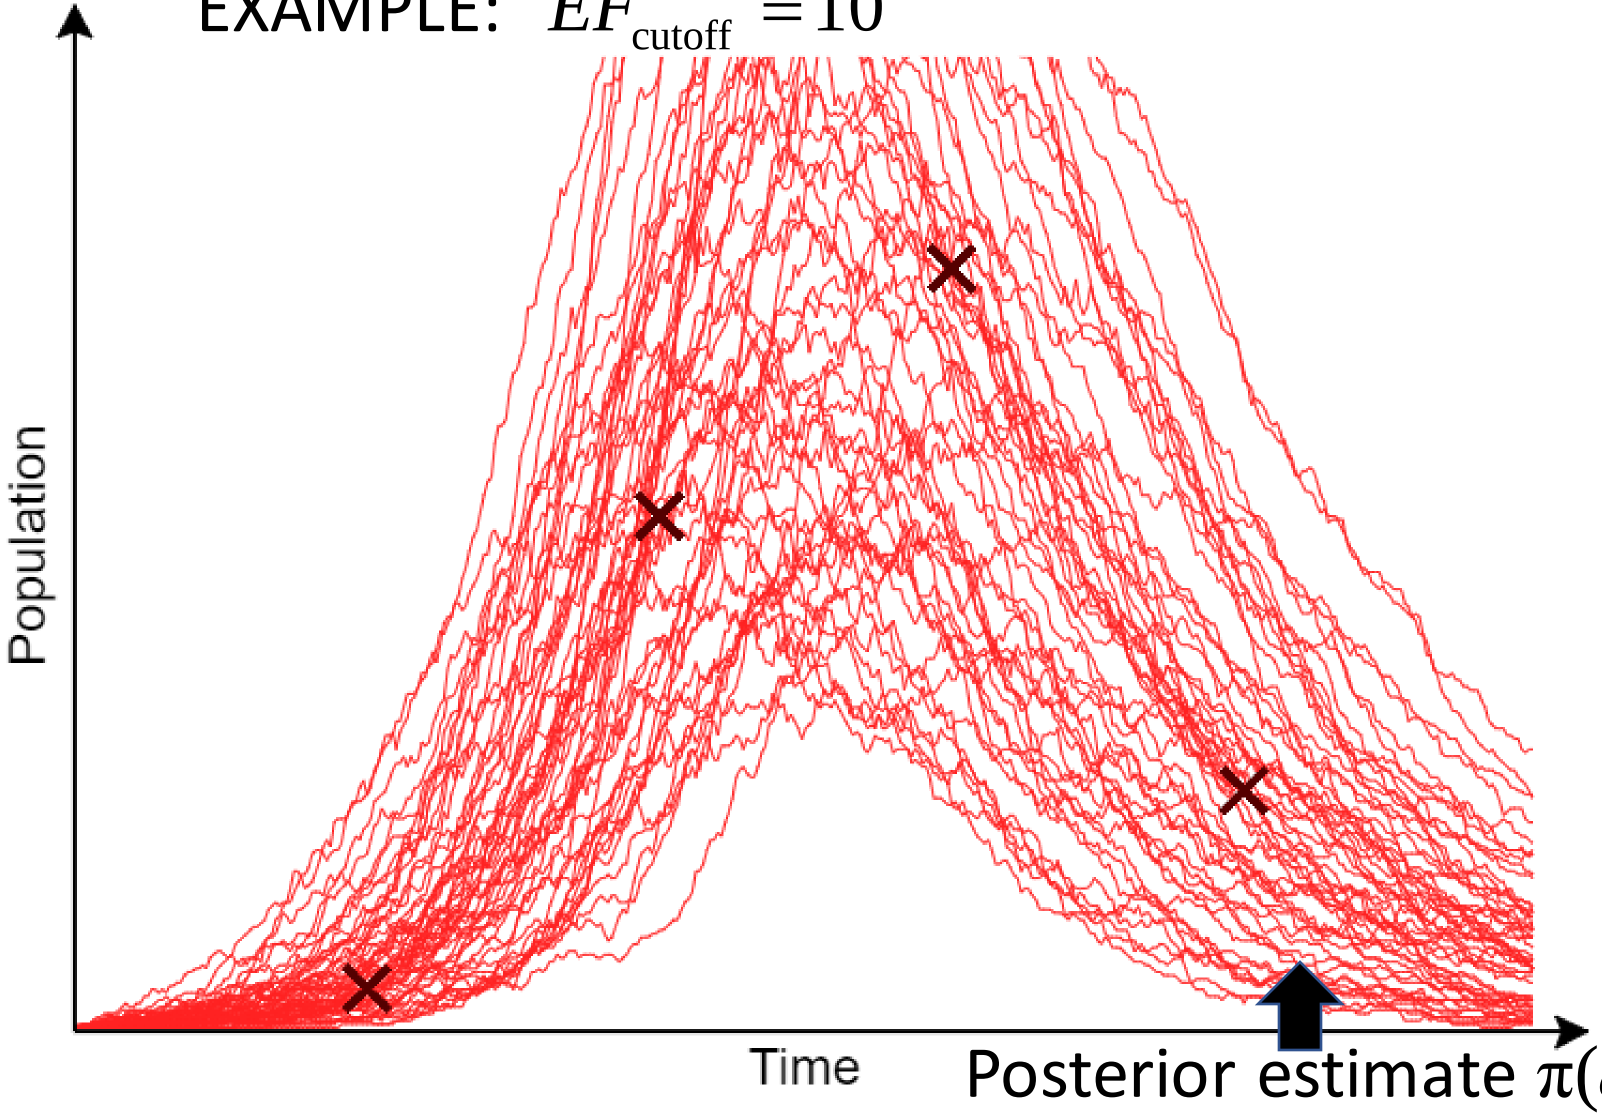

Posterior estimate $\pi(\xi|y)$

Posterior estimate $\pi(\theta|y)$

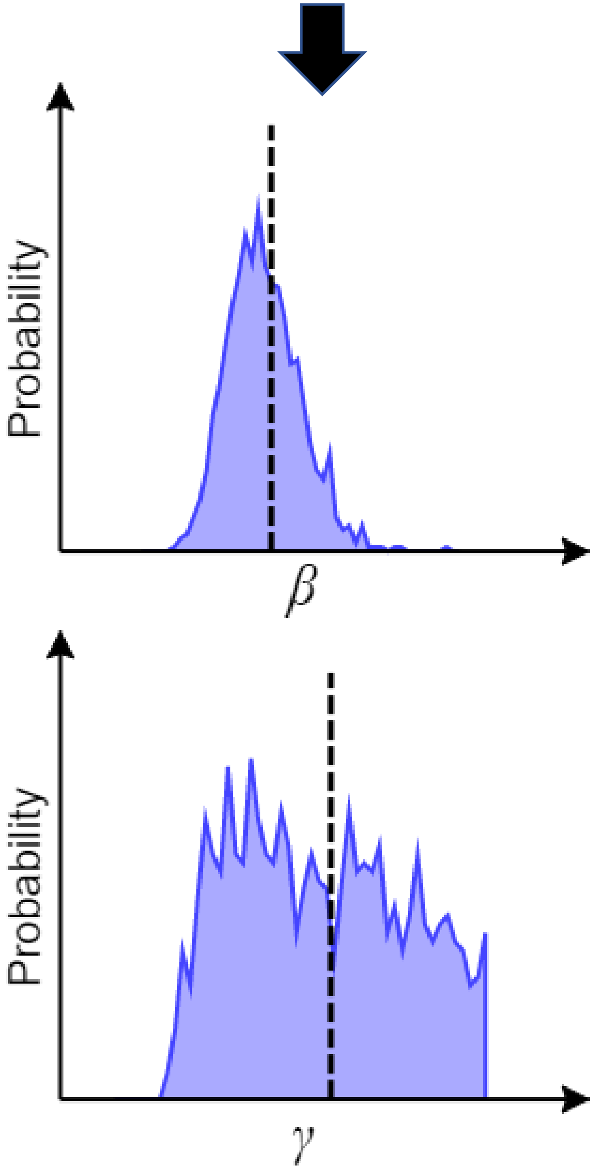

ABC rejection sampling

EXAMPLE: $EF_{\text{cutoff}} = 10^4$

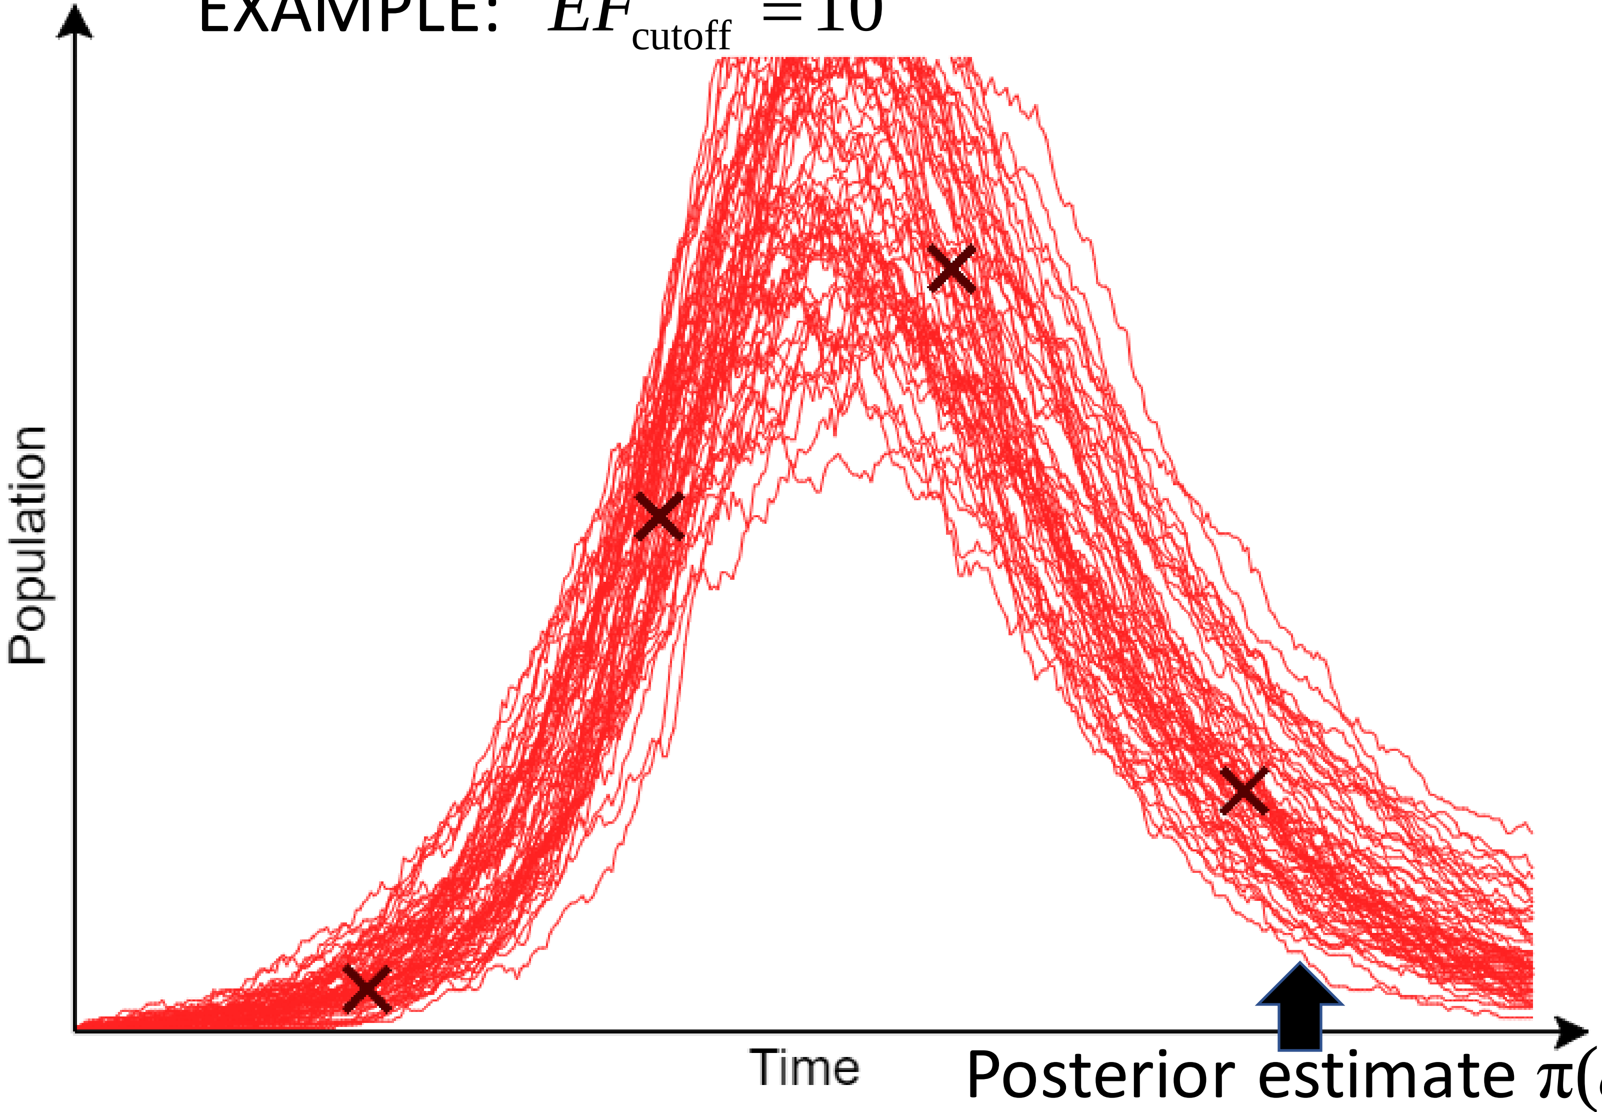

Posterior estimate $\pi(\xi|y)$

Posterior estimate $\pi(\theta|y)$

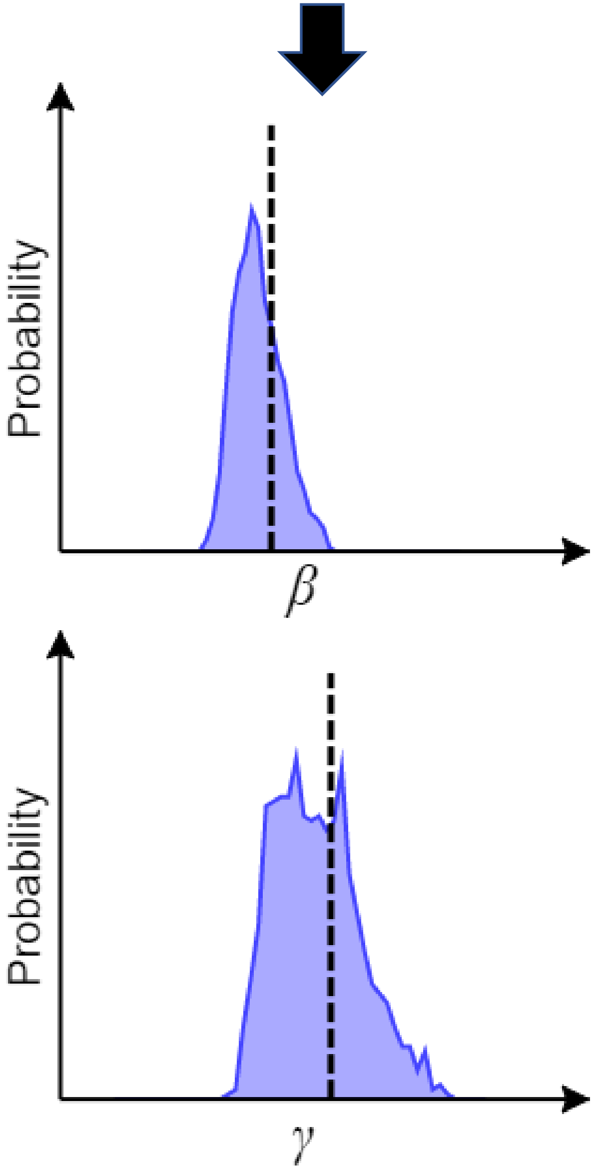

ABC rejection sampling

EXAMPLE: $EF_{\text{cutoff}} = 10^3$

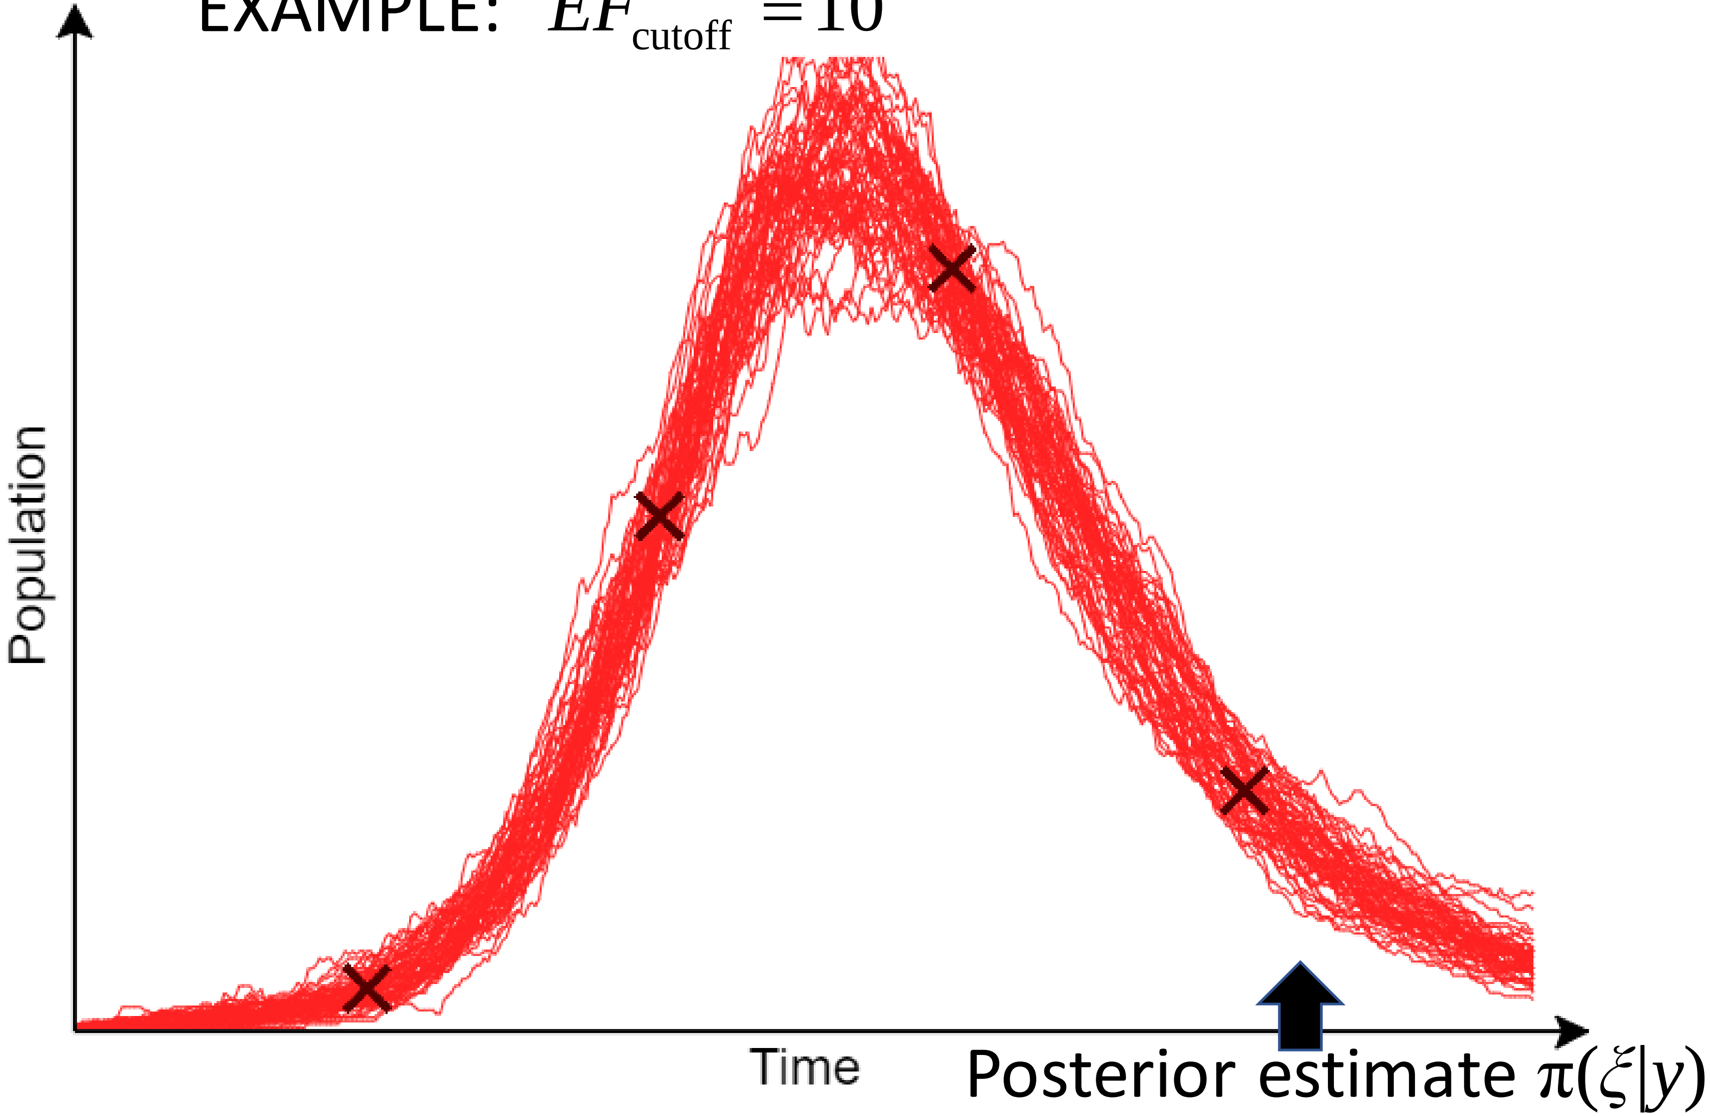

Posterior estimate $\pi(\theta|y)$

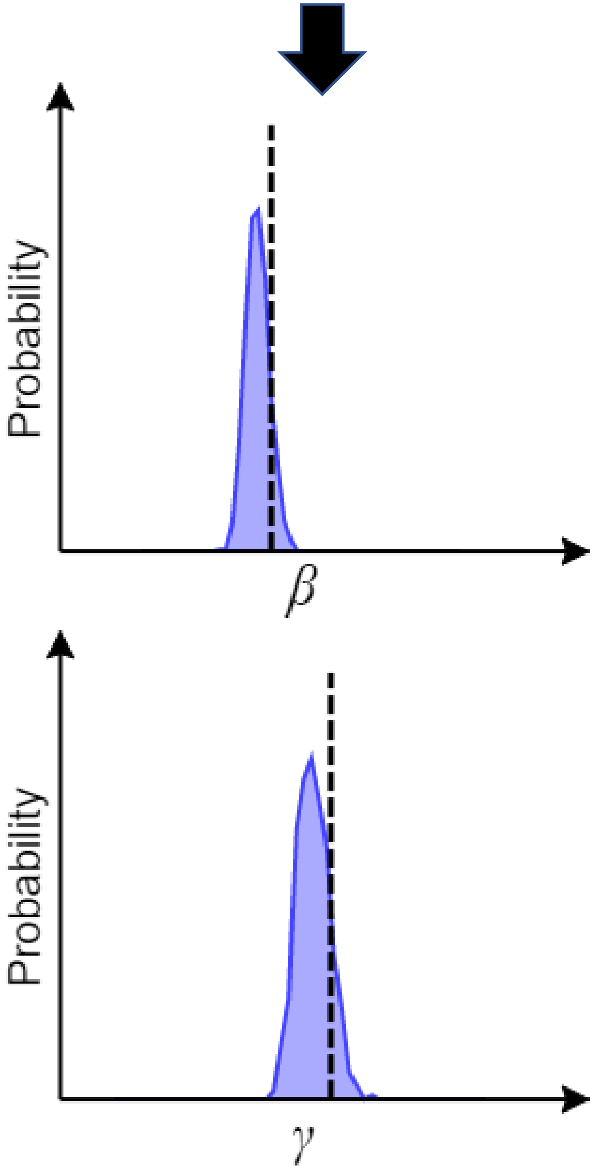

ABC rejection sampling

EXAMPLE: $EF_{\text{cutoff}} = 10^3$

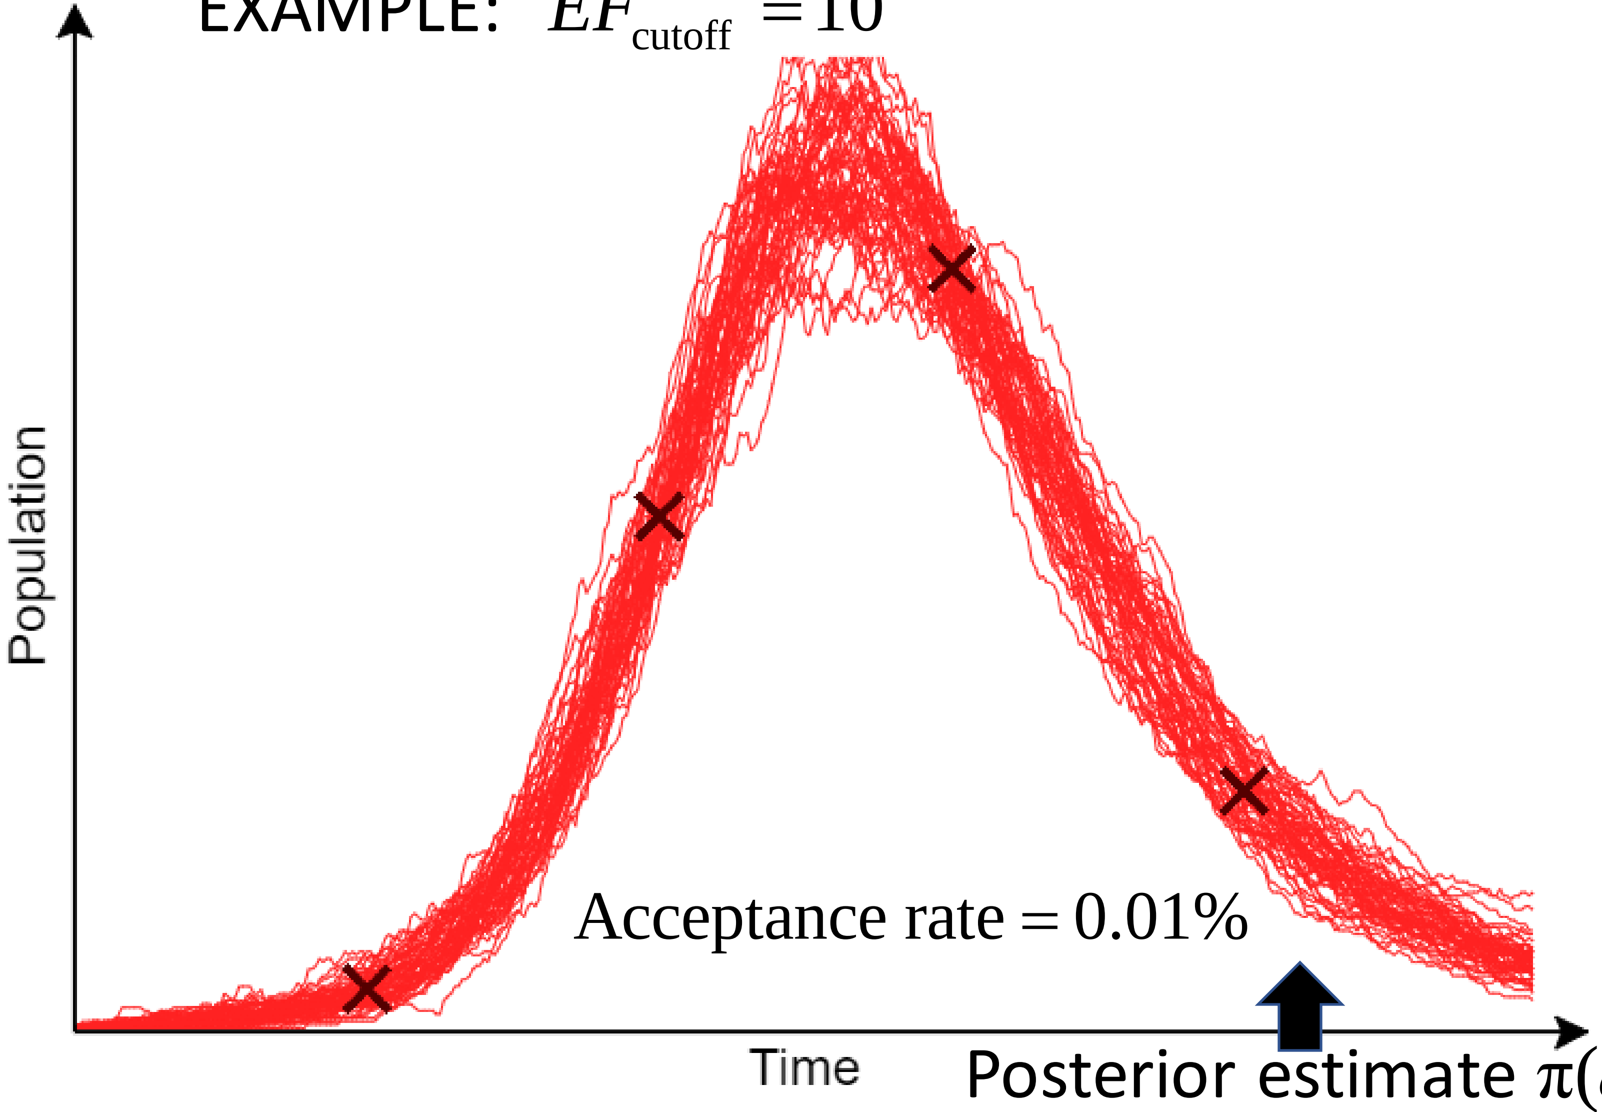

Acceptance rate $= 0.01\%$

Posterior estimate $\pi(\xi|y)$

Posterior estimate $\pi(\theta|y)$

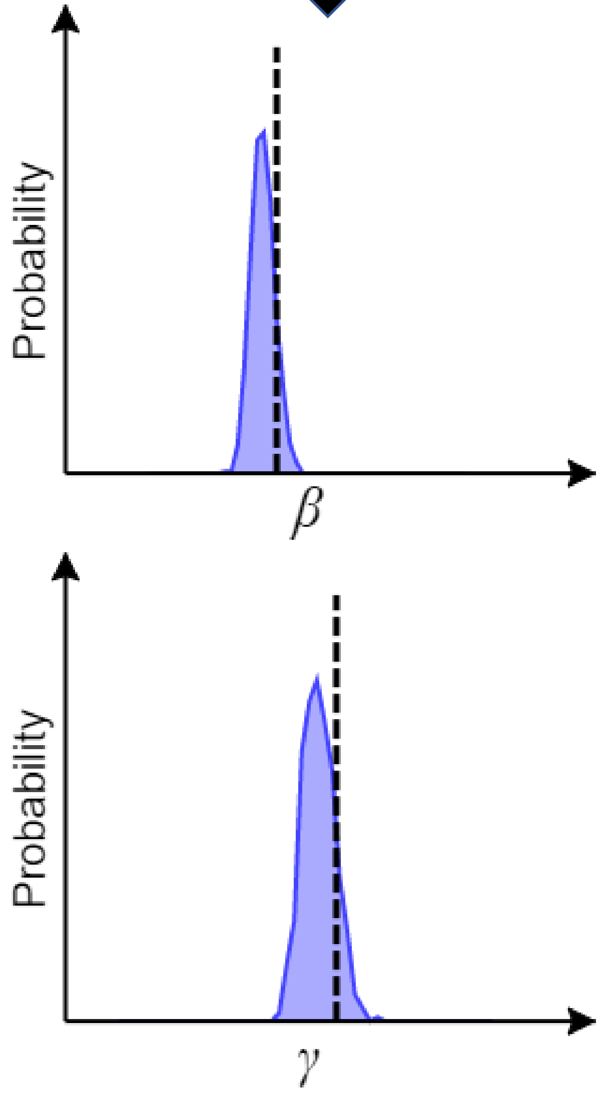

ABC rejection sampling

- Very simple to implement
- Only need to be able to simulate from the model

- Becomes computationally very slow for small EF_{cutoff}
  - For complex system not possible to get good posterior estimate

- WHY?
  - Because most samples from prior $\pi(\theta)$ have a low posterior probability $\pi(\theta|y)$
  - This is the motivation behind ABC-SMC…

ABC Sequential Monte Carlo (ABC-SMC)

- Run over several generations G
- For each generation g a different EF^{g}_{cutoff} is used, such that

$$EF^{1}_{\text{cutoff}} > EF^{2}_{\text{cutoff}} > \ldots > EF^{G}_{\text{cutoff}}$$

Generation 1

1) Sample θ from prior $\pi(\theta)$.
2) Simulate ξ.
3) Calculate EF.
4) If $EF < EF^{1}_{\text{cutoff}}$ accept.

Repeat.

Generation *g*

1) Sample θ using previous generation outputs as an importance sampler (with associated weight).
2) Simulate ξ.
3) Calculate EF.
4) If $EF < EF^{g}_{\text{cutoff}}$ accept.

Repeat.

$g = g + 1$

ABC Sequential Monte Carlo (ABC-SMC)

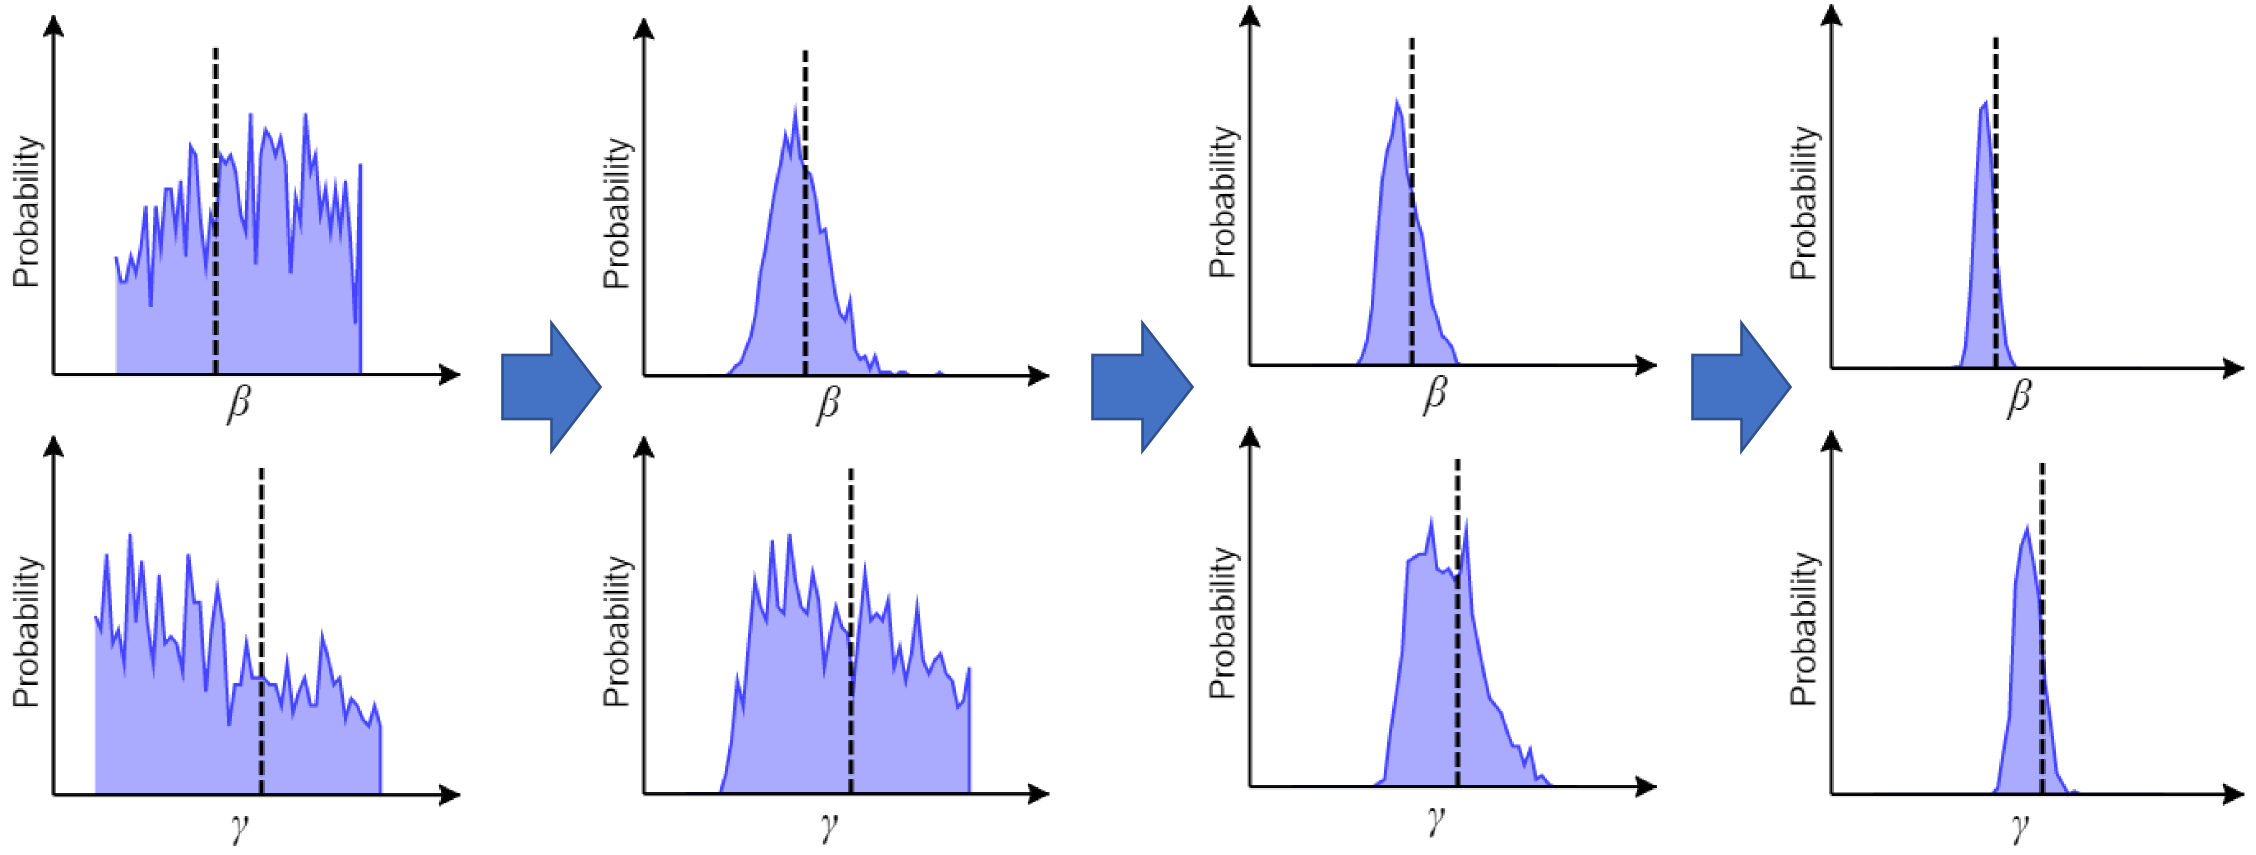

	Generation 1	**Generation 2**	**Generation 3**	**Generation 4**
$\boldsymbol{EF}_{\textbf{cutoff}}$	$EF^{1}_{\text{cutoff}} = 10^{6}$	$EF^{2}_{\text{cutoff}} = 10^{5}$	$EF^{3}_{\text{cutoff}} = 10^{4}$	$EF^{4}_{\text{cutoff}} = 10^{3}$
Acceptance:	38%	33%	20%	10%

ABC Sequential Monte Carlo (ABC-SMC)

- Relatively simple to implement
  - Requires careful optimisation of EF^{g}_{cutoff}
- Only need to be able to simulate from the system
- Usually much faster than ABC rejection sampling

- Can become computationally slow for small EF_{cutoff}
- WHY?

  1) Large number of parameters
     - Importance weights results in small effective sample size

  2) Significant stochasticity
     - Even correct parameters leads to a poor probability for simulating the data

ABC-MBP

STEP 1: Sample θ_p from the prior $\pi(\theta)$ for p=1…P particles

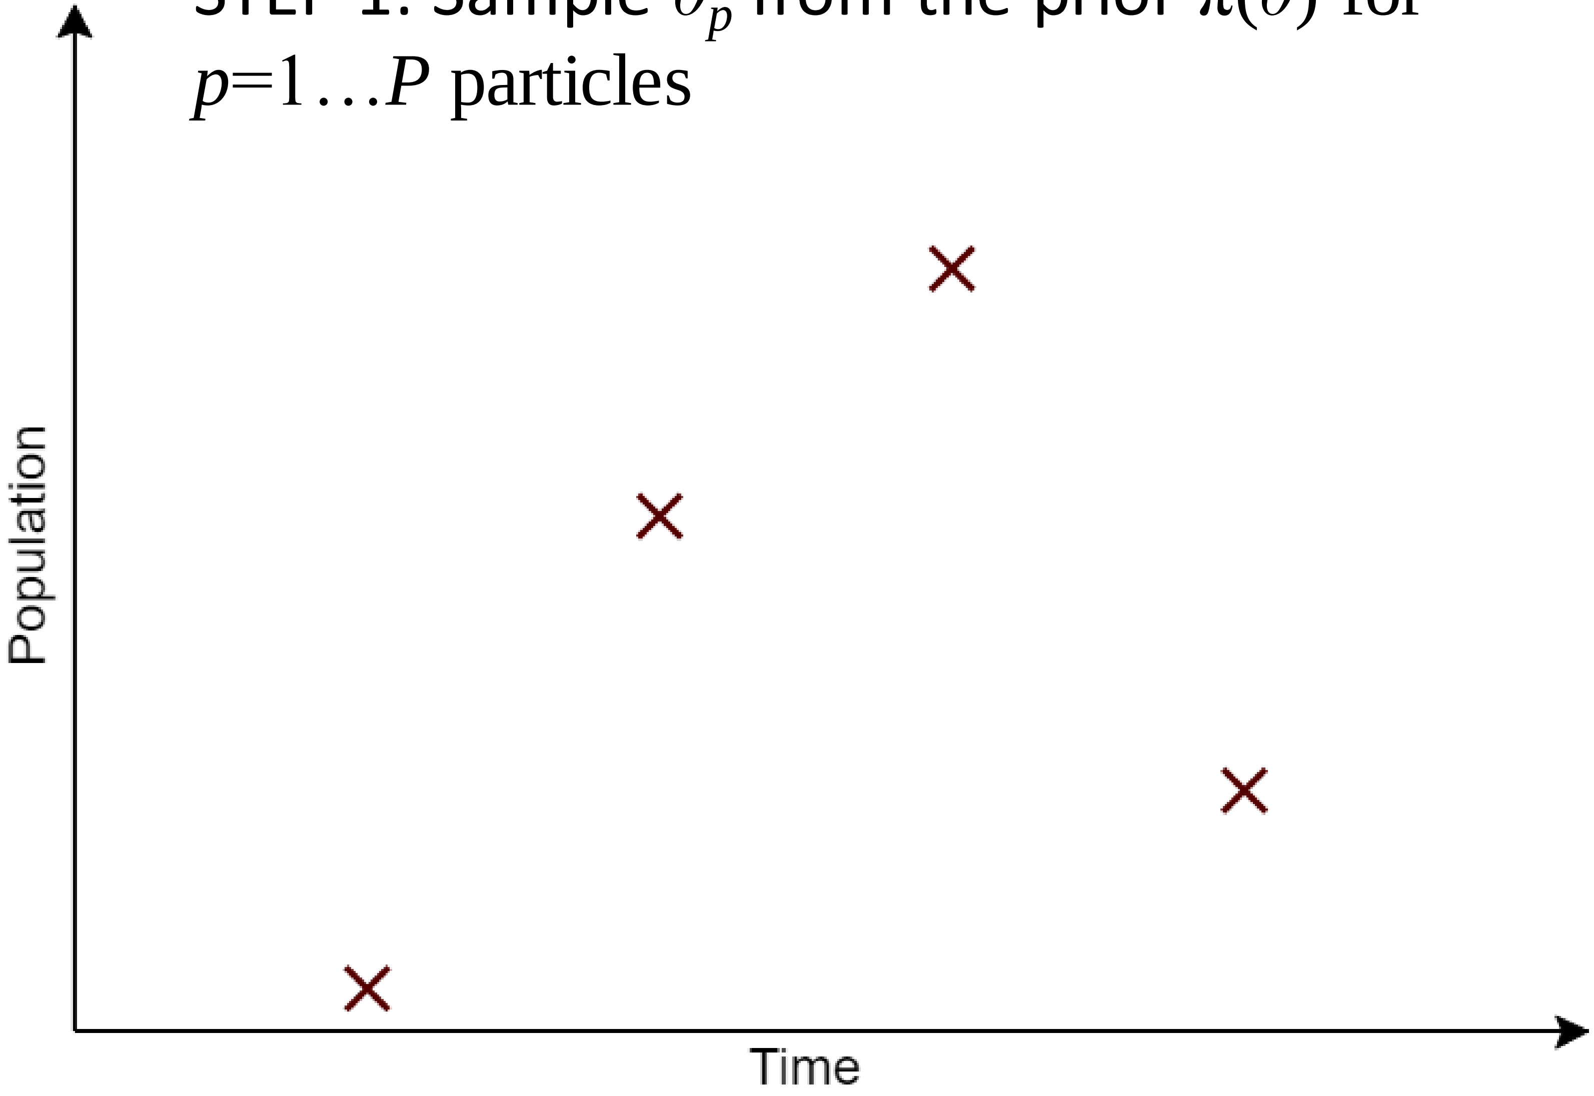

Sample from prior (grey)

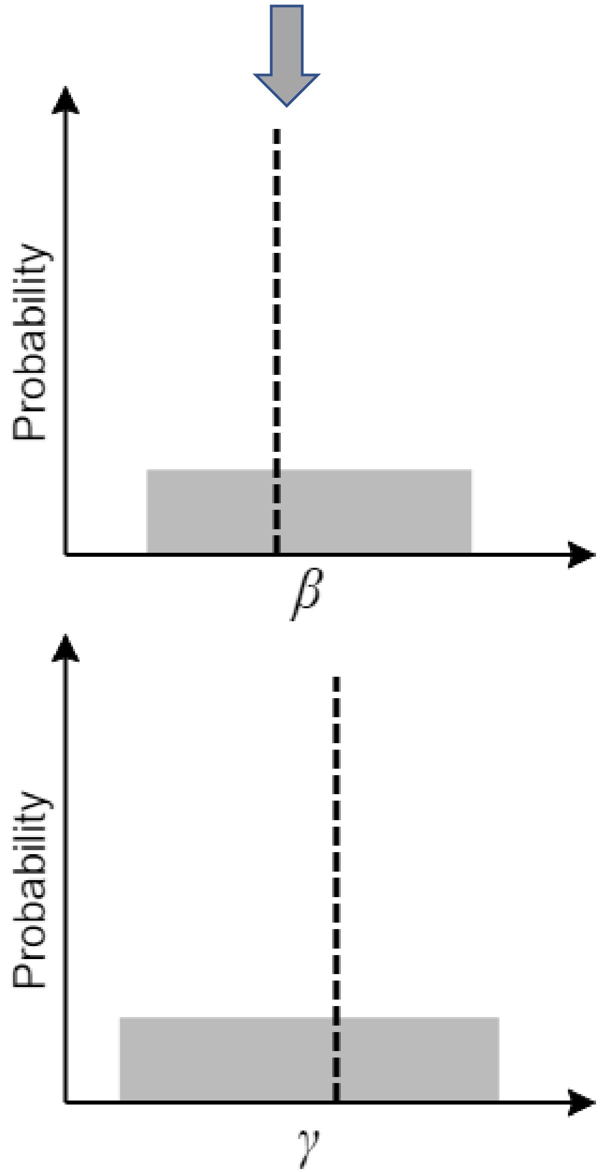

ABC-MBP

STEP 2: Simulate ξ_p from model

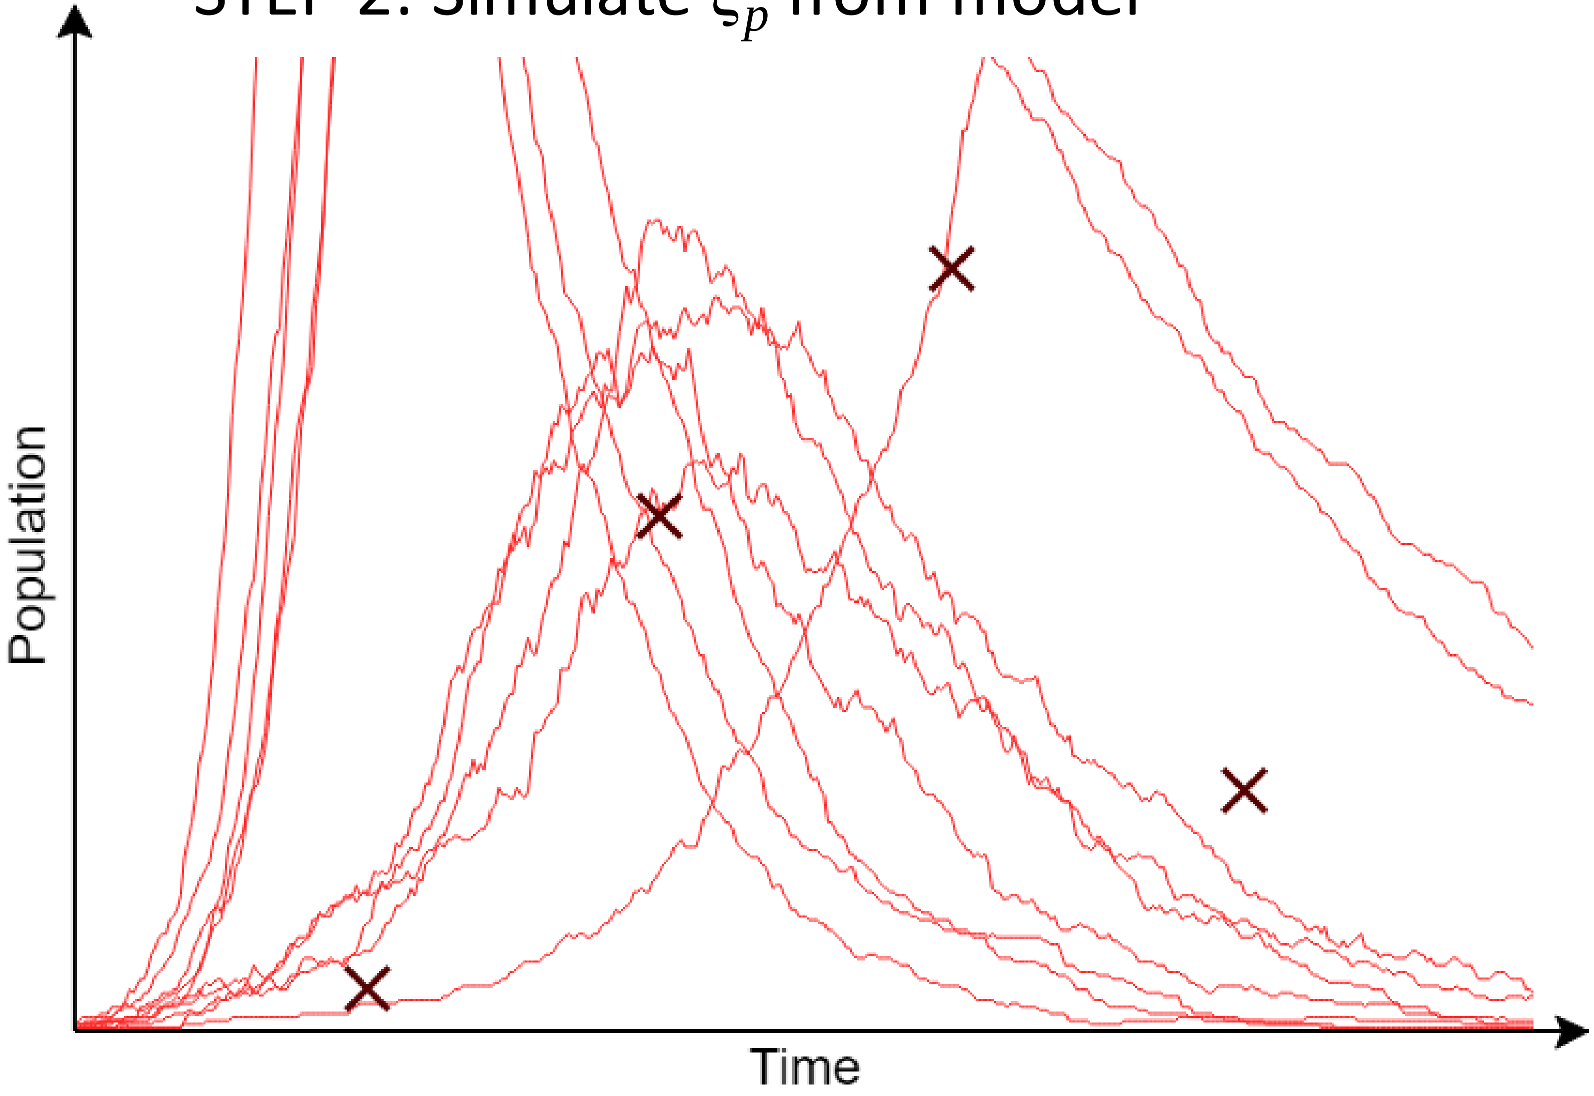

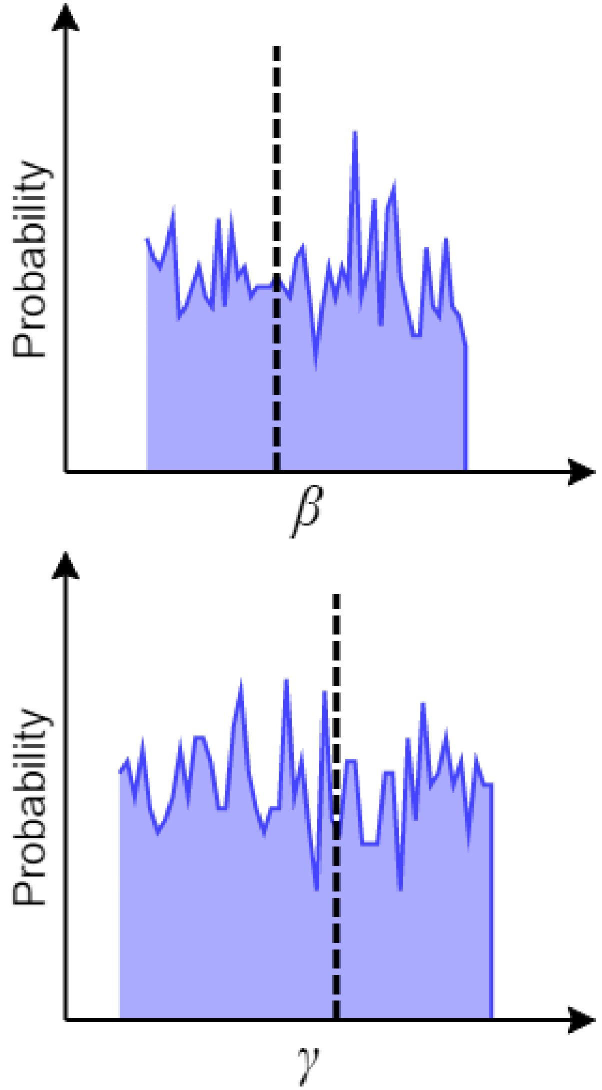

ABC-MBP

STEP 3: Calculate error function EF_p

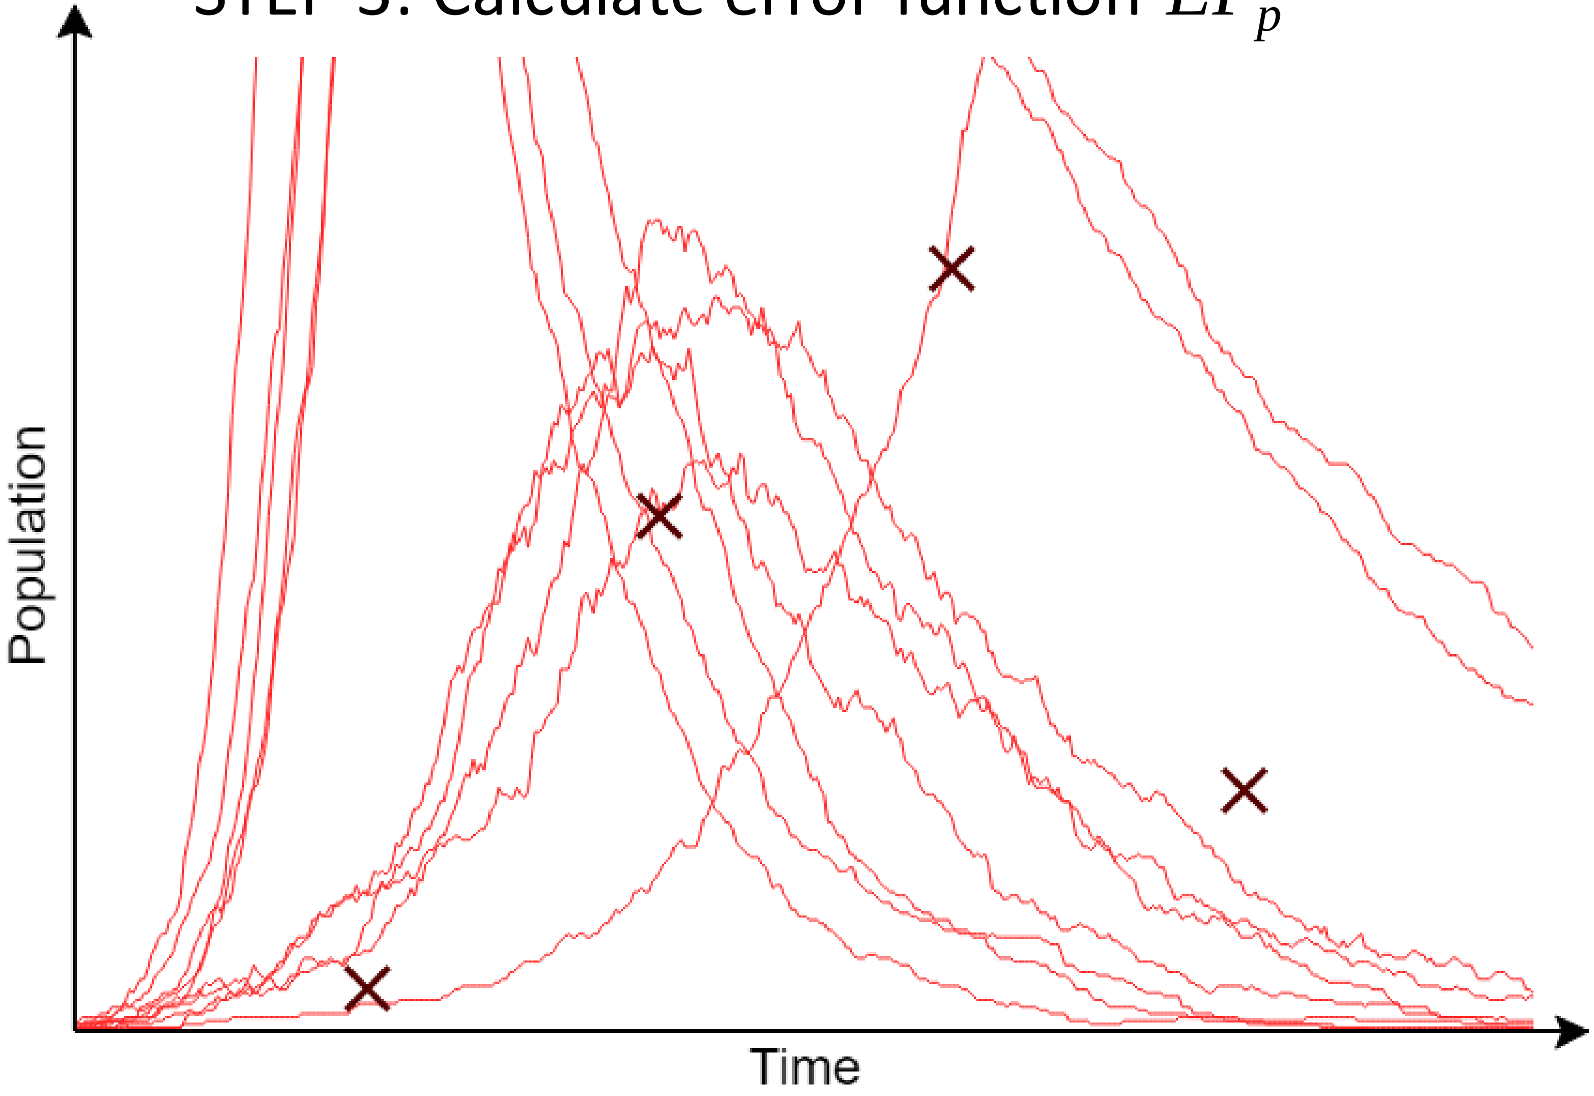

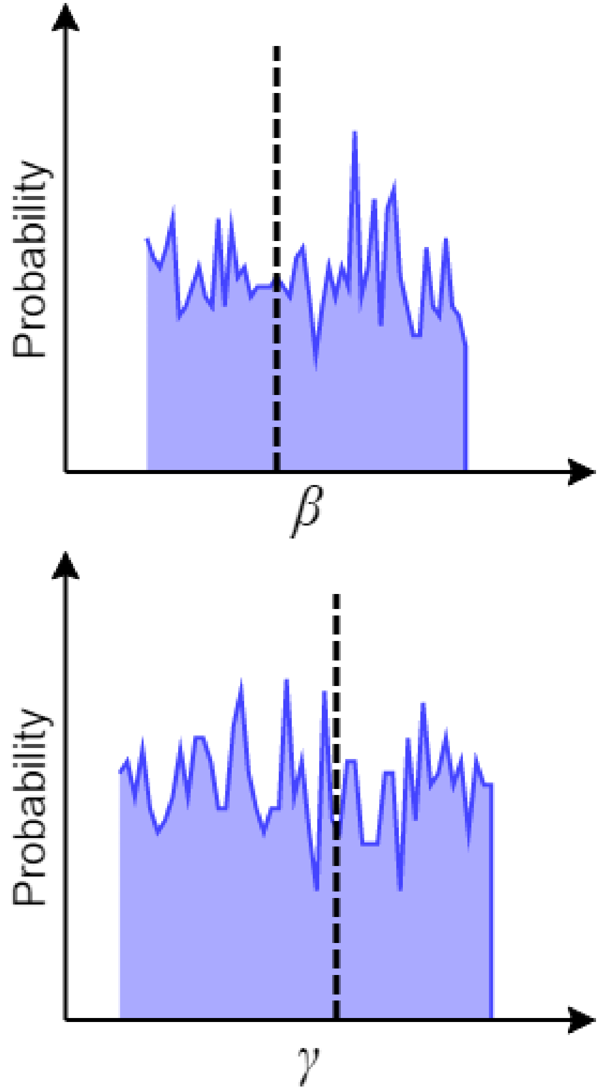

ABC-MBP

STEP 4: Set EF^1_{cutoff} such that half particles have $EF_p < EF^1_{\text{cutoff}}$

$$EF^1_{\text{cutoff}} = 1.5 \times 10^5$$

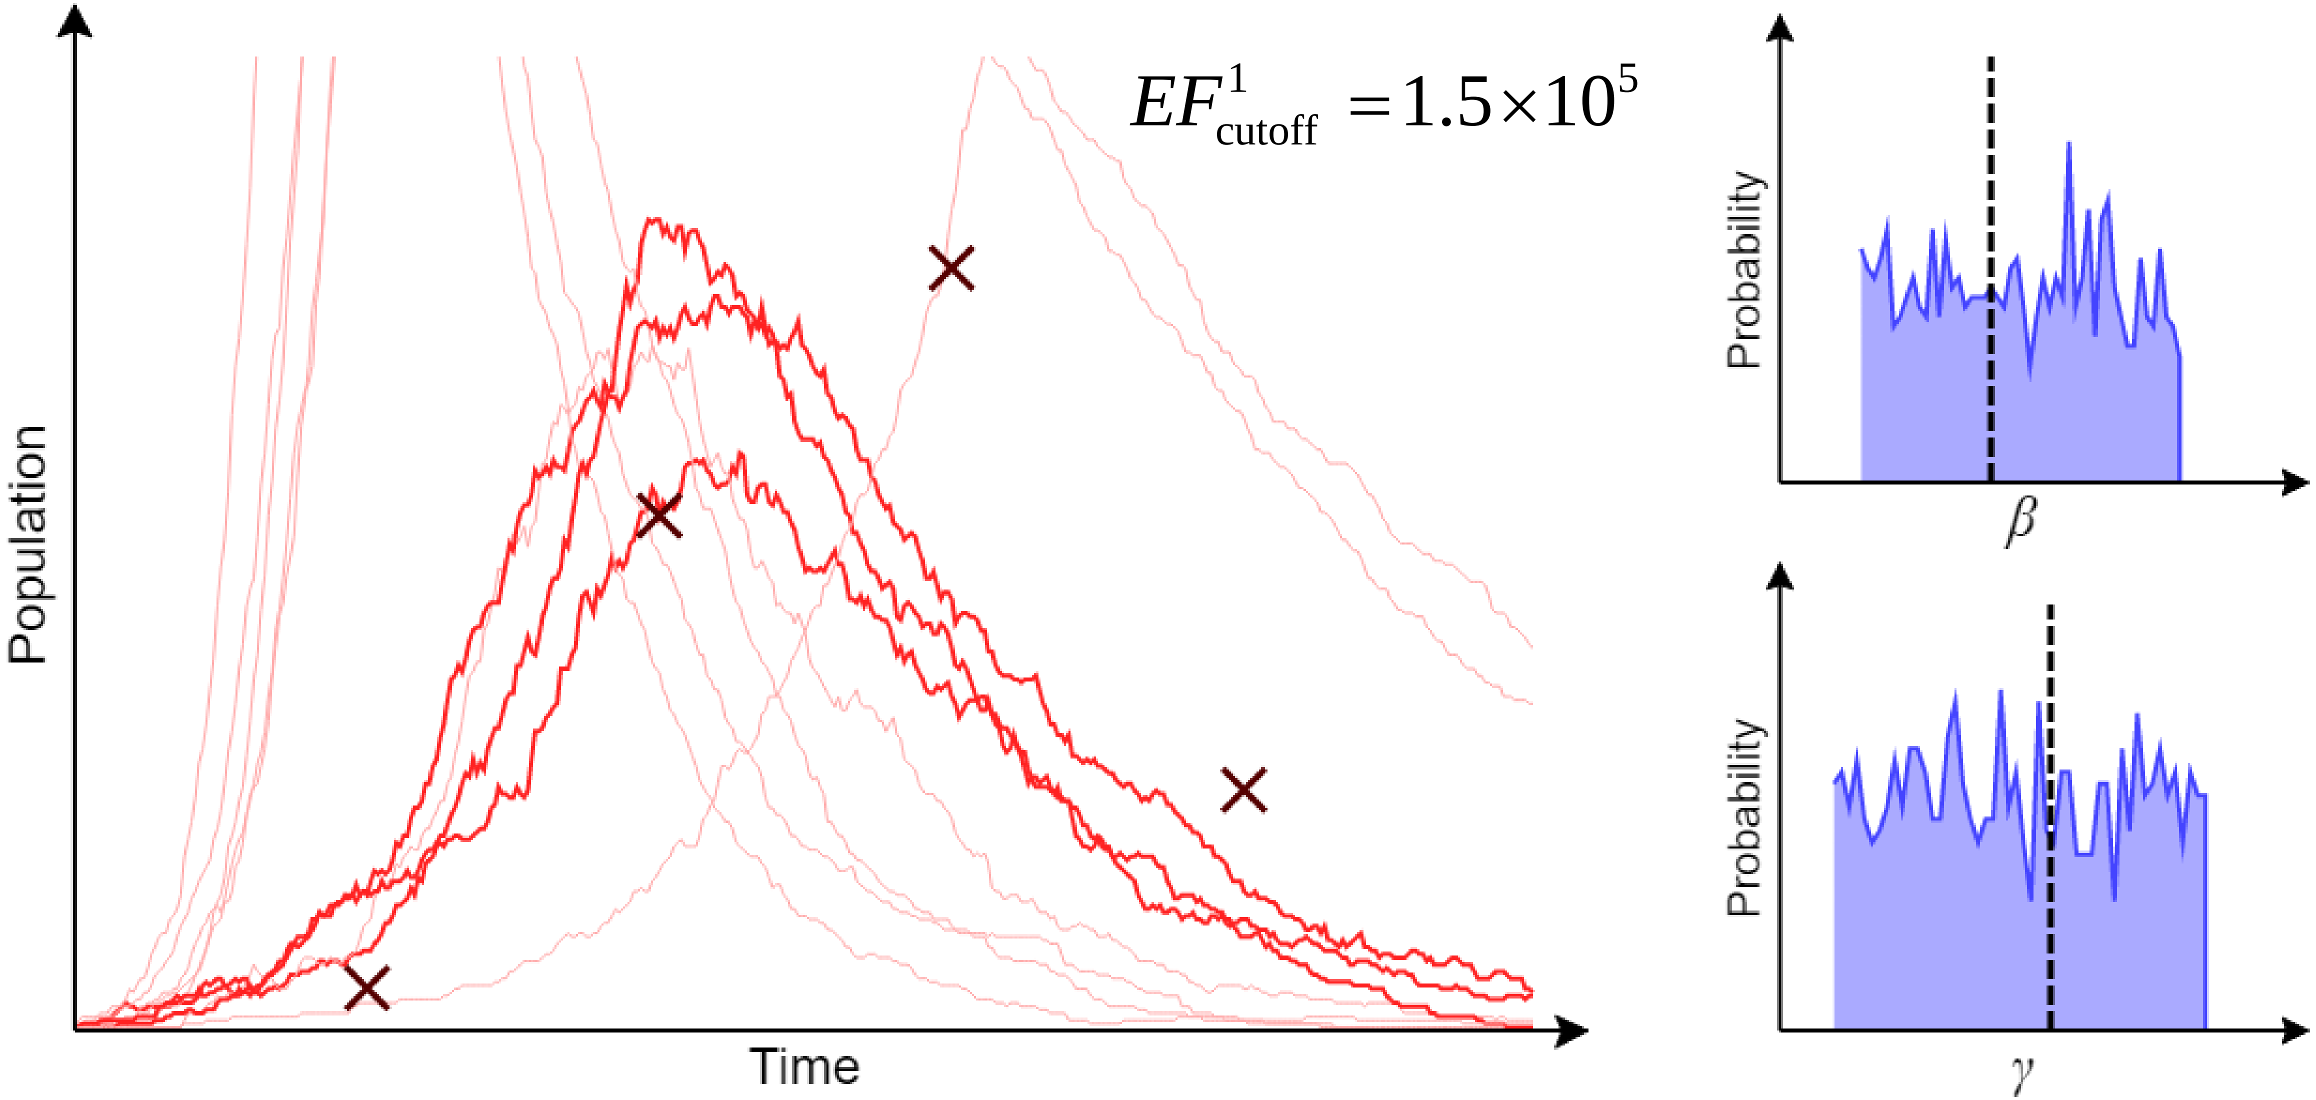

ABC-MBP

STEP 5: Particles with $EF_p < EF^1_{\text{cutoff}}$ are duplicated and others culled

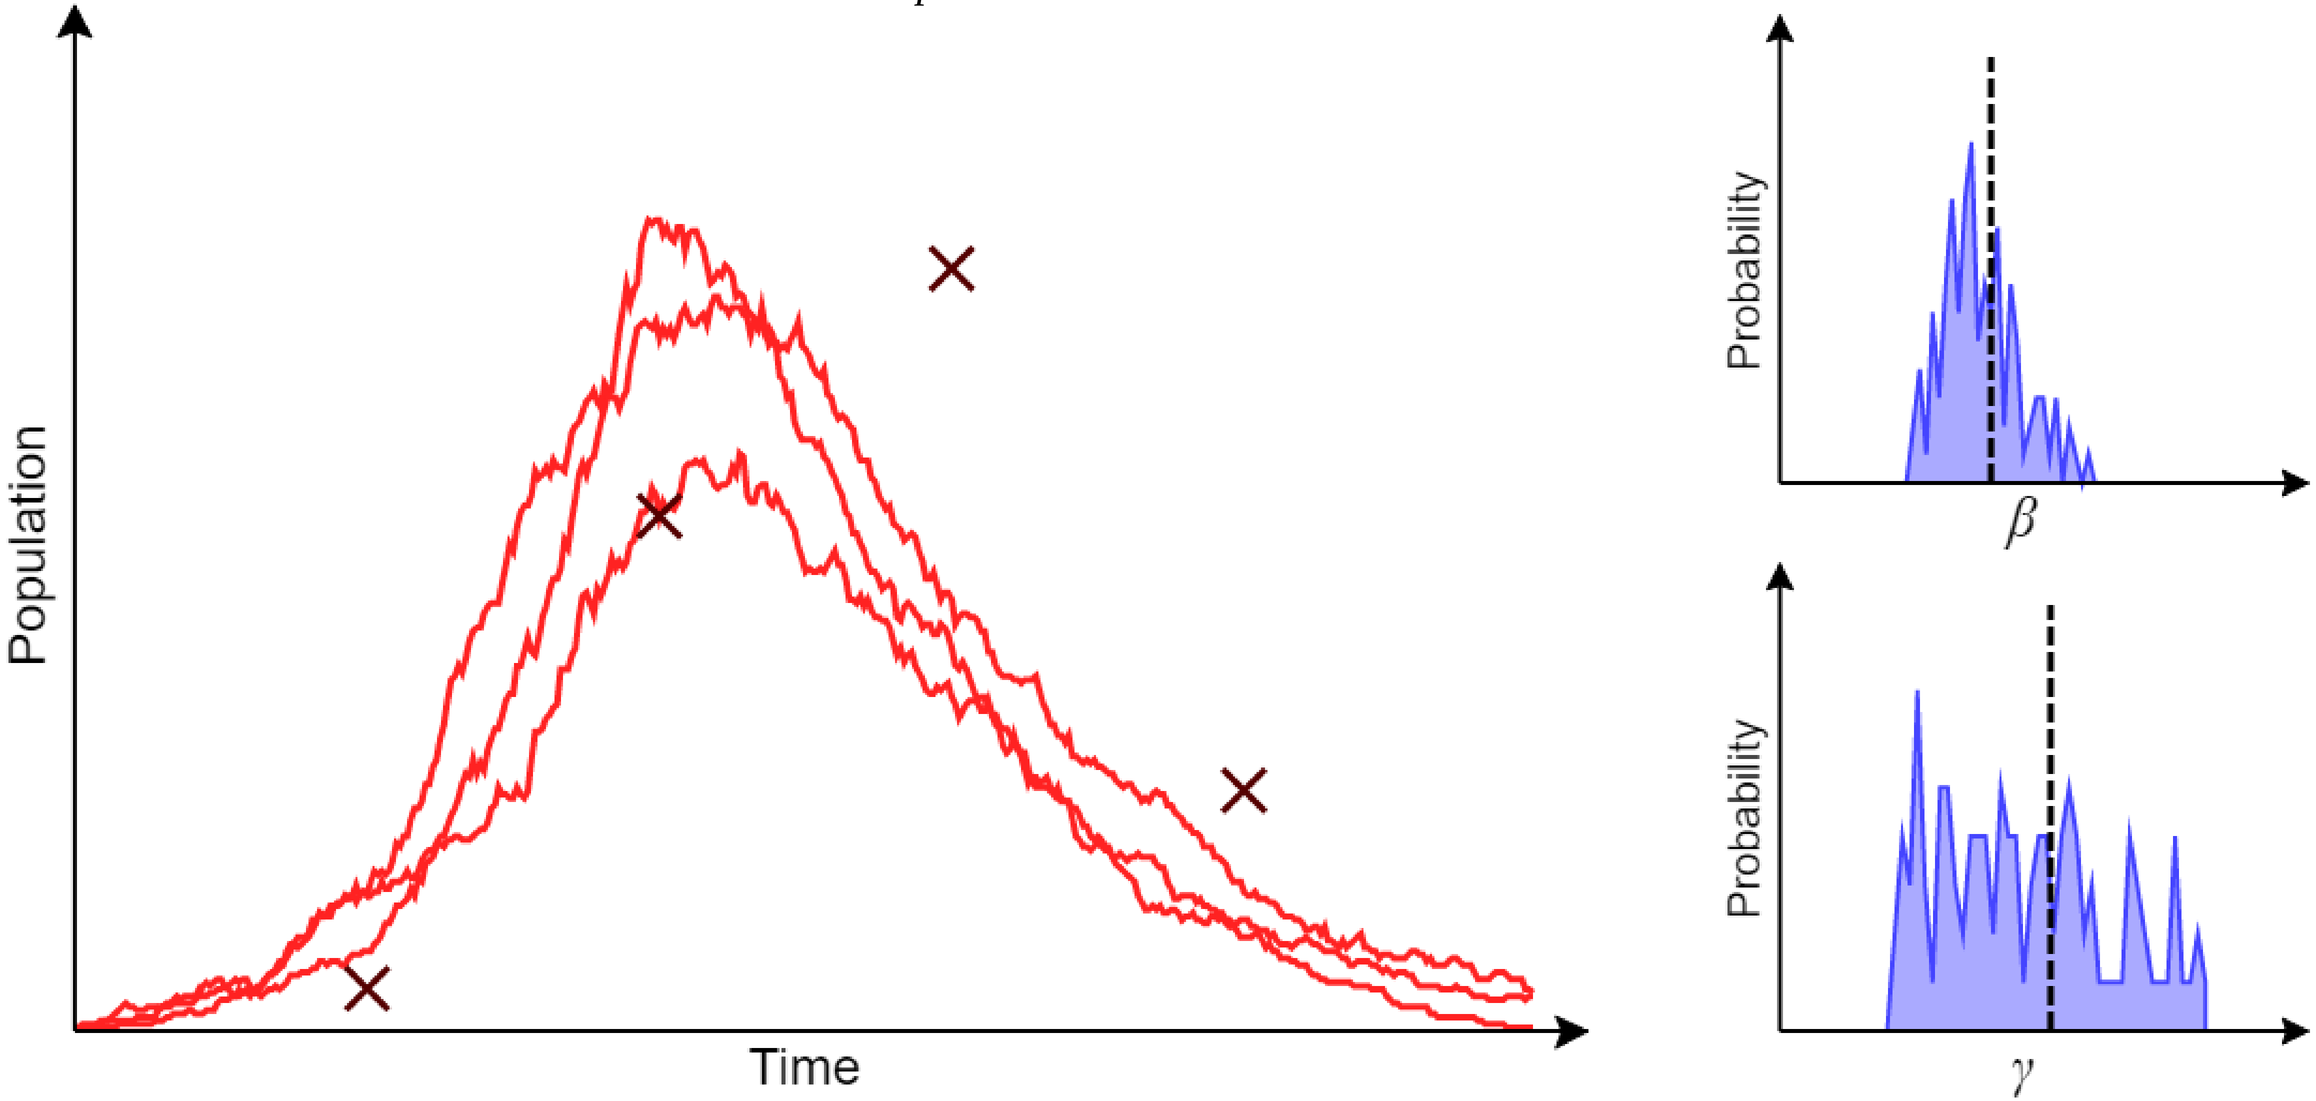

ABC-MBP

STEP 6: MBPs are applied to allow particles to explore posterior space

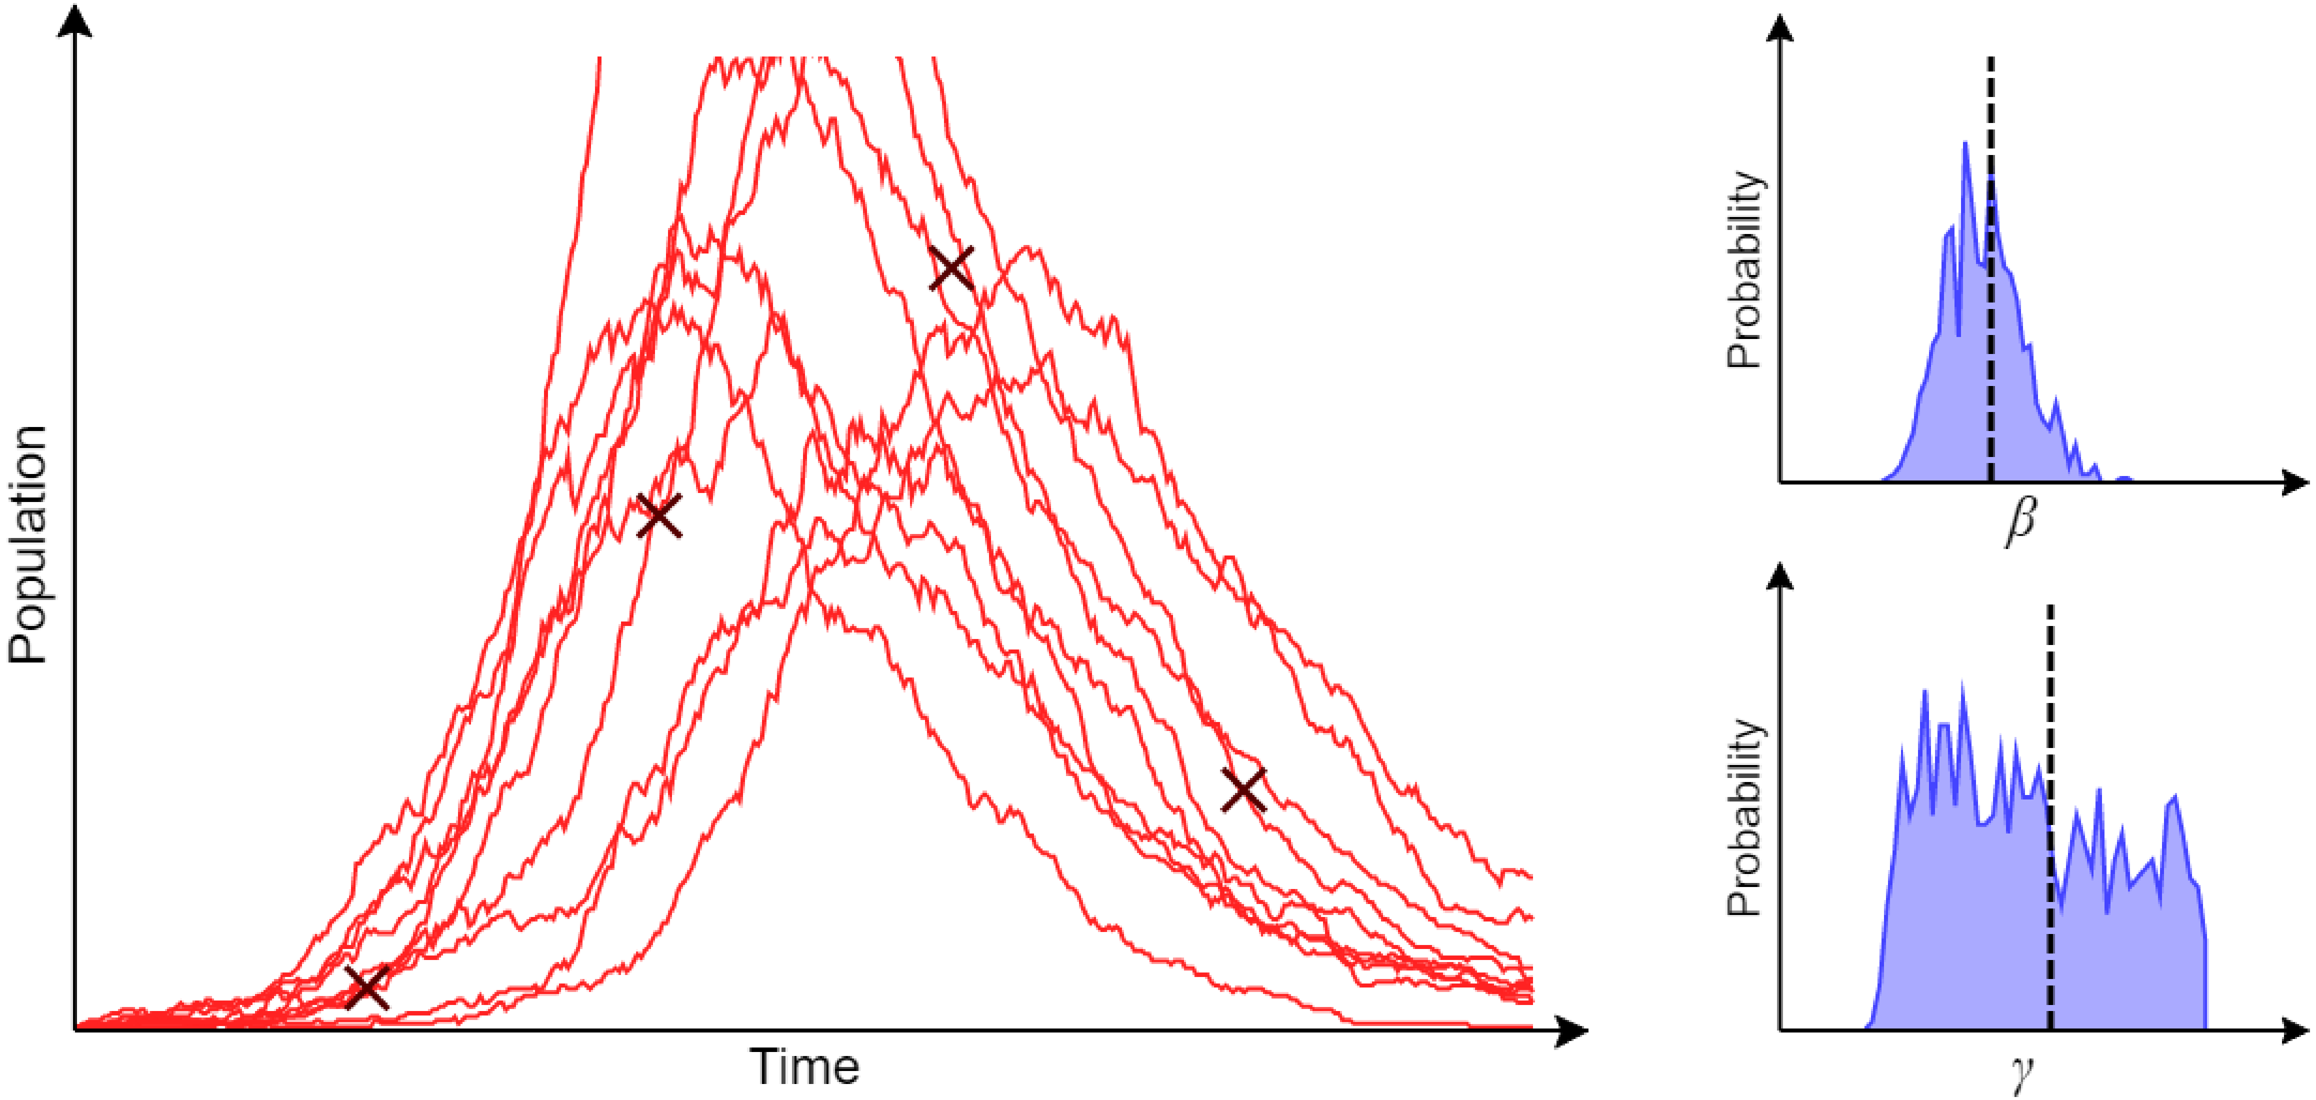

ABC-MBP

This completes the first generation. Repeat over more generations…

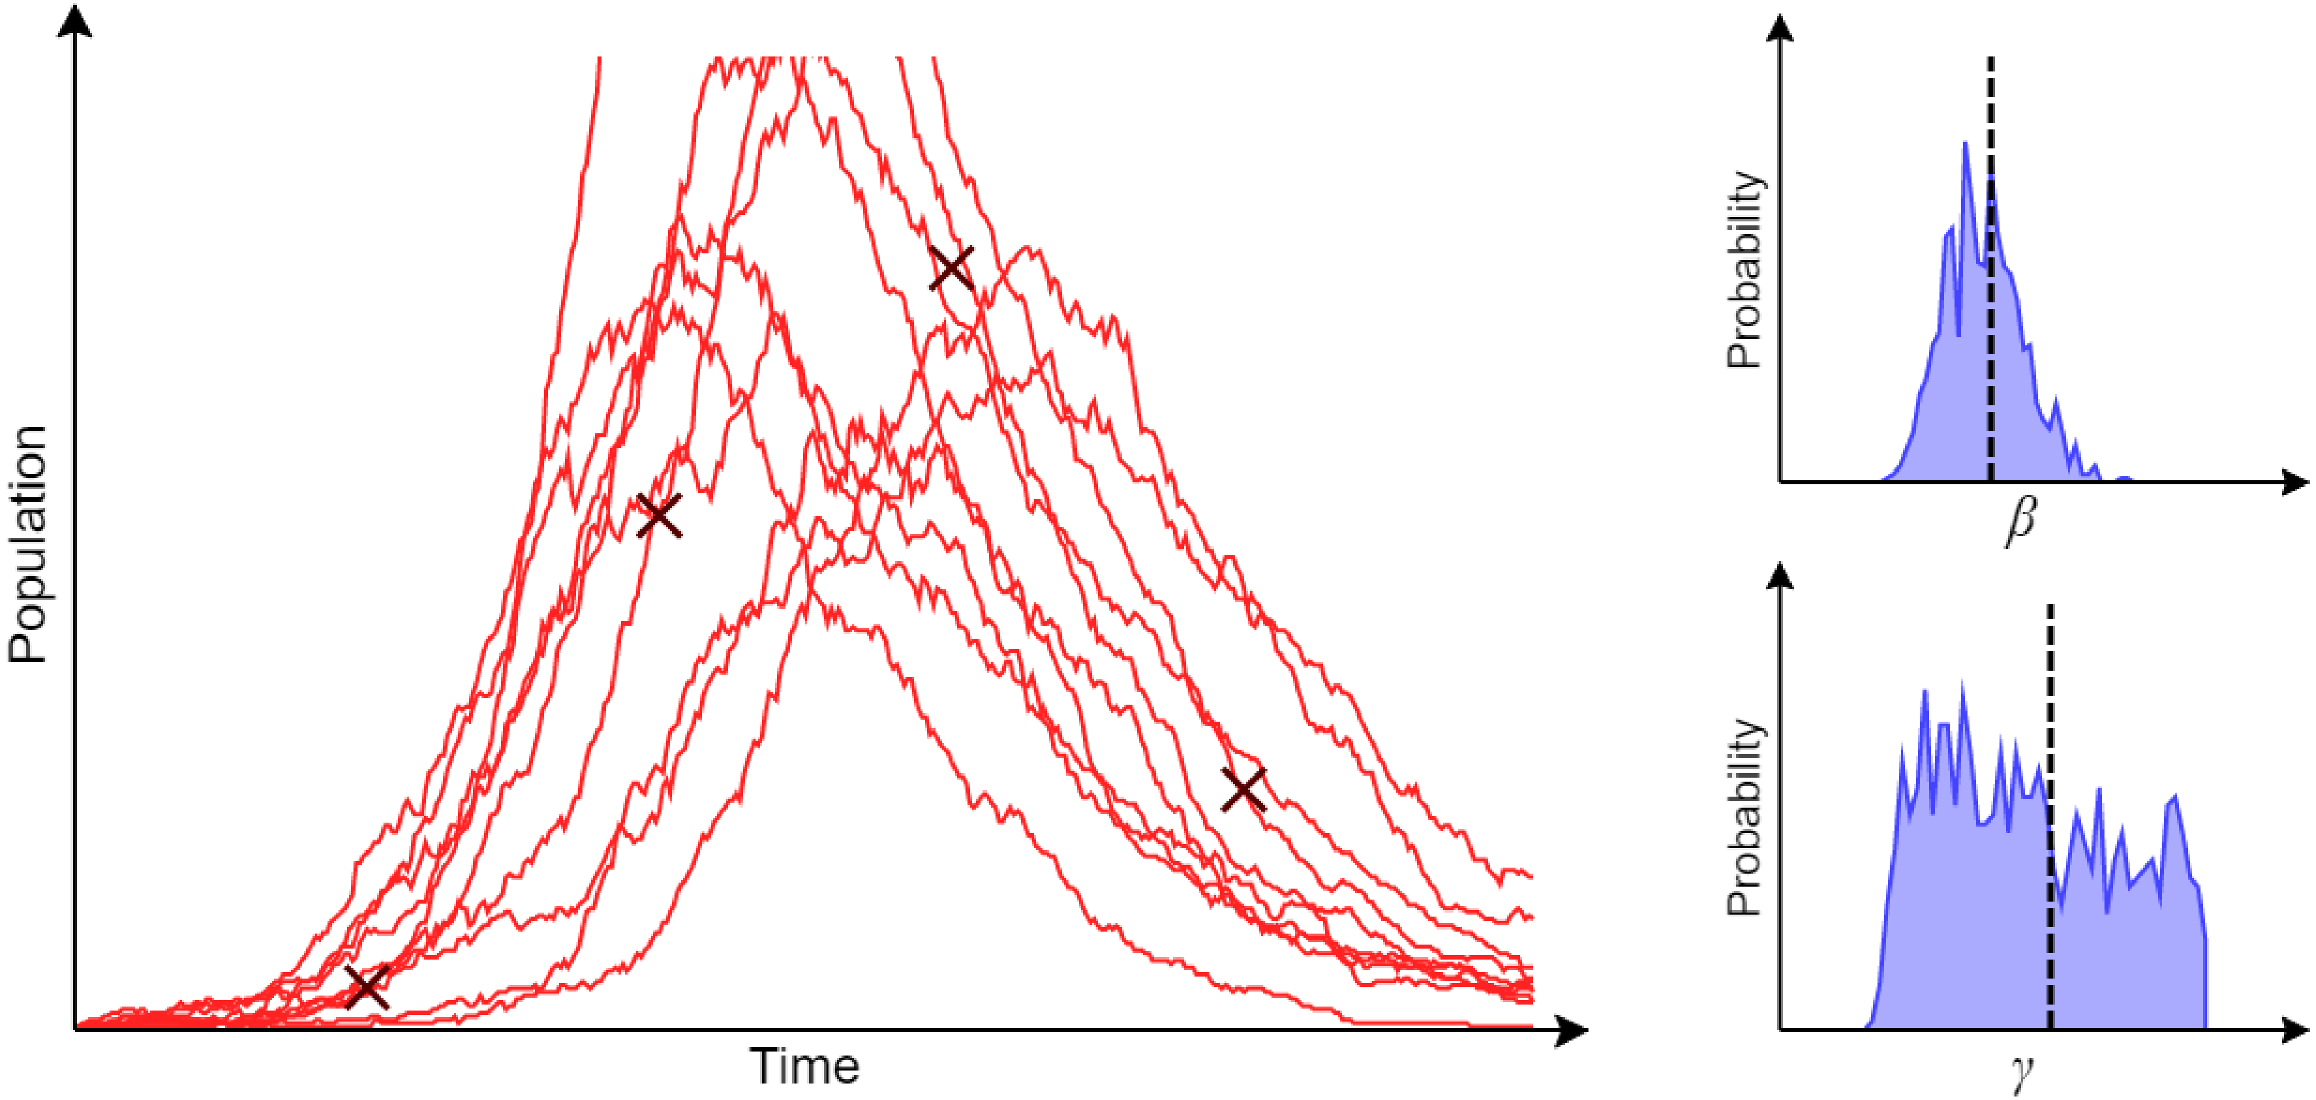

ABC-MBP

Generation 2: Calculate error function EF_p

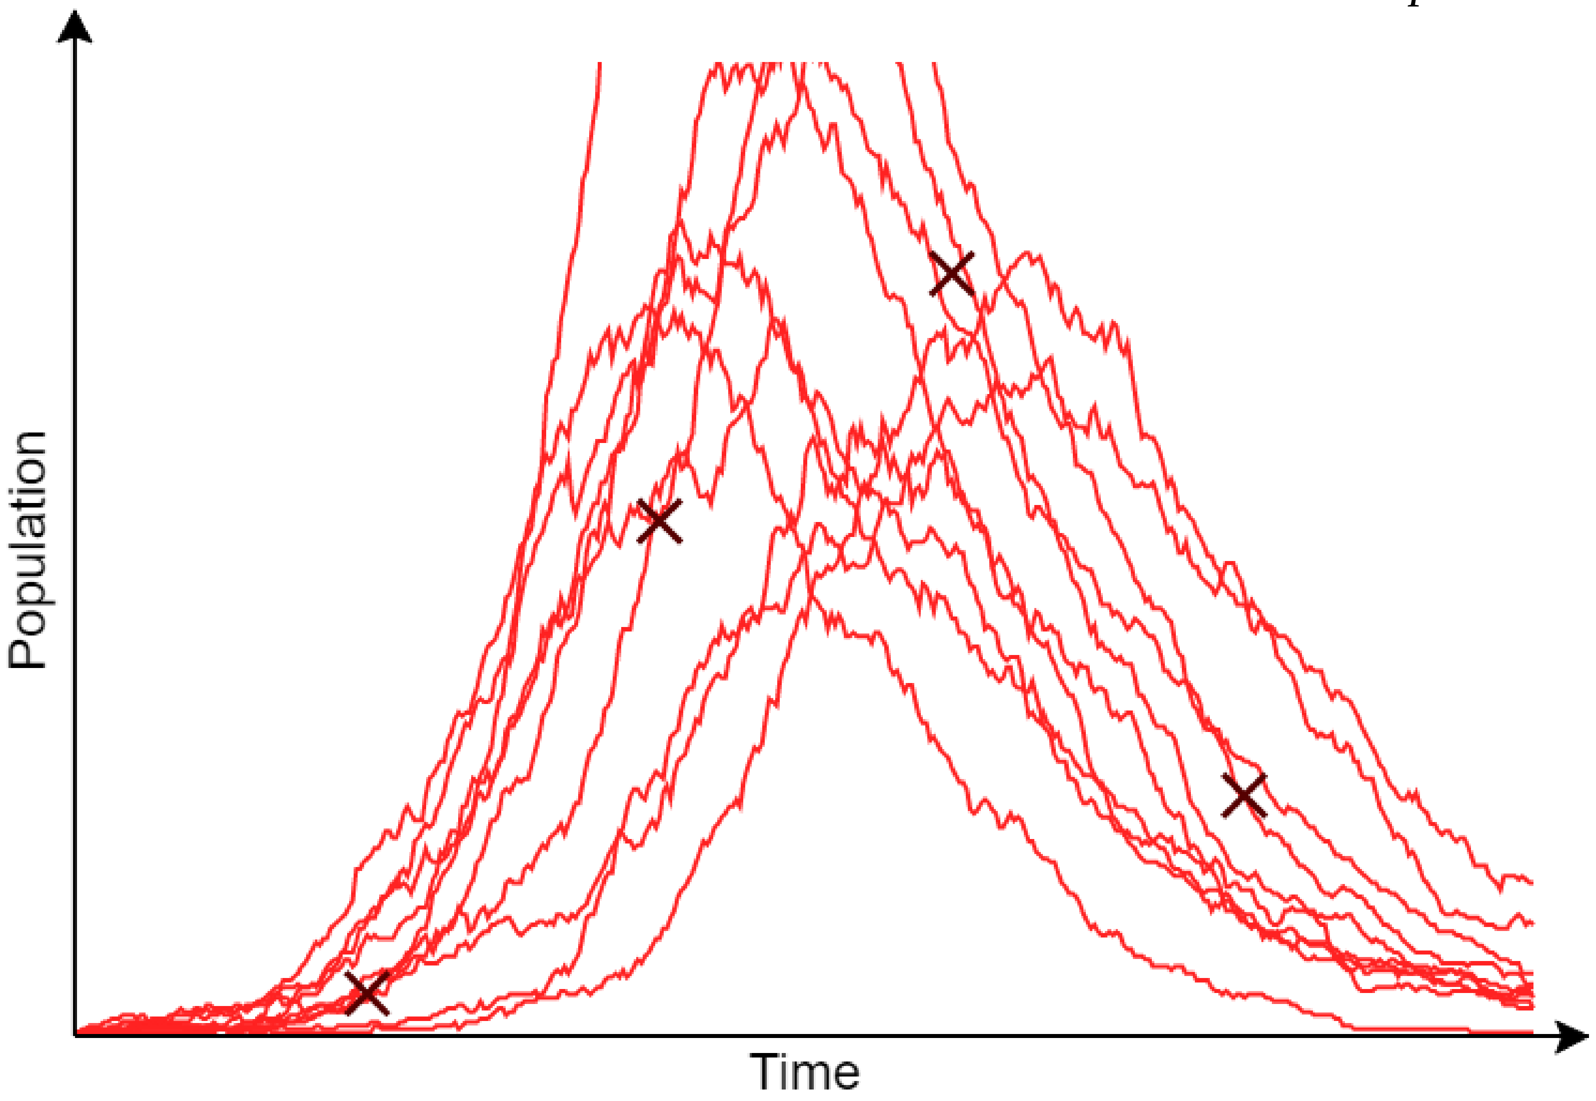

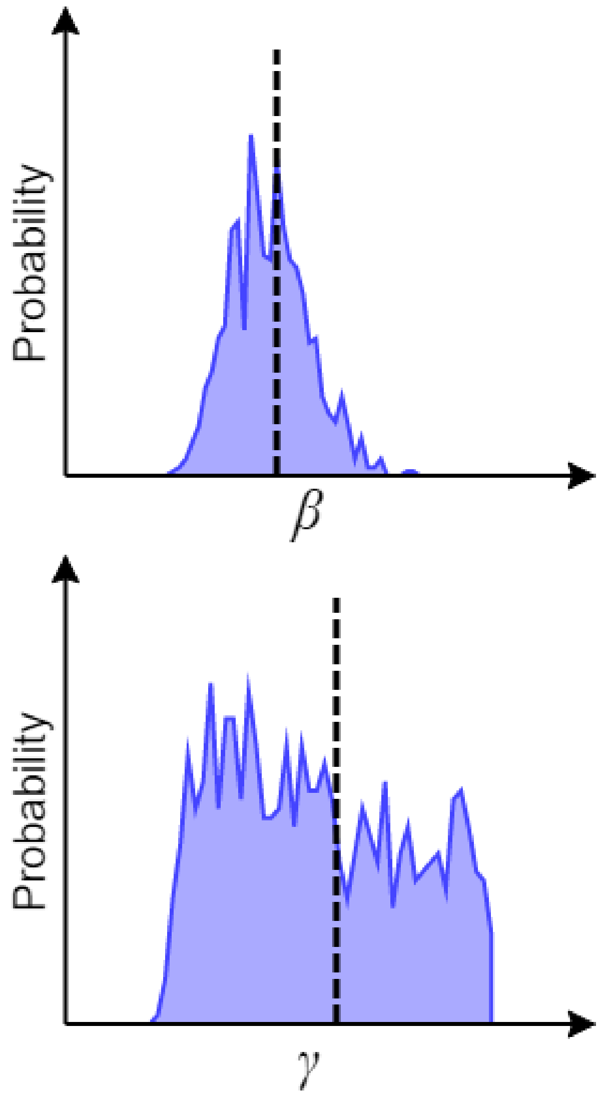

ABC-MBP

Generation 2: Set EF^2_{cutoff} such that half particles have $EF_p < EF^2_{\text{cutoff}}$

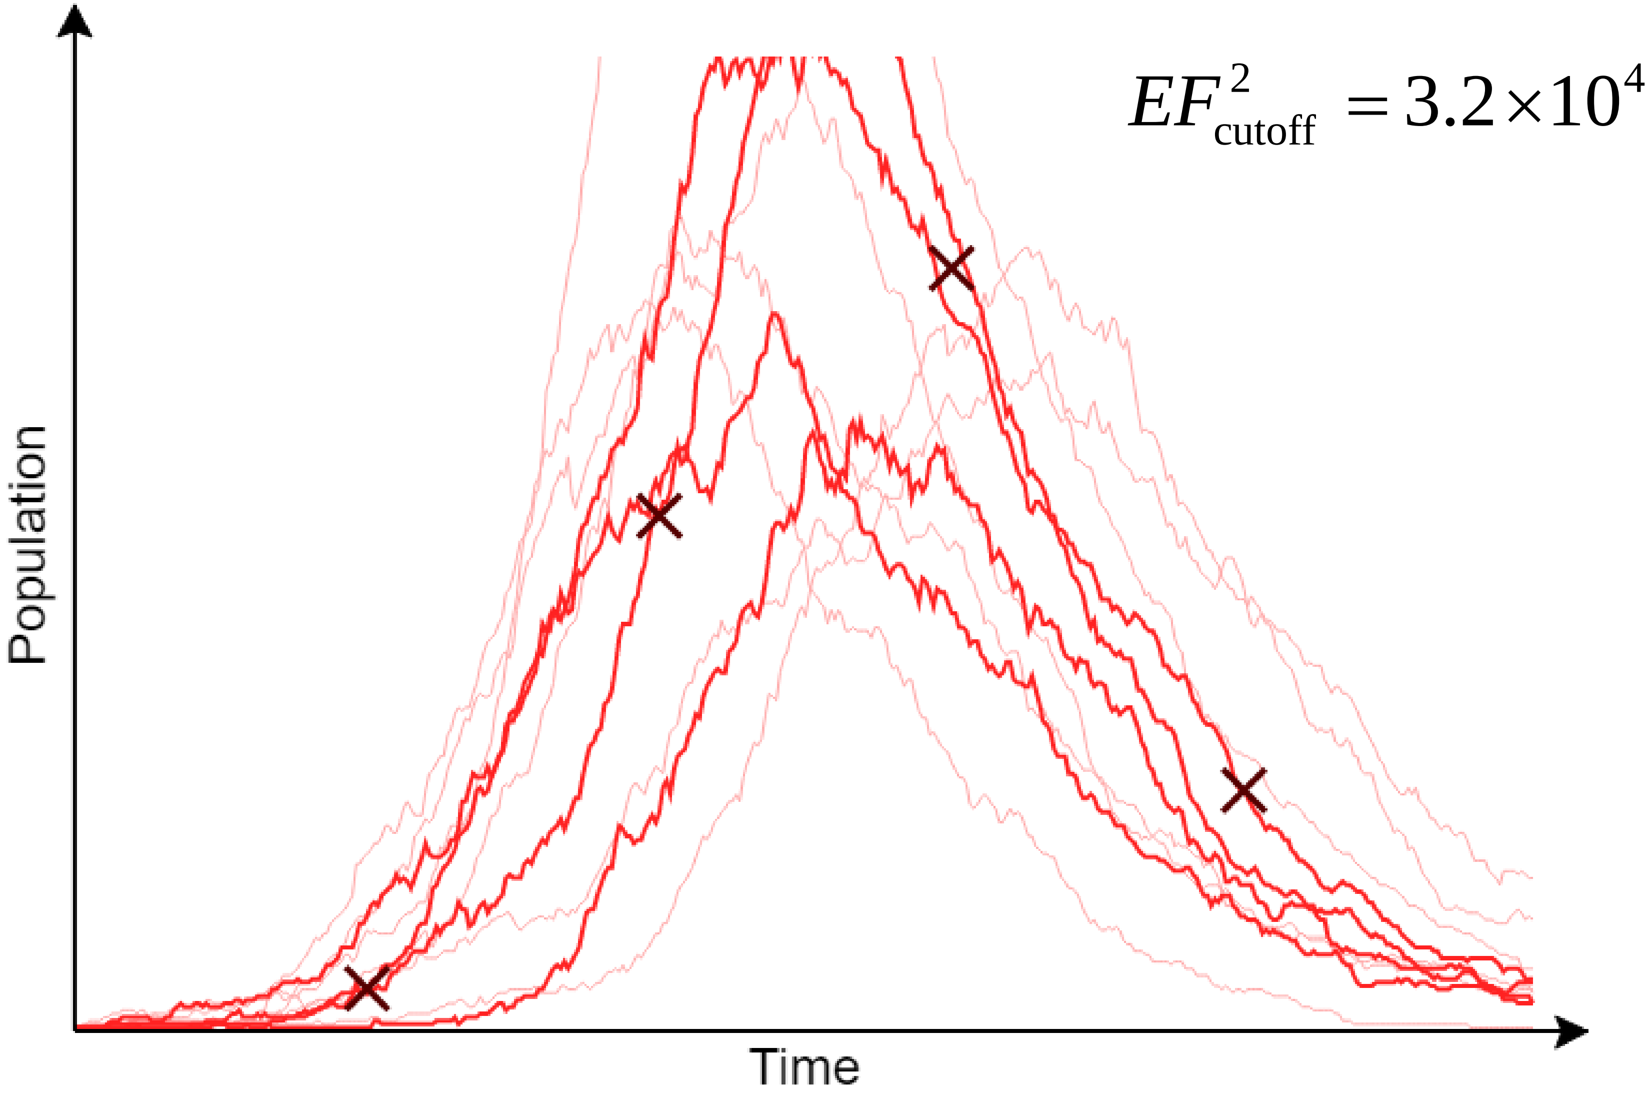

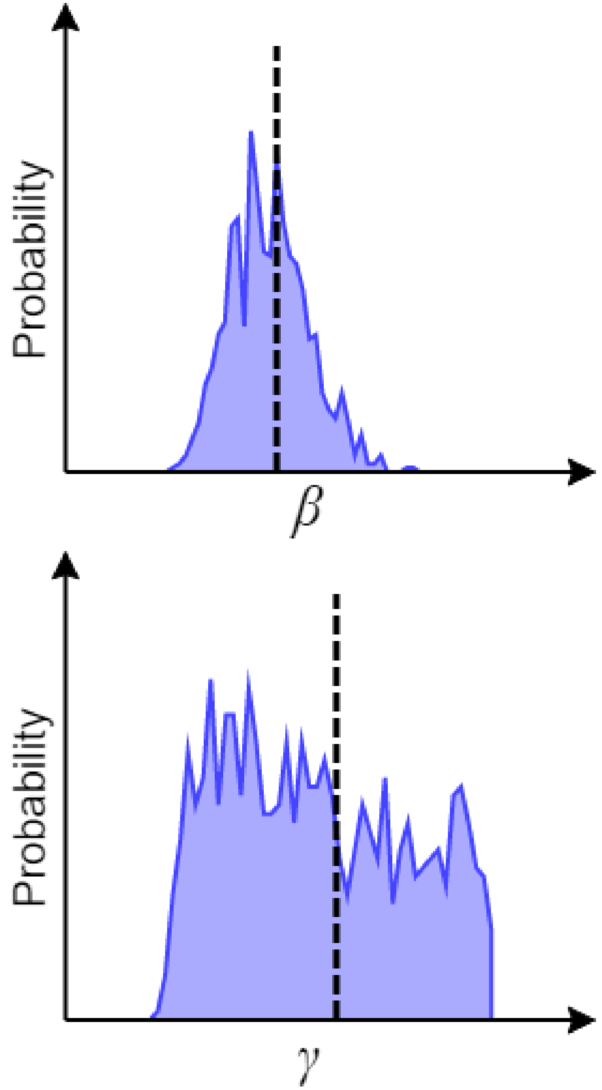

ABC-MBP

Generation 2: Particles with $EF_p < EF^2_{\text{cutoff}}$ are duplicated, others culled.

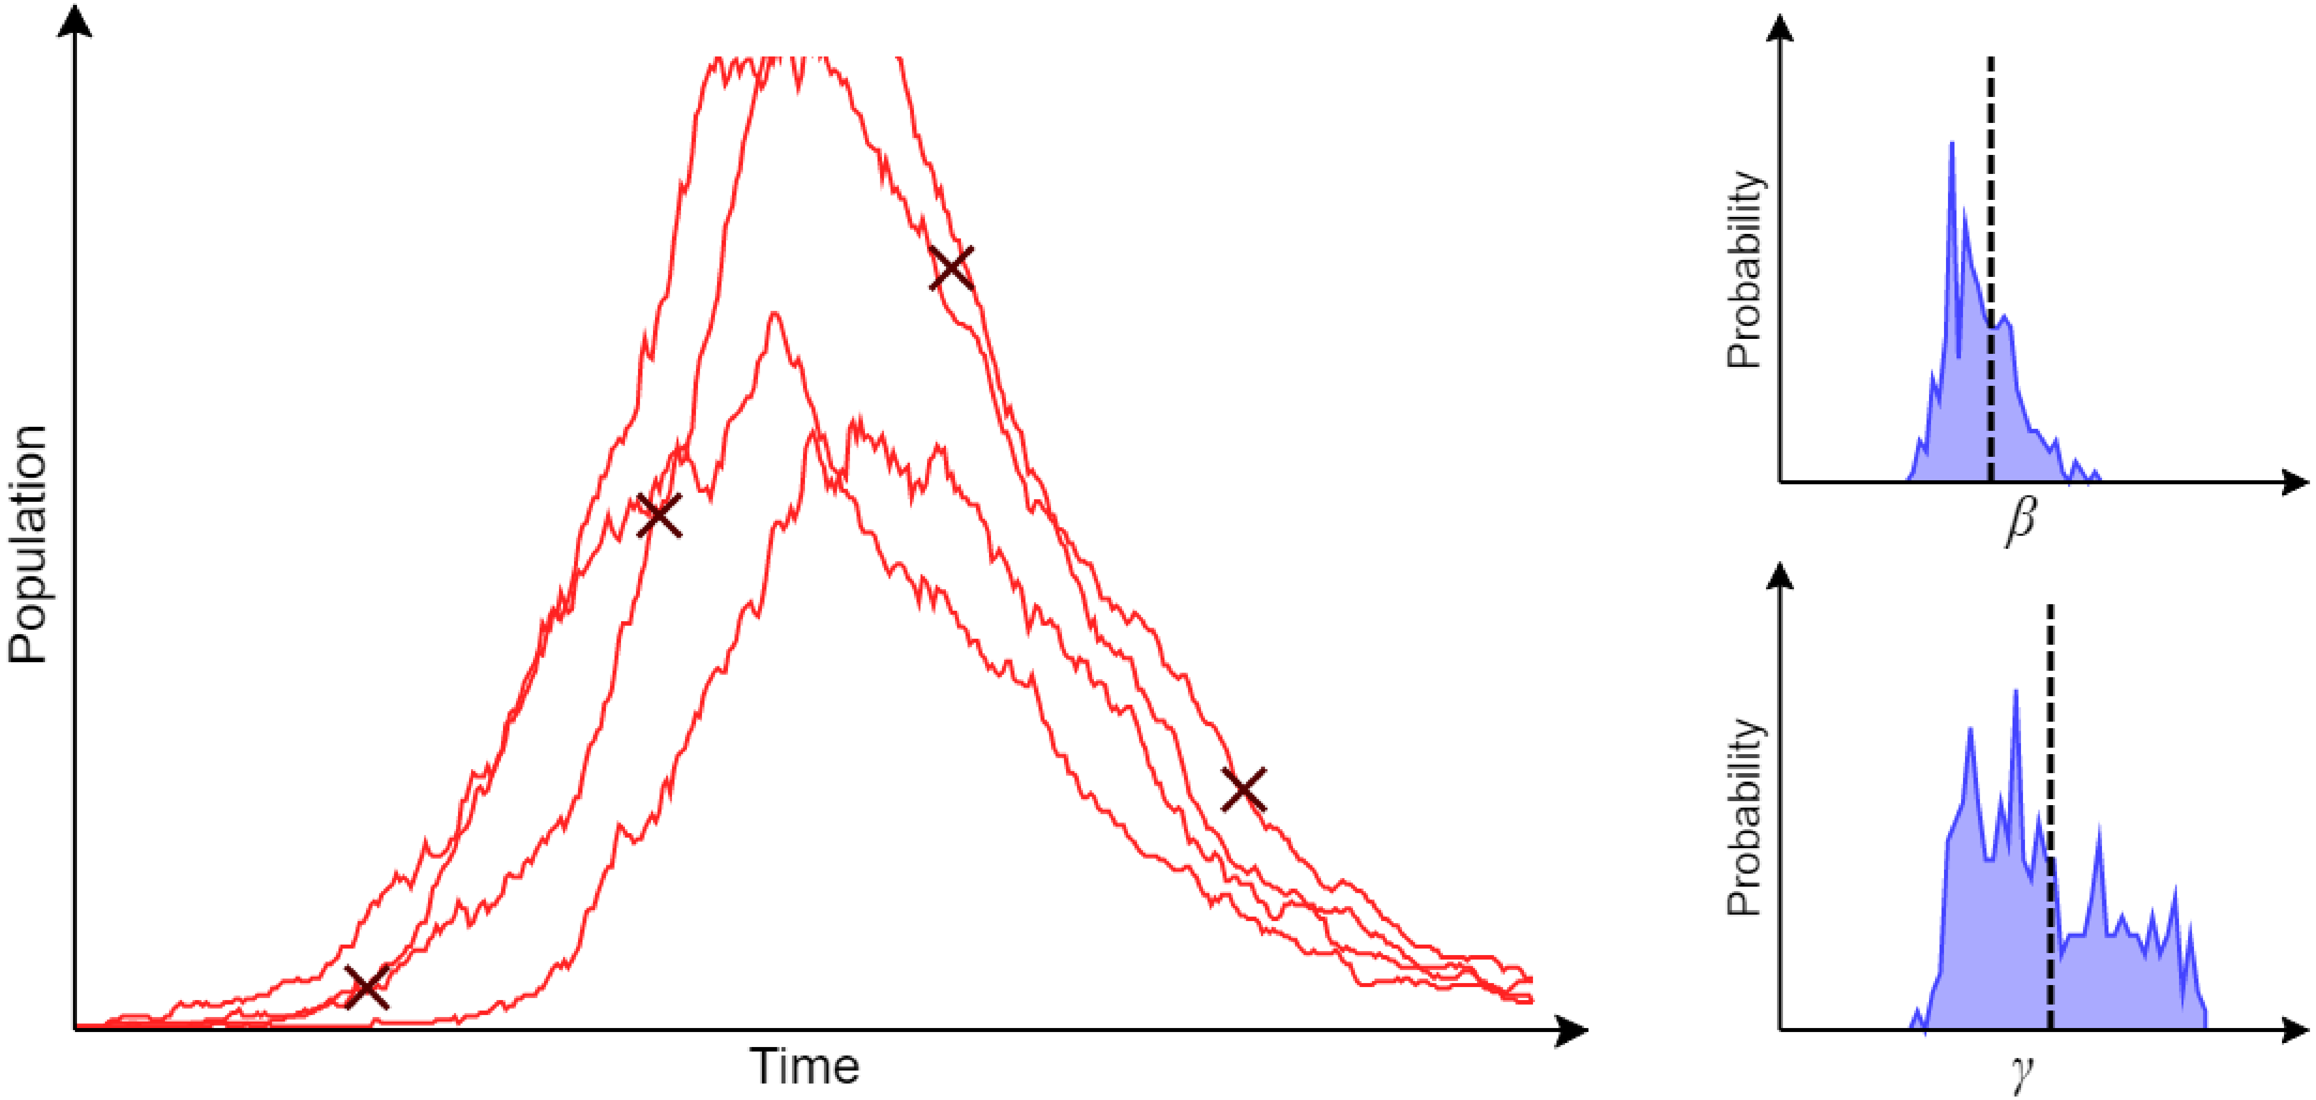

ABC-MBP

Generation 2: MBPs applied to allow particles to explore posterior space

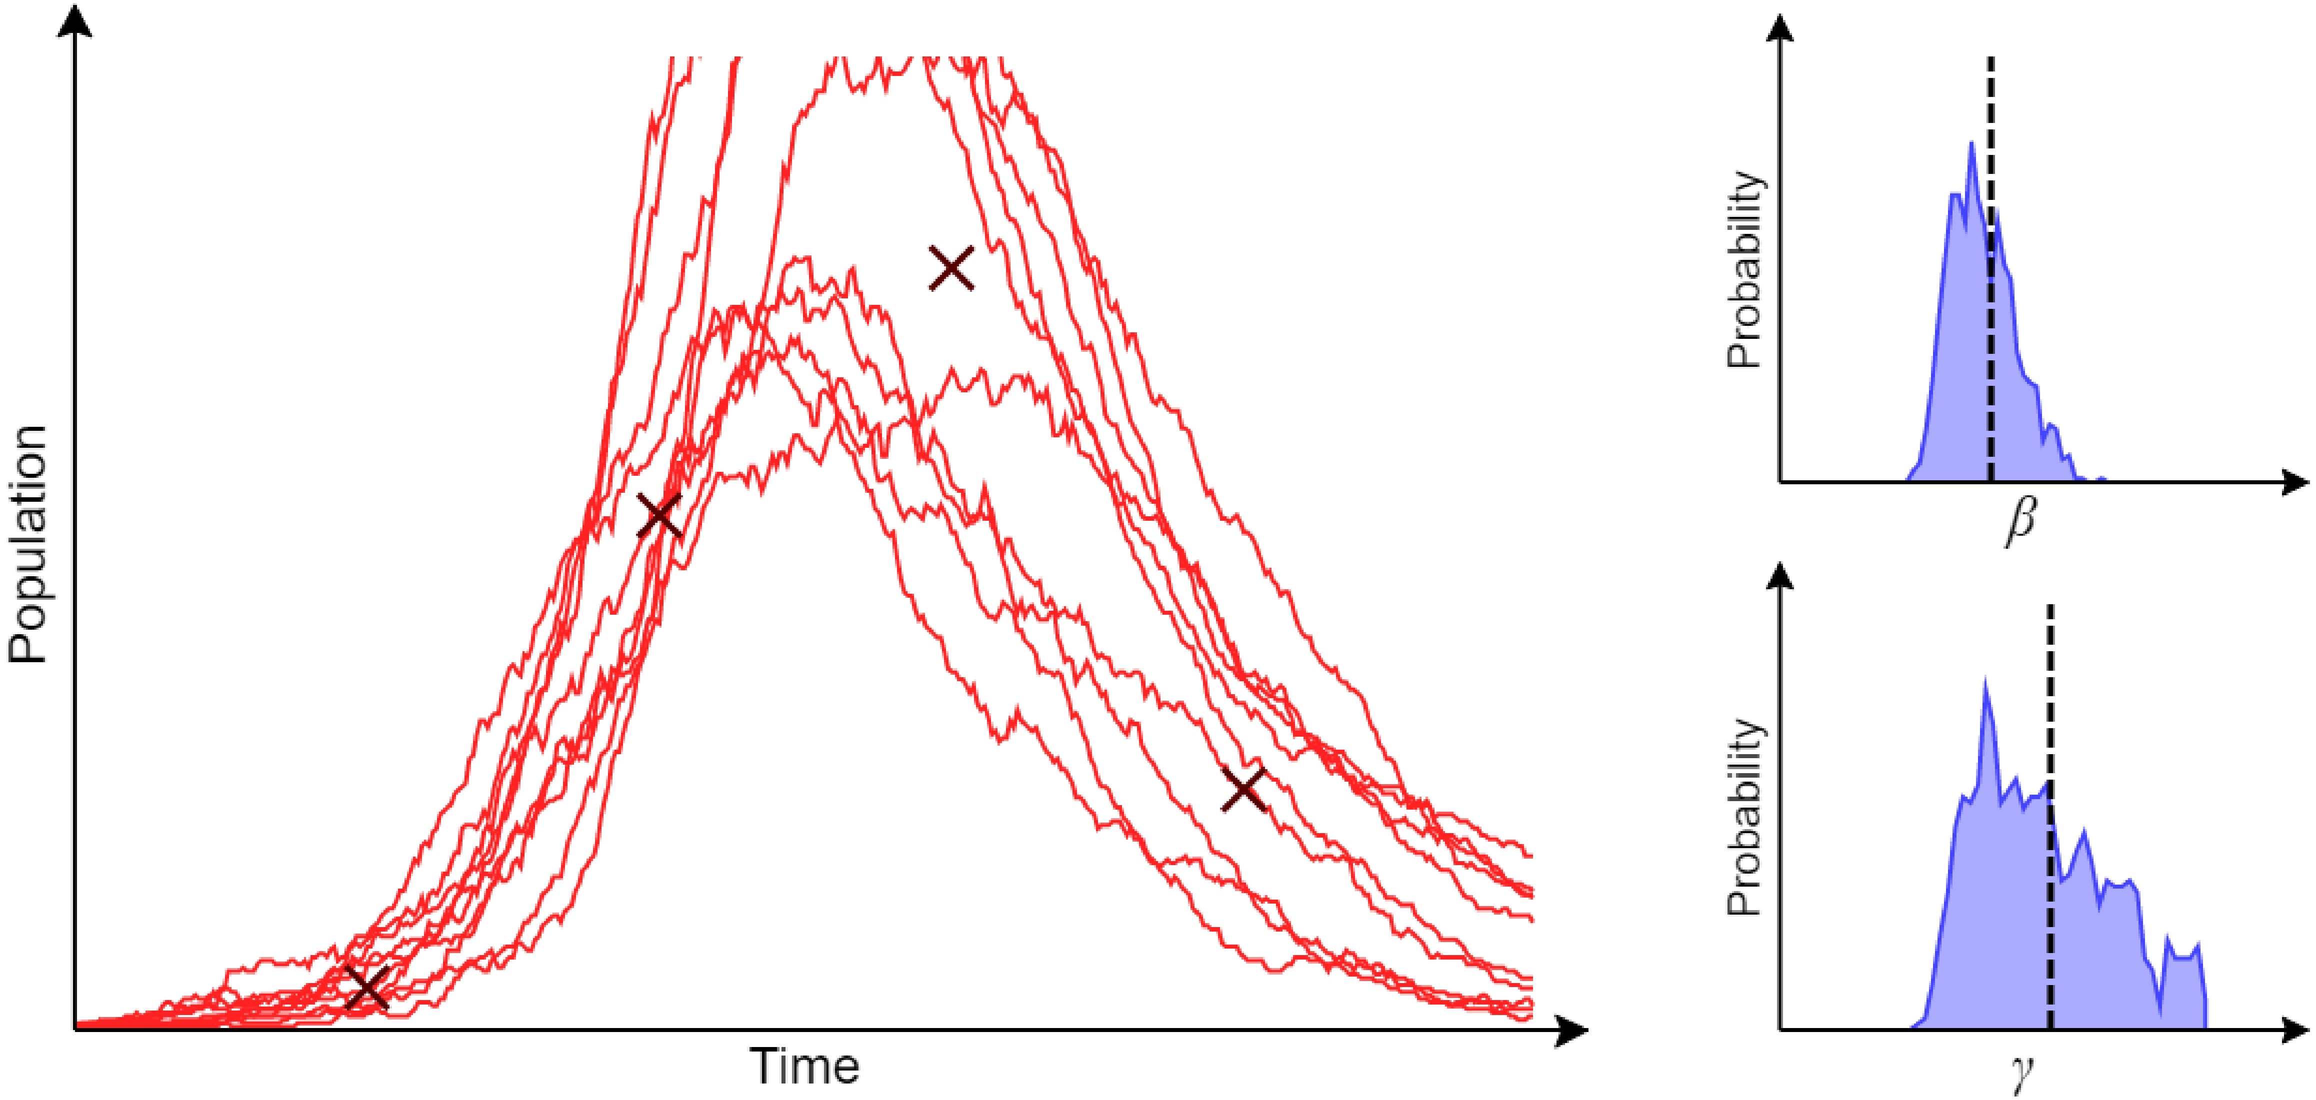

ABC-MBP

Generation 3:

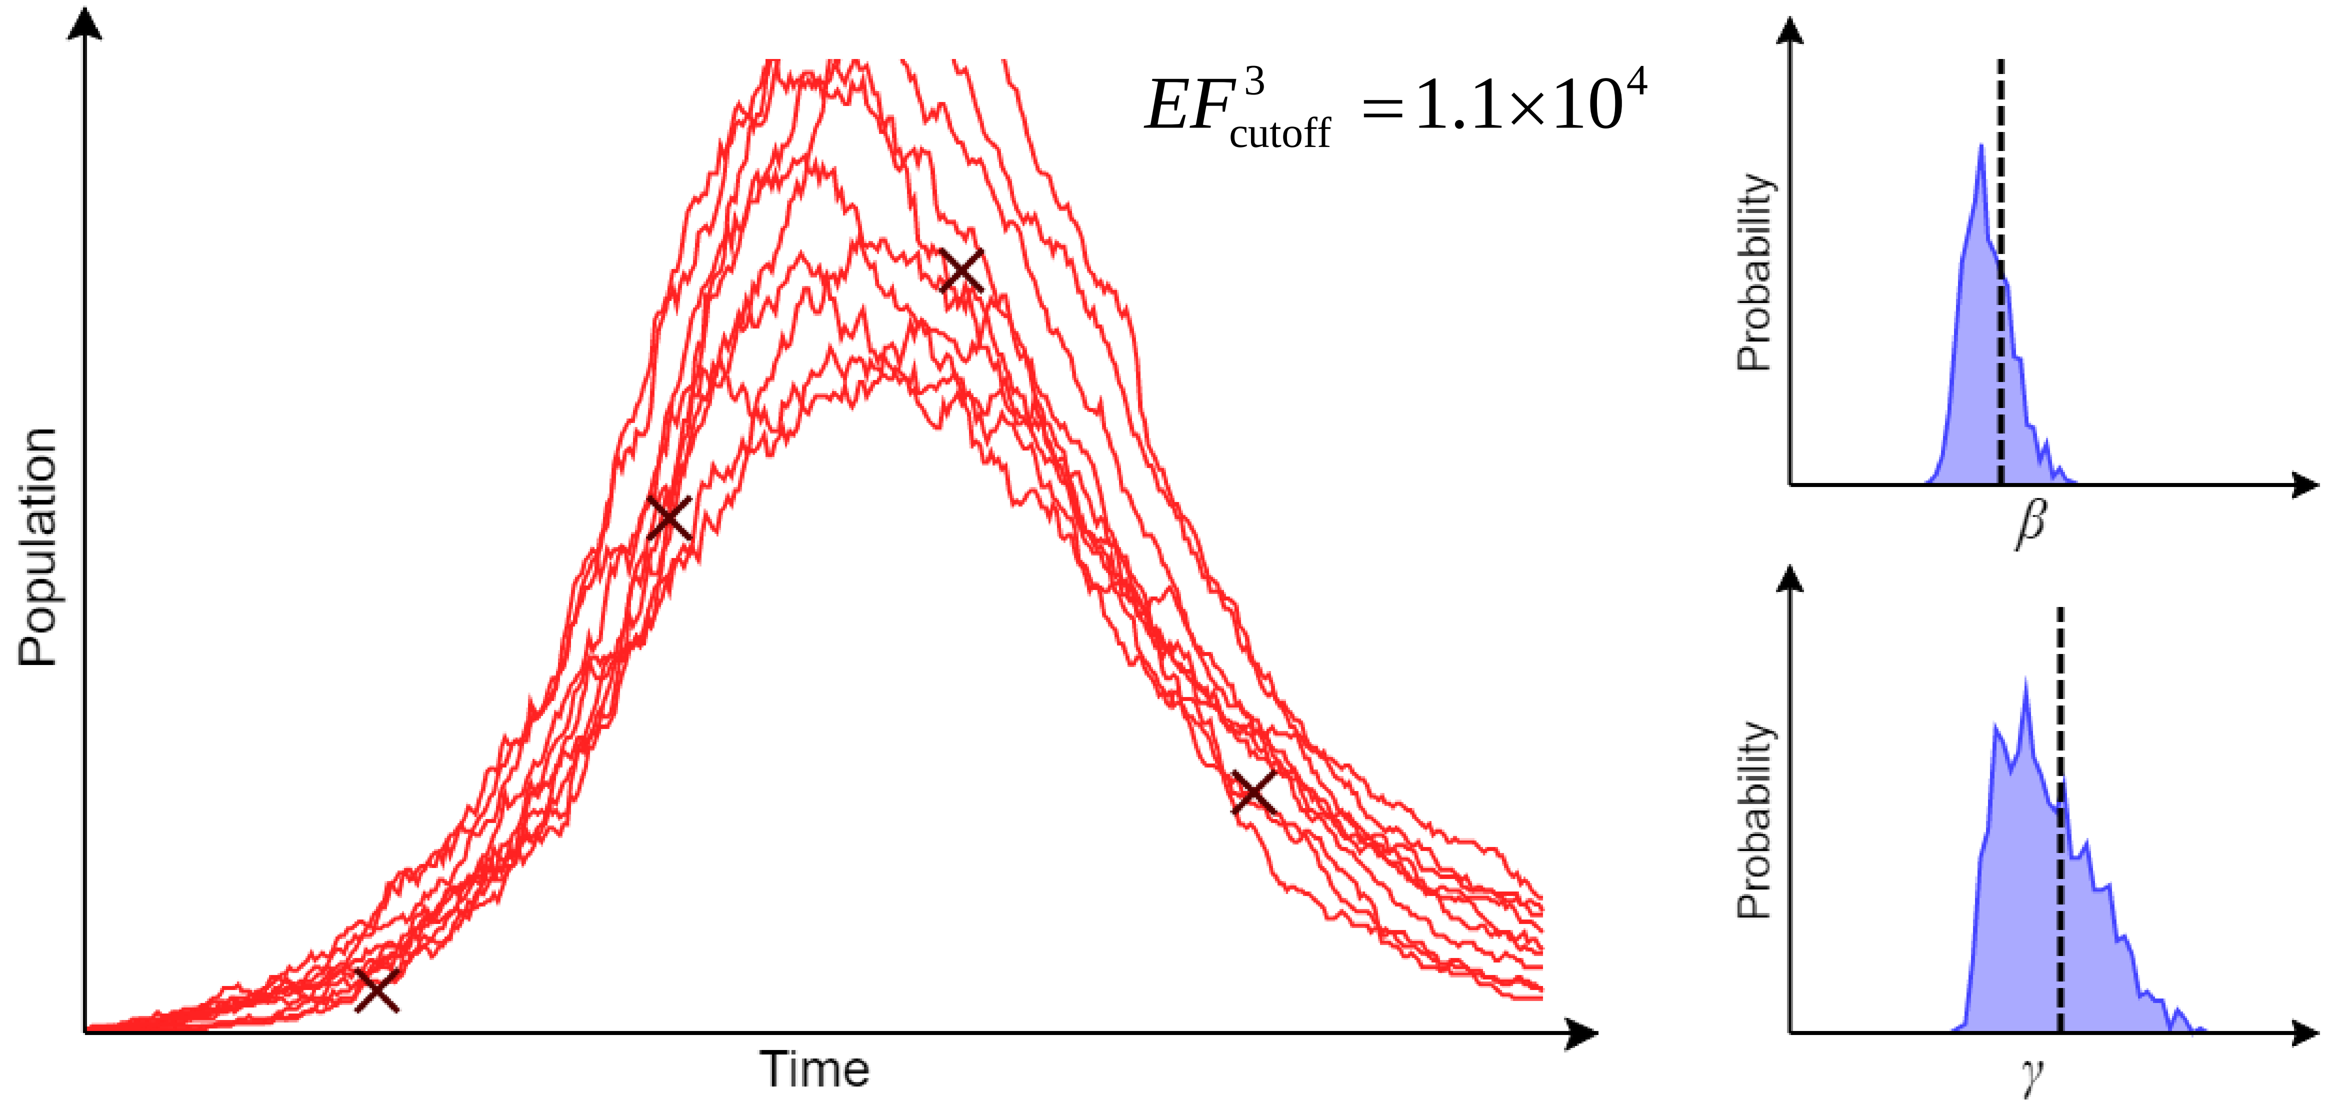

ABC-MBP

Generation 4:

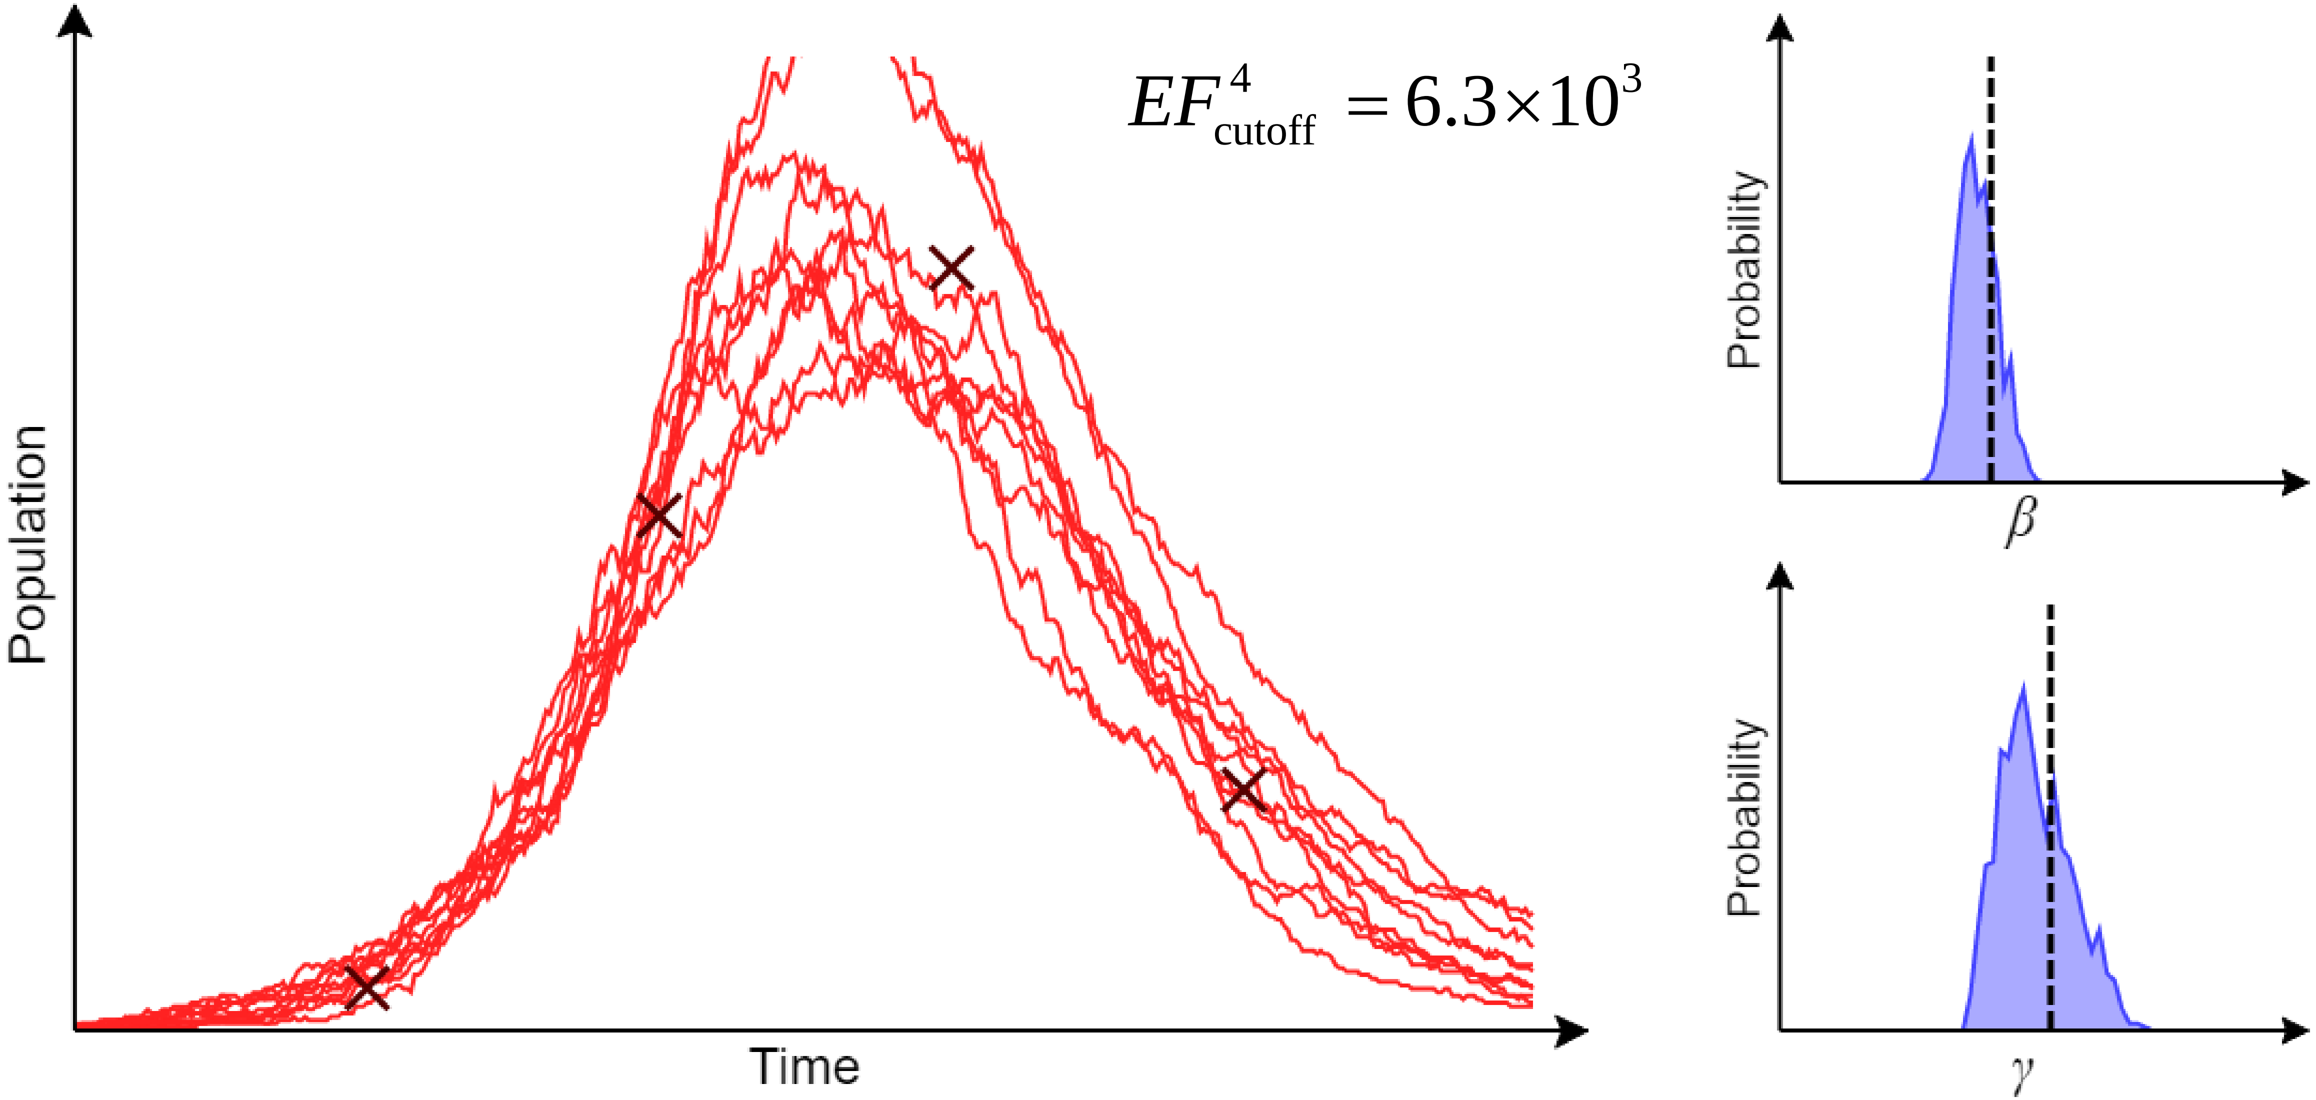

ABC-MBP

Generation 5:

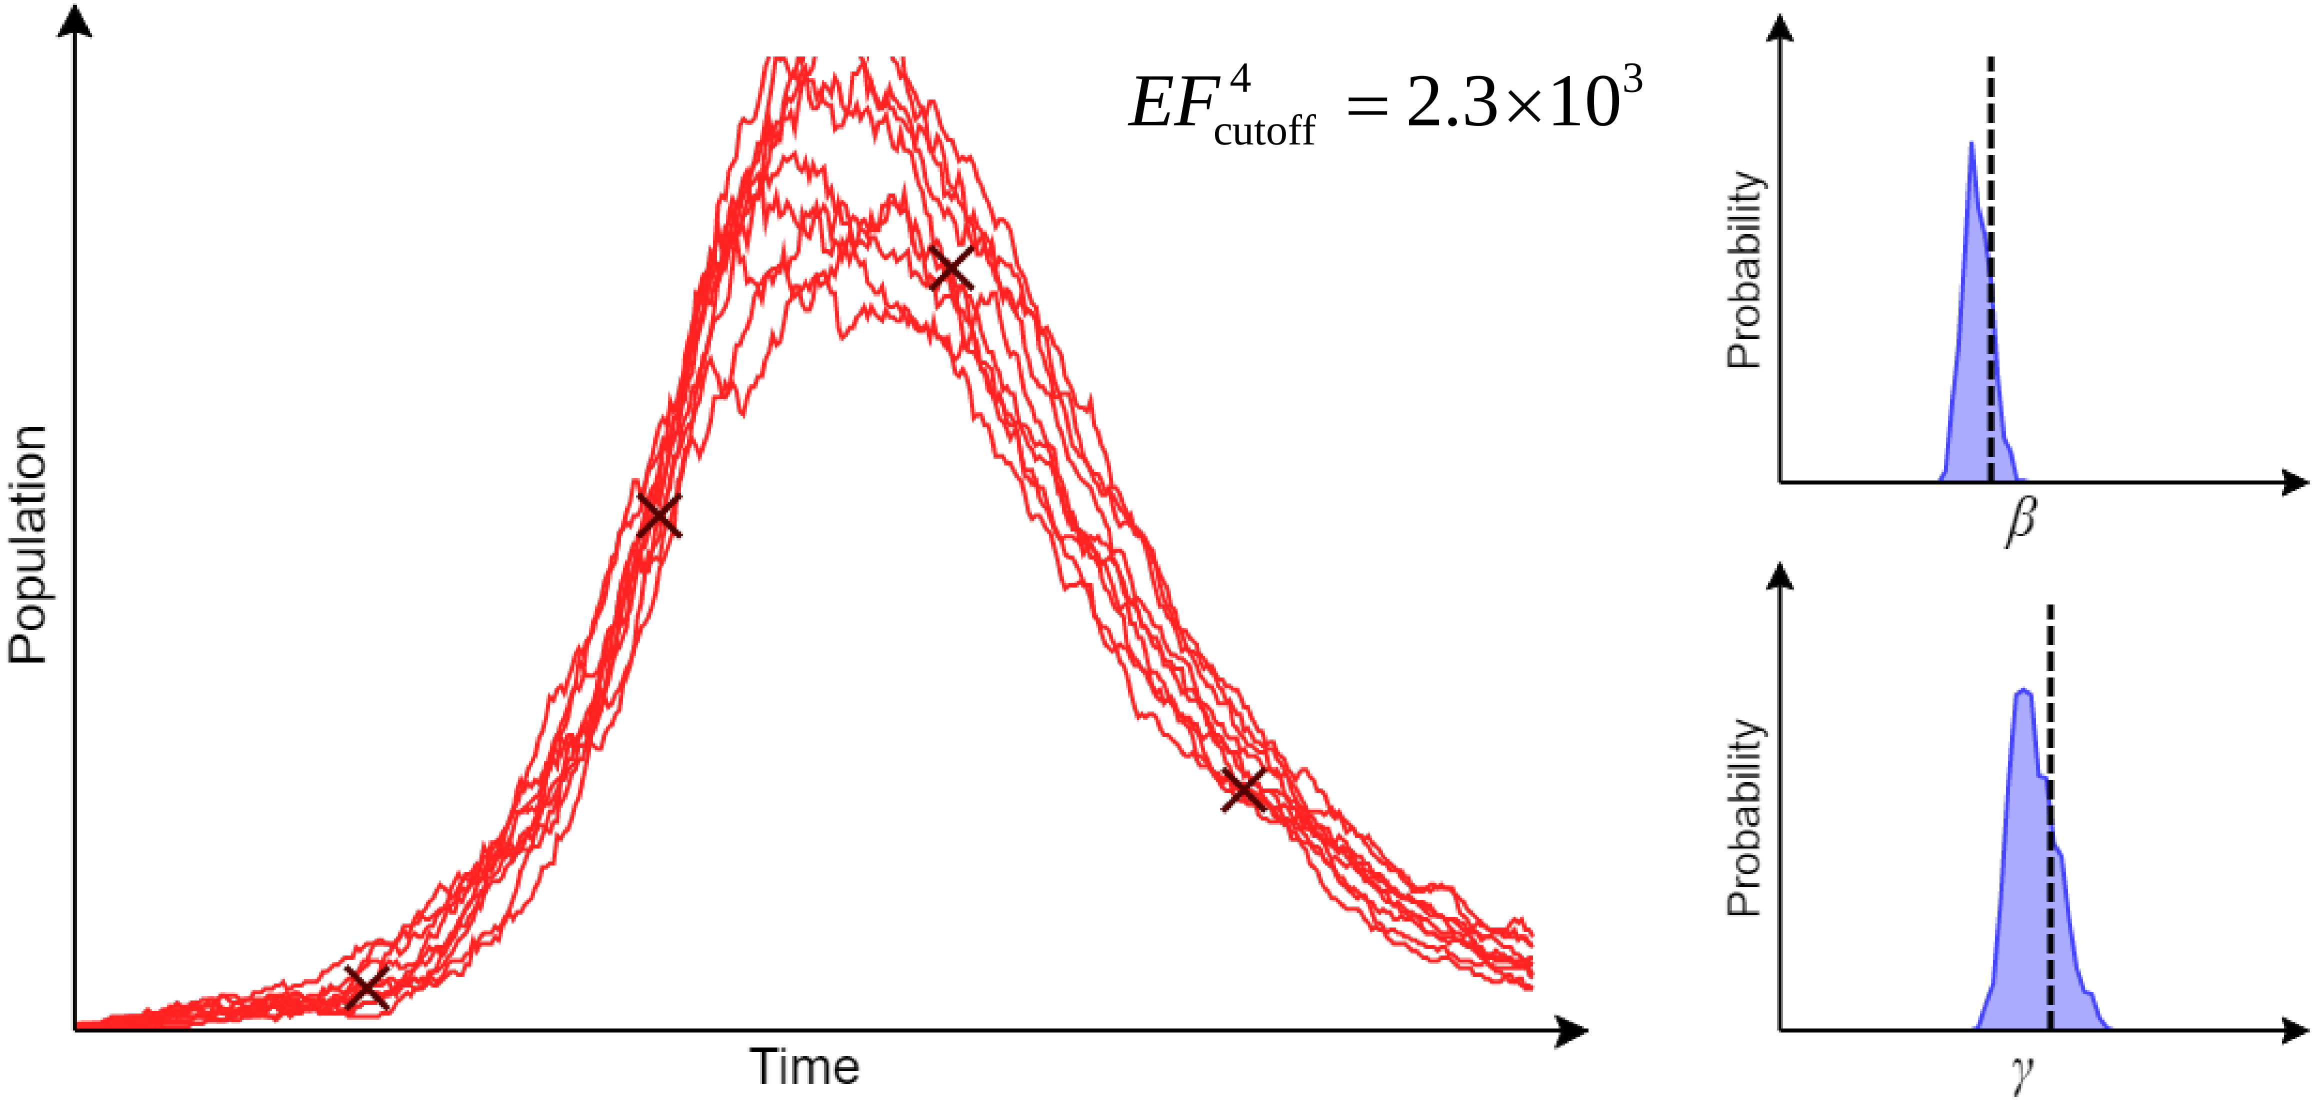

Model based proposals (MBPs)

Covariance estimate obtained from all particles

- STEP 1: Propose a new parameter set $\theta'_p \sim \mathrm{MVN}(\theta_p, s^g \Sigma^g),$

Tuneable constant

- STEP 2: *Modify* ξ_p to generate ξ'_p
- STEP 3: Calculate EF'_p
- STEP 4: If $EF'_p \geq EF_{cut}$ reject,

  otherwise accept with probability: $\max\left\{\frac{\pi(\theta'_p)}{\pi(\theta_p)}, 1\right\}$

Posterior-based proposals for speeding up Markov chain Monte Carlo

C. M. Pooley, S. C. Bishop, A. Doeschl-Wilson and G. Marion (2019) Royal Society Open Science

BEEPmbp

- "Bayesian Estimation of Epidemic Parameters using Model-Based Proposals"
  - Open source software tool for fitting epidemiological models
  - Parallel implementation suitable for HPC

- Supports potential spatial and demographic stratification

- Different data types:
  - Time-series transition data (*e.g*. daily cases)
  - Time-series population data (*e.g.* numbers in hospital)
  - Marginal distributions (*e.g.* age distribution of deaths)

github.com/ScottishCovidResponse/BEEPmbp

Benchmark models for speed comparison

Complexity

Stochasticity

Simple

$$\lambda_t = \frac{R}{N\mu_I} I_t + \eta$$

- 3 parameters
- R fixed
- Non-spatial
- 2 time-series data

Simple Stochastic $N = 10^3$	**Simple Continuum** $N = 10^5$

Medium

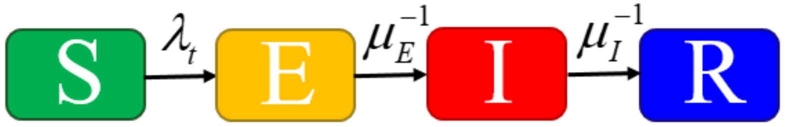

$$\lambda_t = \frac{R_t}{N\mu_I} I_t + \eta$$

- 13 parameters
- R_t time-varying
- Non-spatial
- 3 time-series data

Medium Stochastic $N = 10^4$	**Medium Continuum** $N = 10^6$

Hard

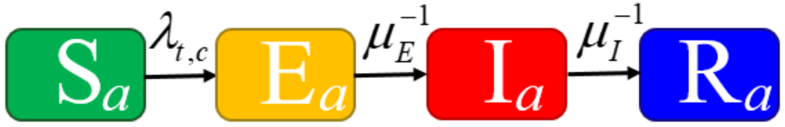

$$\lambda_{t,a} = \frac{R_t e^{\alpha_a}}{N_a \mu_I} \sum_{a'} M_{a',a} I_{t,a} + \eta$$

- 36 parameters
- R_t time-varying
- 10 spatial locations
- 30 time-series data

Hard Stochastic $N = 10^5$	**Hard Continuum** $N = 10^7$

Simple Model: SIR

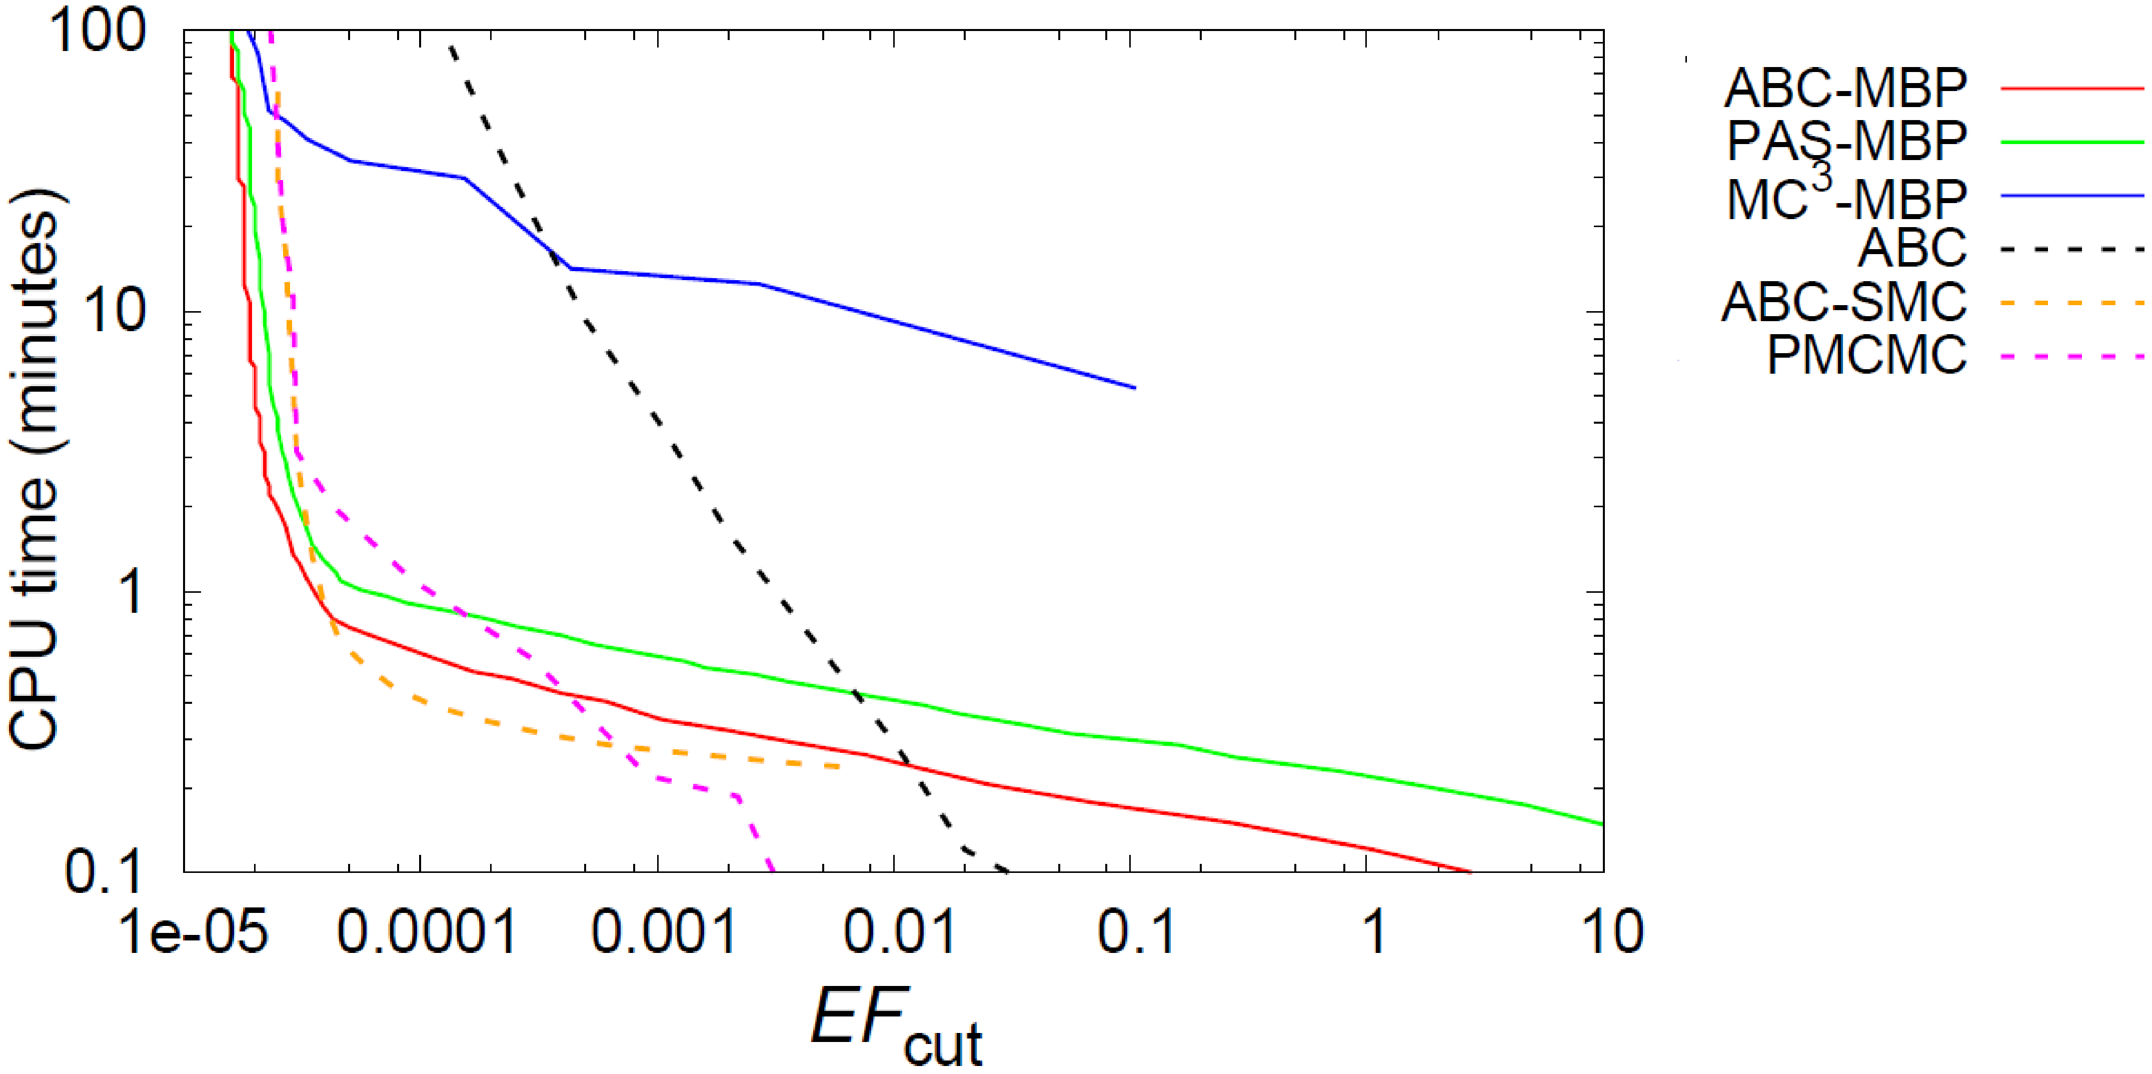

Simple Model: SIR

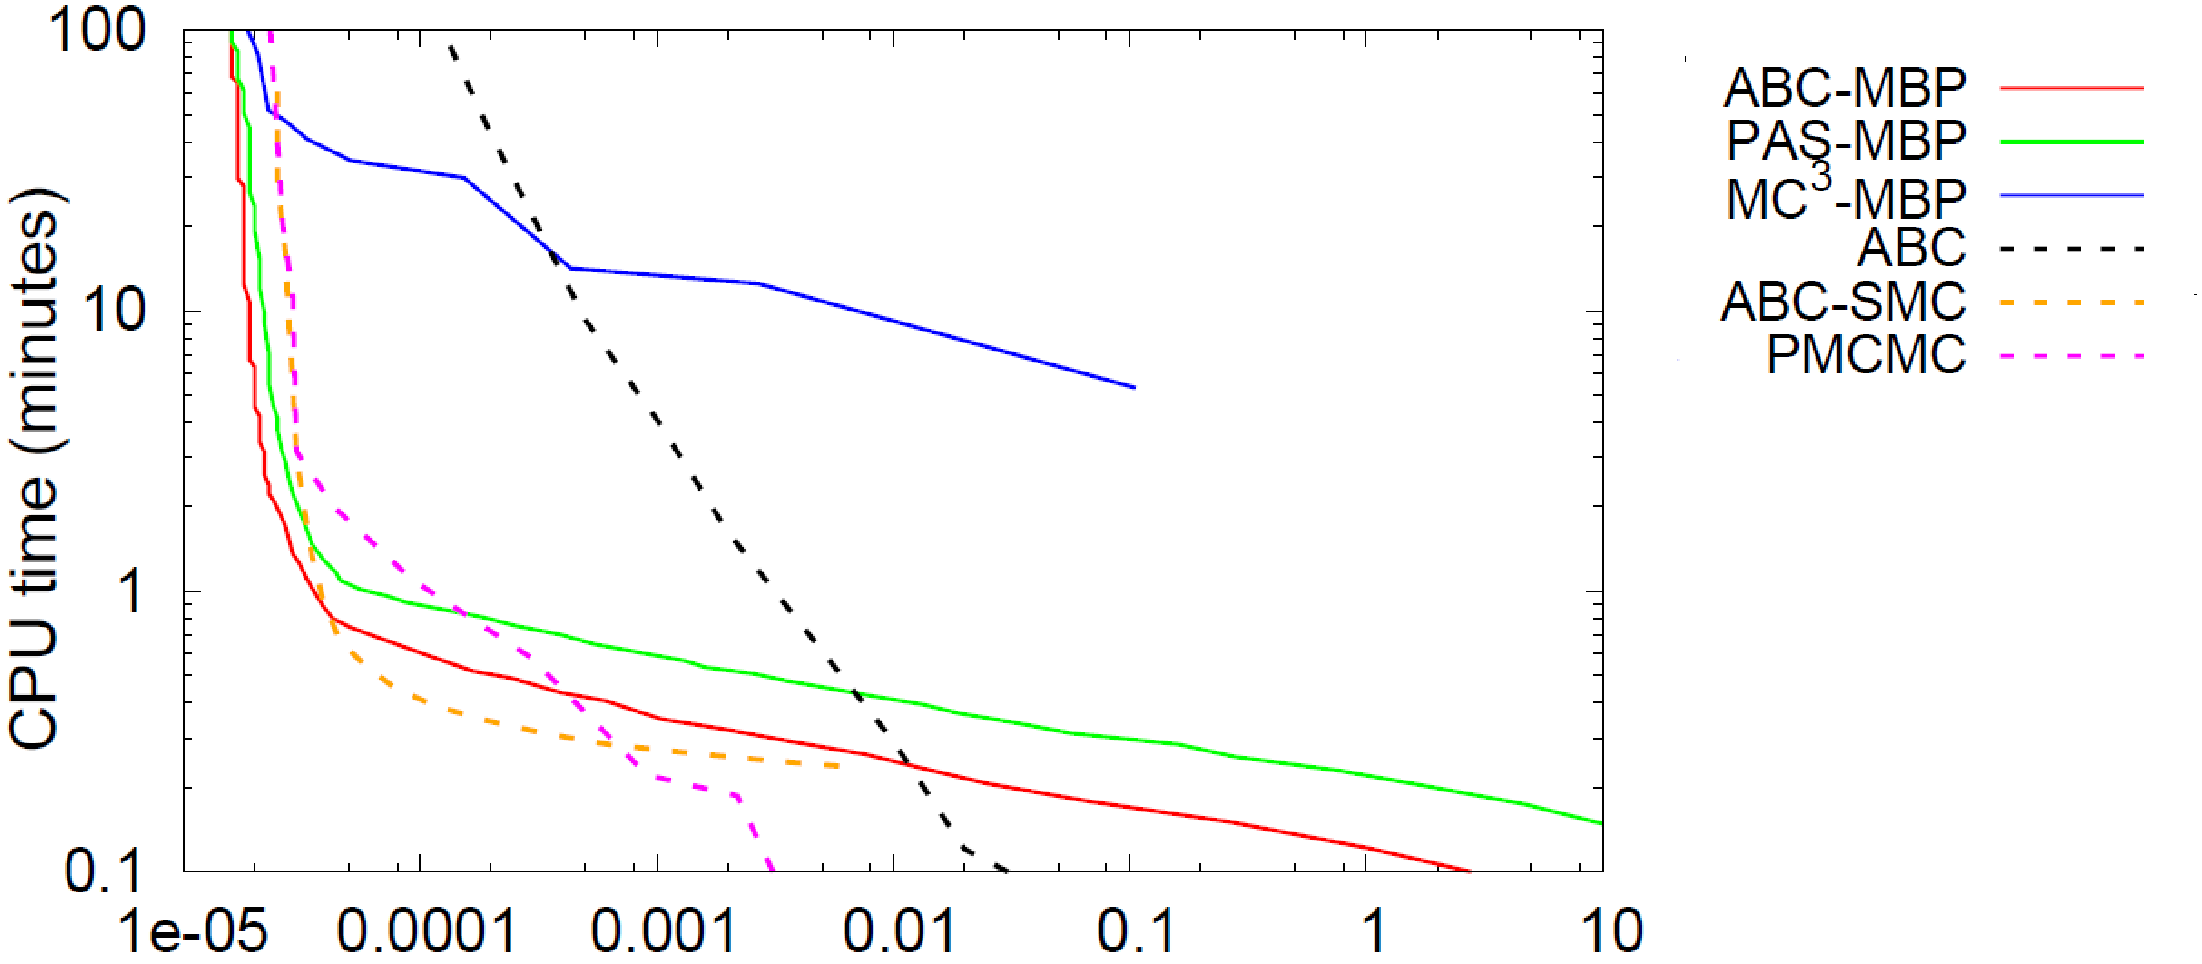

Good posterior estimate

EF_{cut}

Poor posterior estimate

Simple Model: SIR

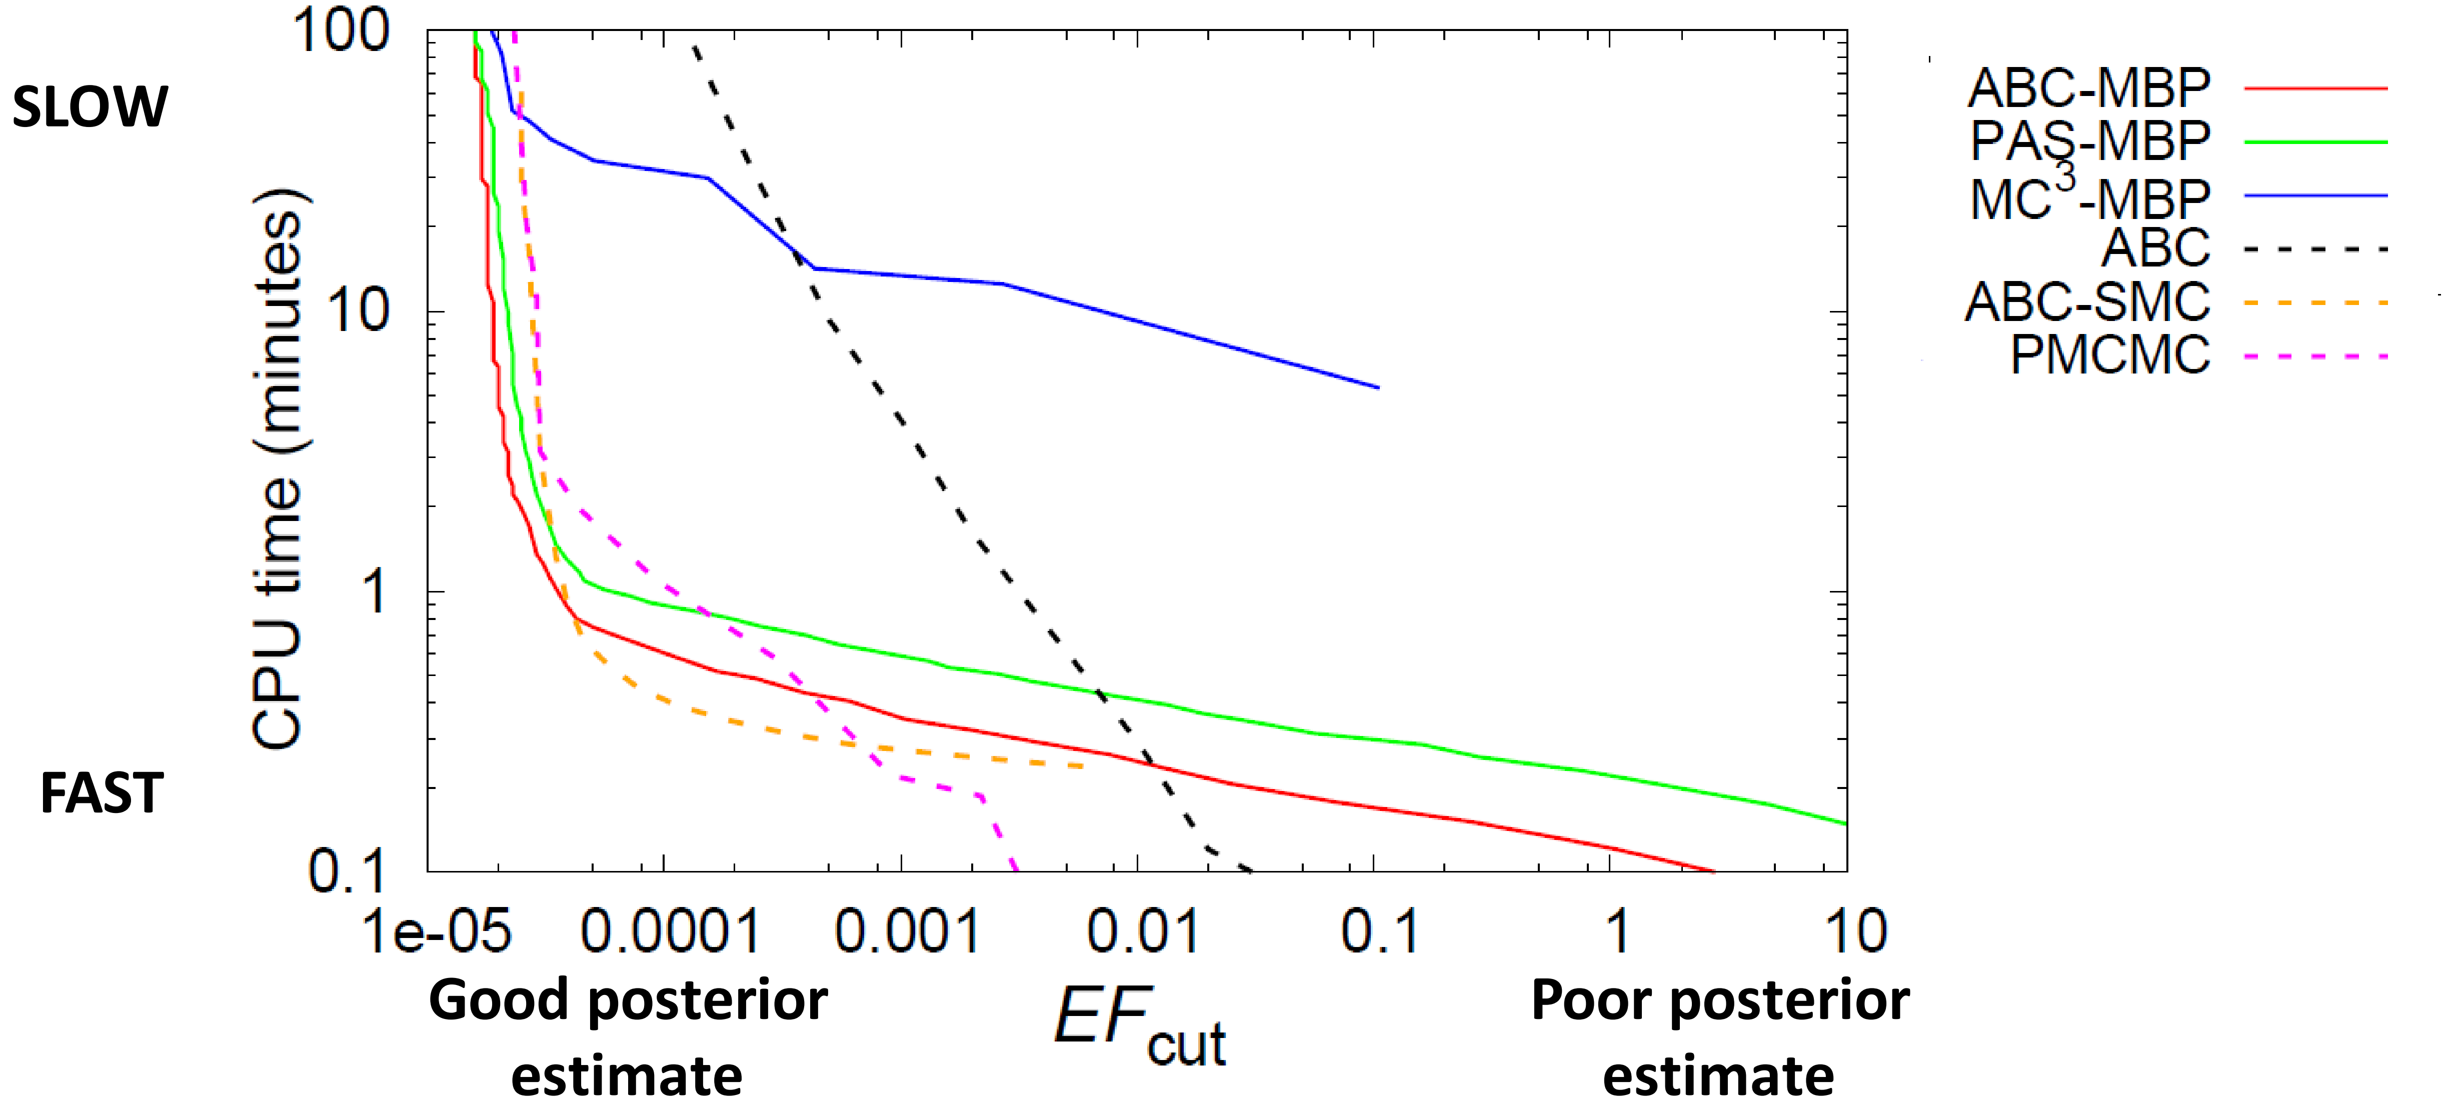

Complex Model: Spatial SEIR with time variation

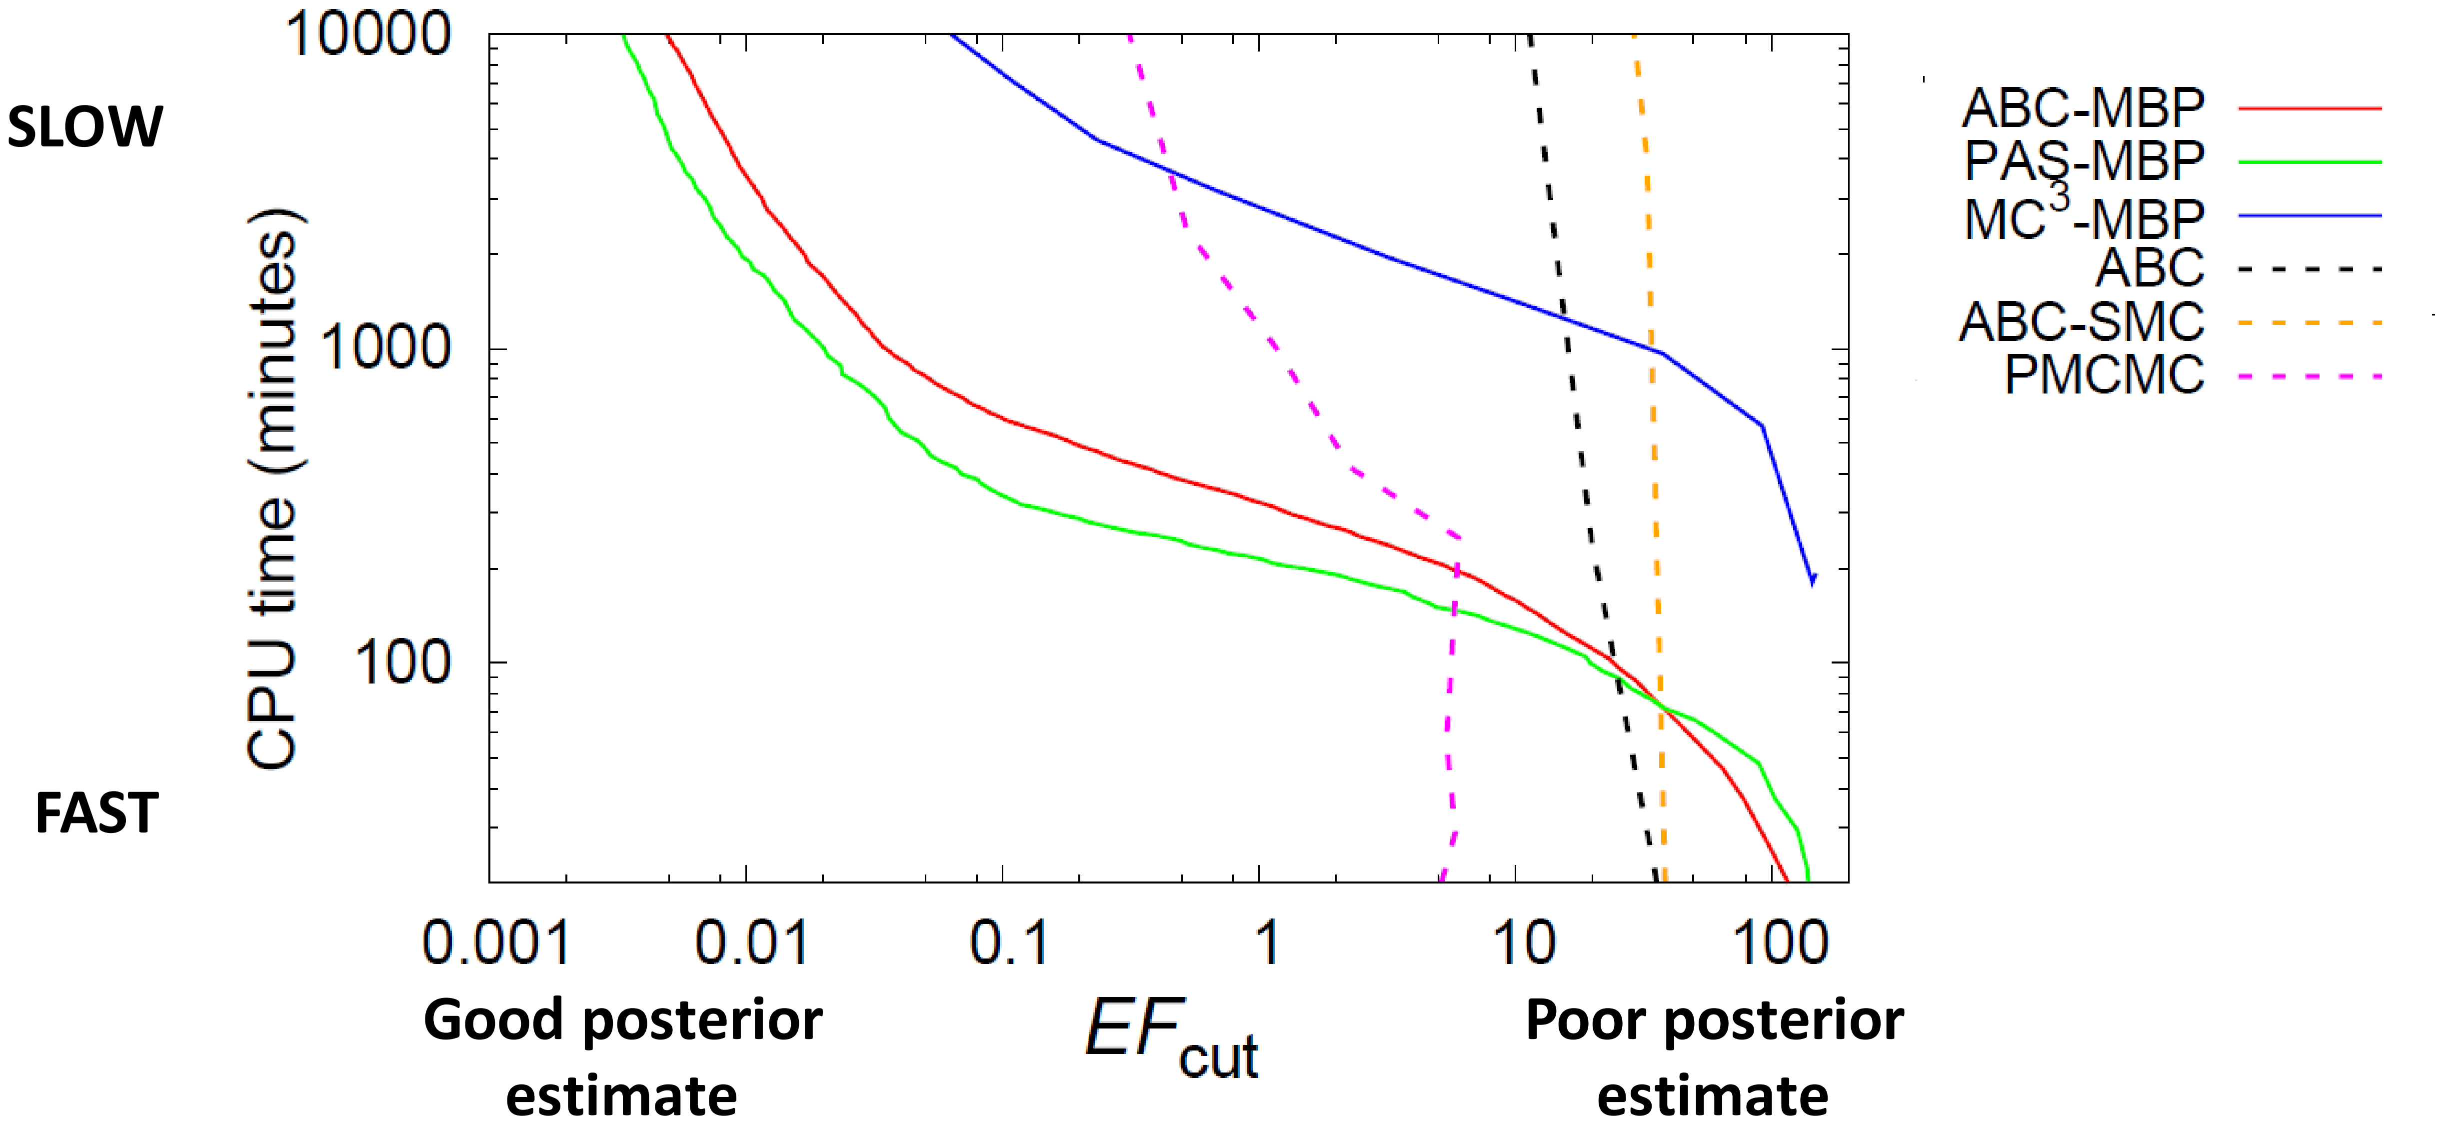

Application to COVID-19

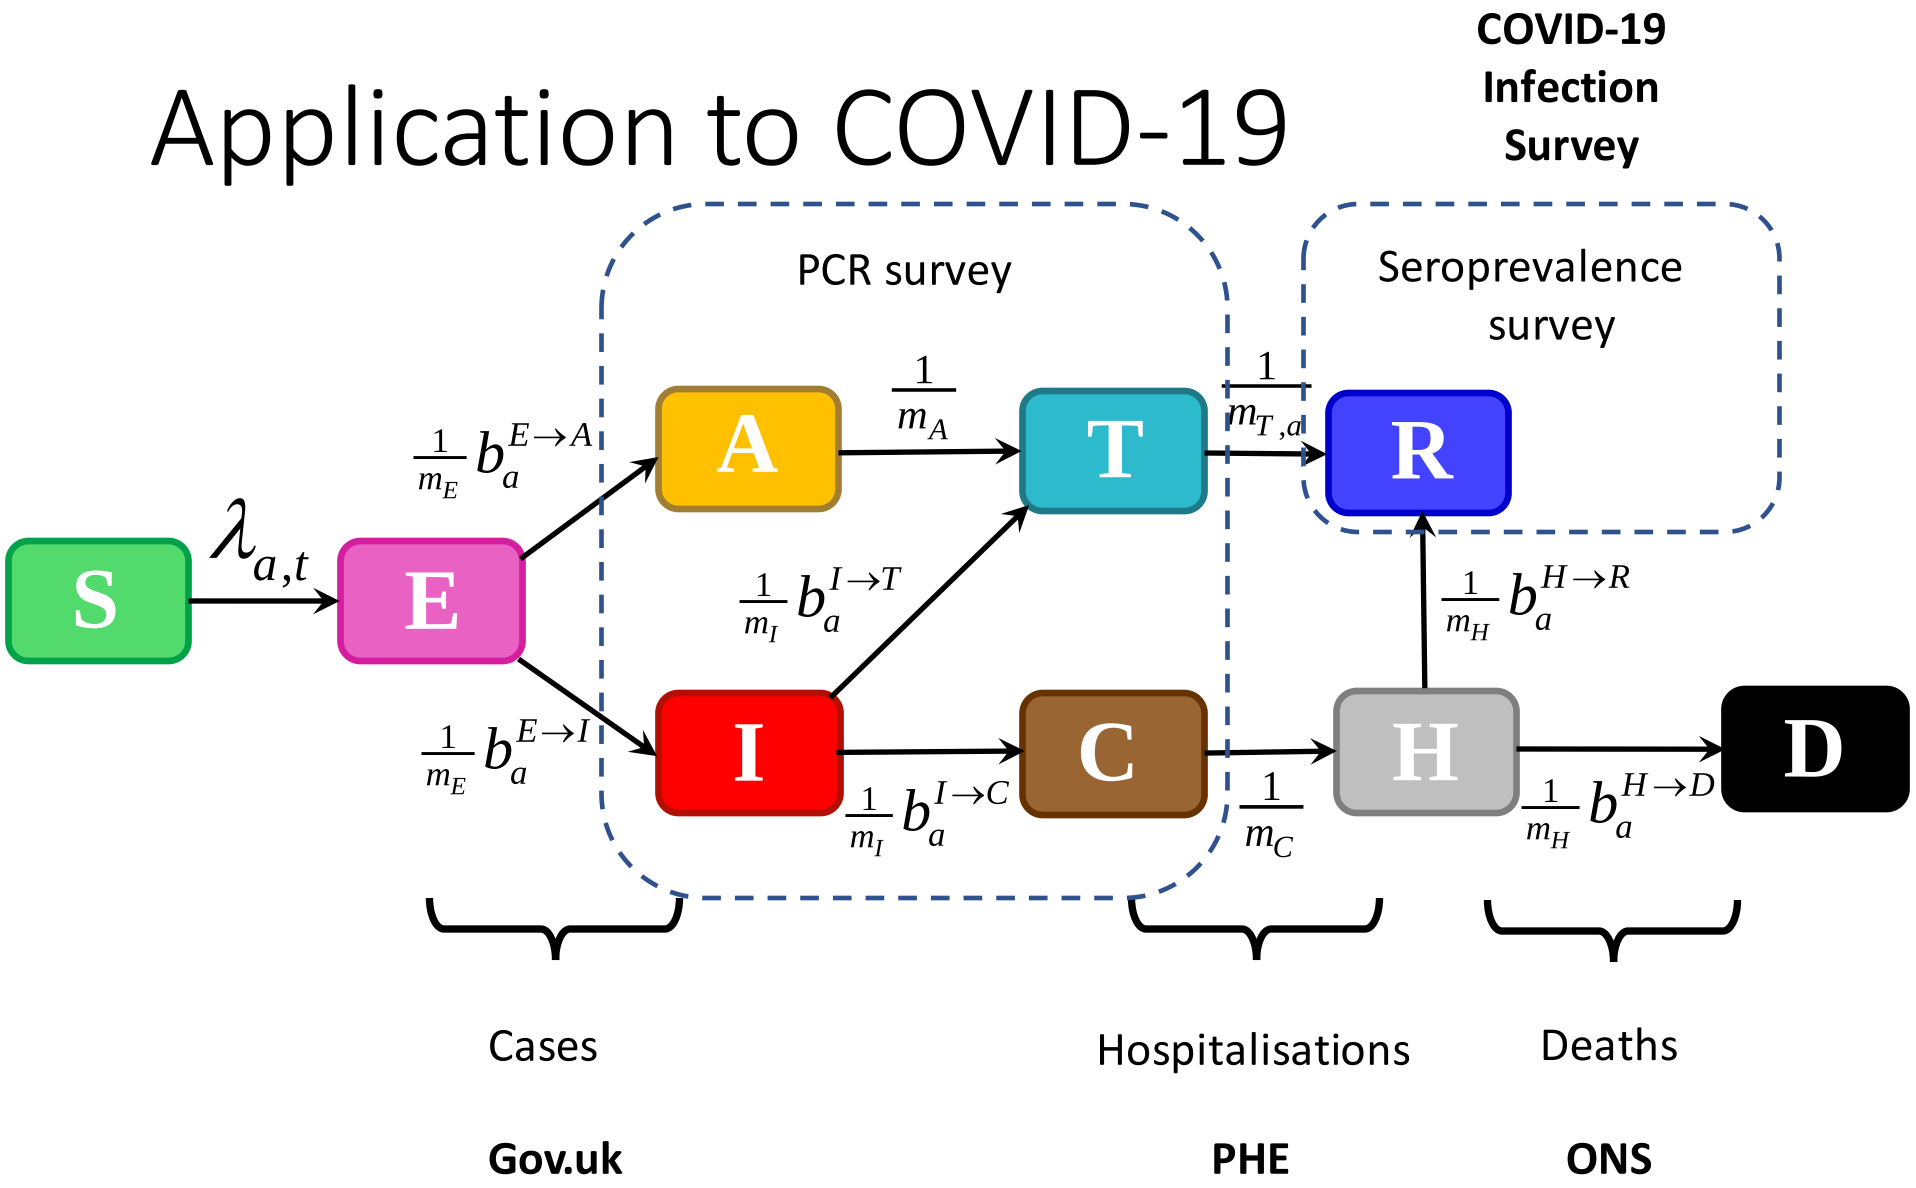

S: Susceptible
E: Exposed
A: Asymptomatic
I: Symptomatic
T: Test-sensitive
R: Recovered
H: Hospitalised
C: Critically ill
D: Dead

18 age categories a={0-4, 5-9,…, 75-79, 80+, Care homes}

Application to COVID-19

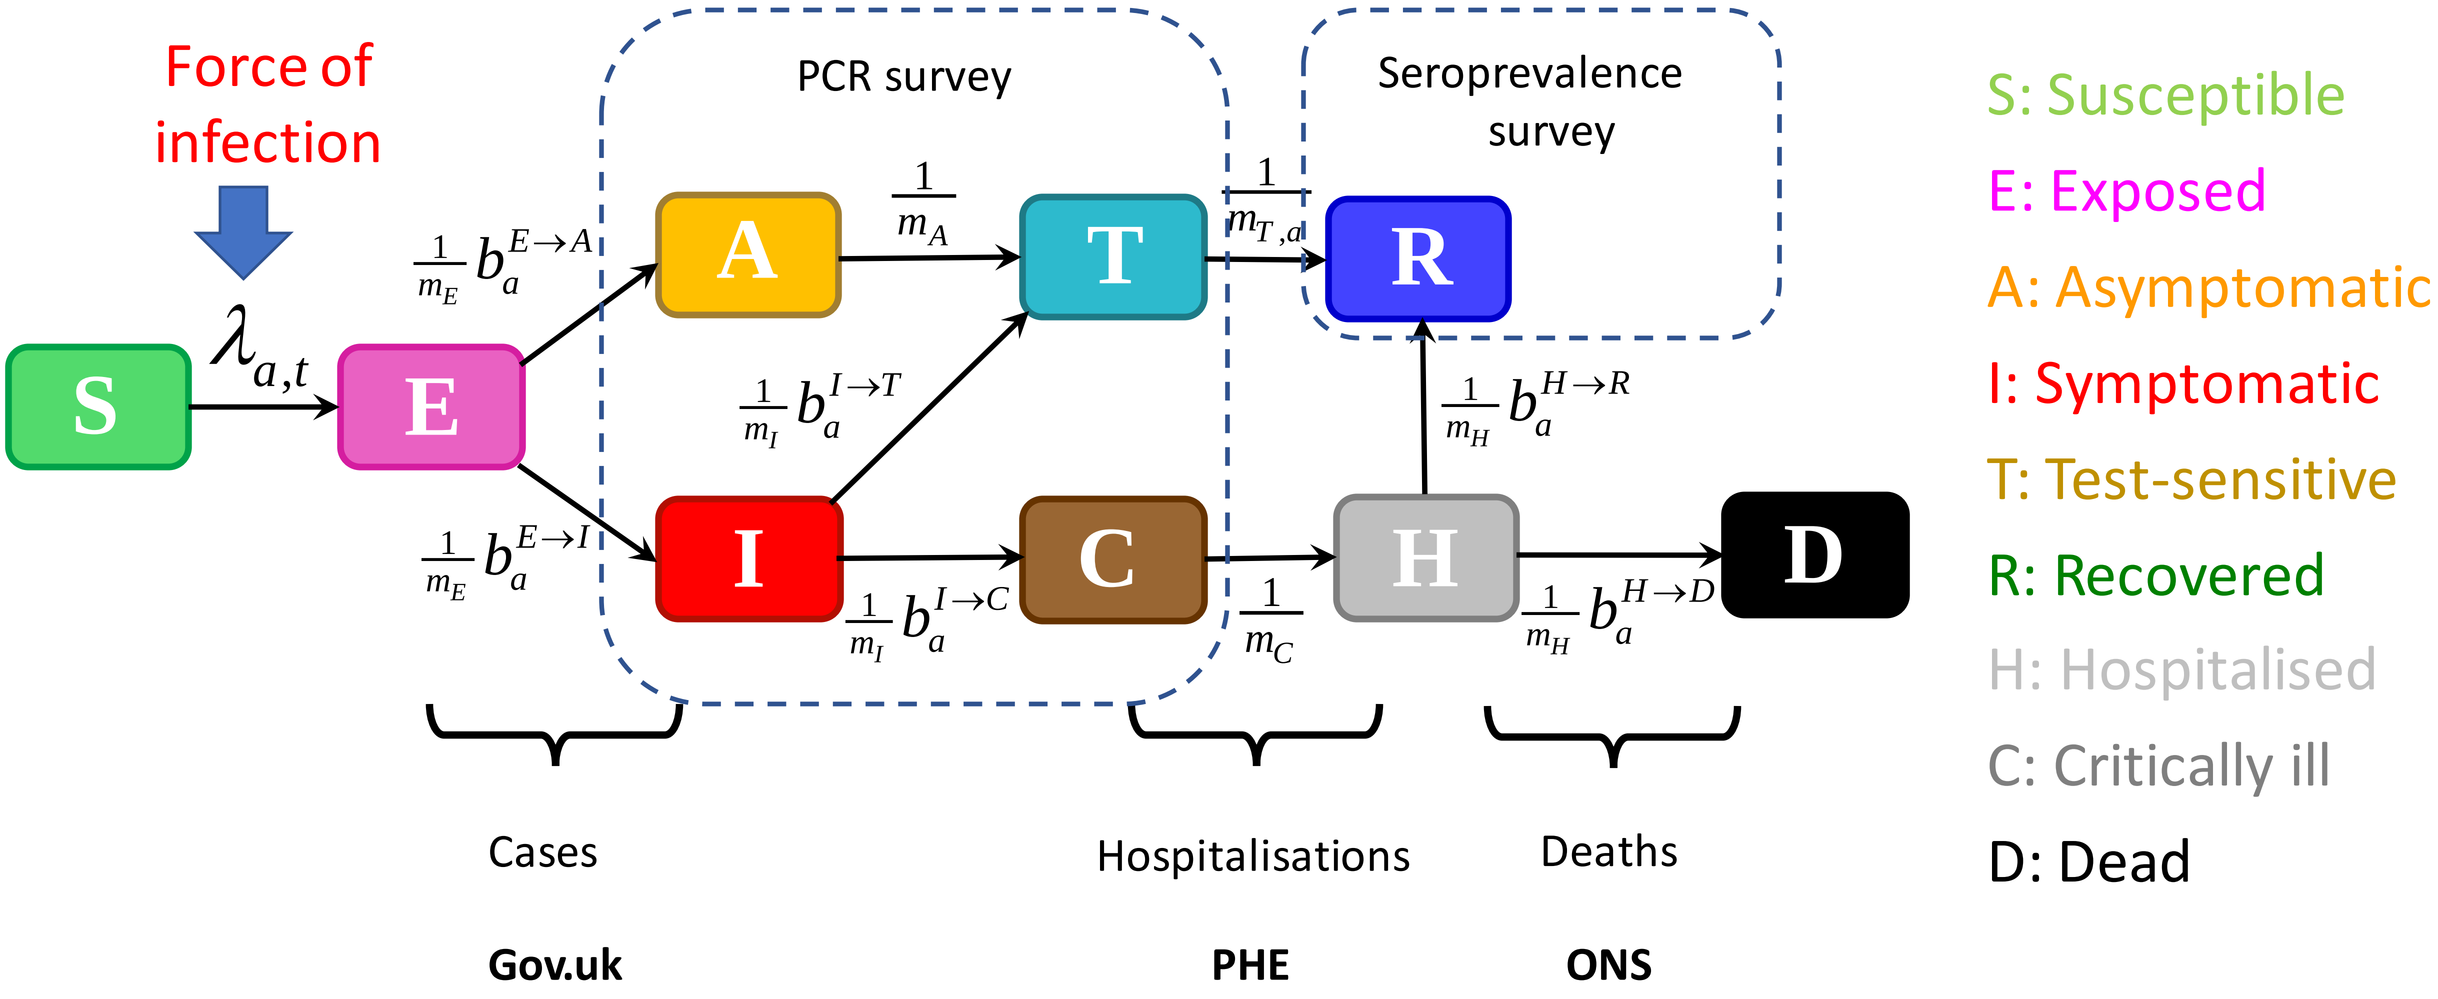

18 age categories a={0-4, 5-9,…, 75-79, 80+, Care homes}

Force of infection

Transmission rate

Infectivity of state c

individuals in state c in age group a' at time t

$$\lambda_{a,t} = \beta\left[\frac{1}{P_a}\sum_{a',c} C_{t,a,a'}\varphi_c N_{c,a',t}\right] + \frac{1}{f}\eta_t$$

Force of infection in age group a at time t

Age contact matrix

External source of infection

Age contact matrix C

- A baseline contact matrix C^0 is taken from the BBC Pandemic! study

$$C_{t,a,a'} = f_t v_a C^0_{a,a'} v_{a'}.$$

- This is altered by
  - A time varying factor which accounts for government interventions and changing social behaviour
  - Age modifying factor

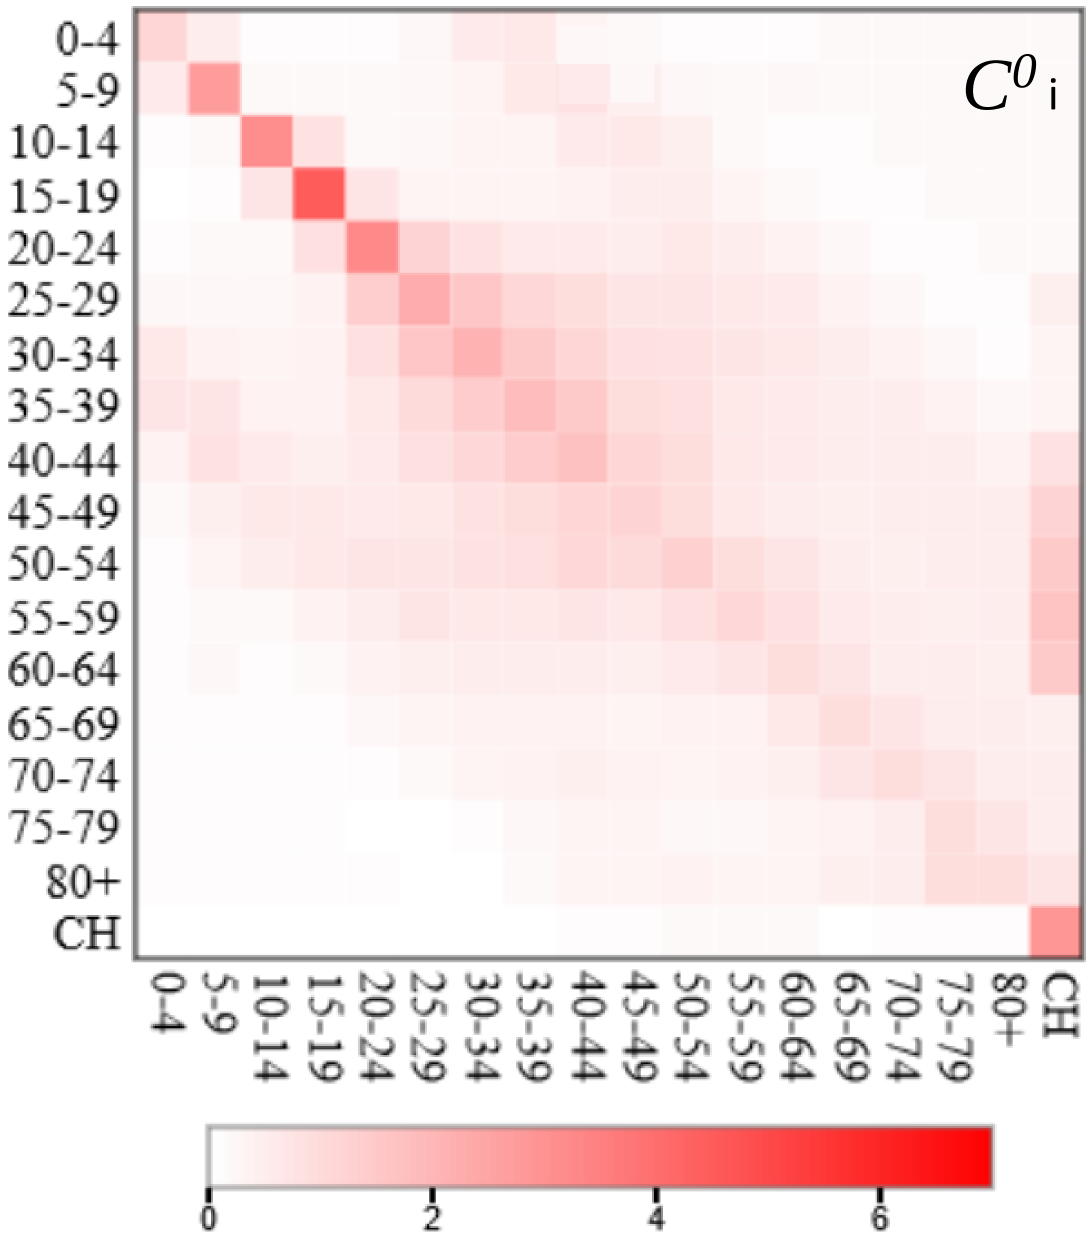

External force of infection η

- COVID-19 enters the UK via individuals moving to and from other countries

- Estimate external force of infection by using:
  - Global flight data from CAA
  - Global Covid-19 data from Johns Hopkins University

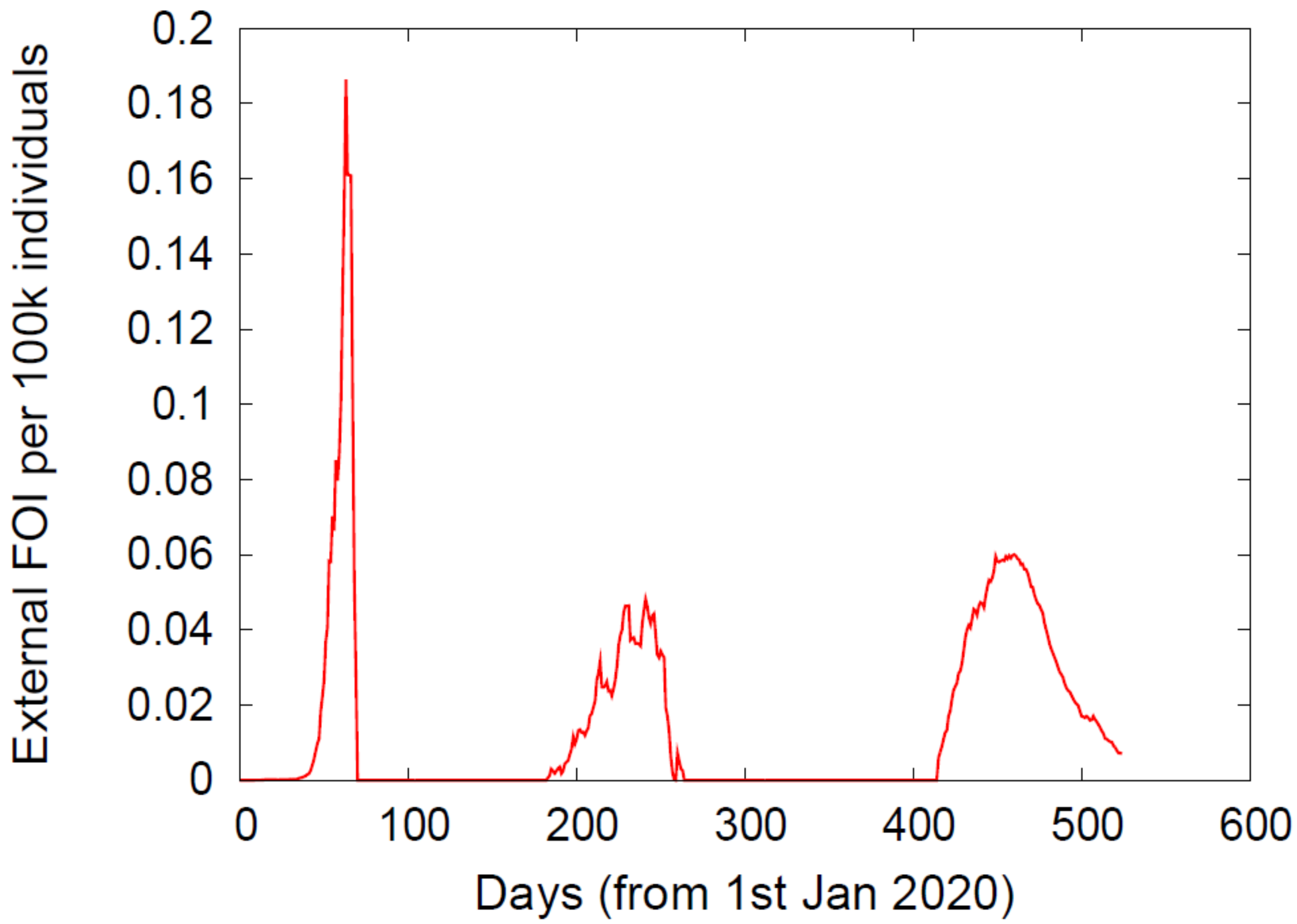

Inference

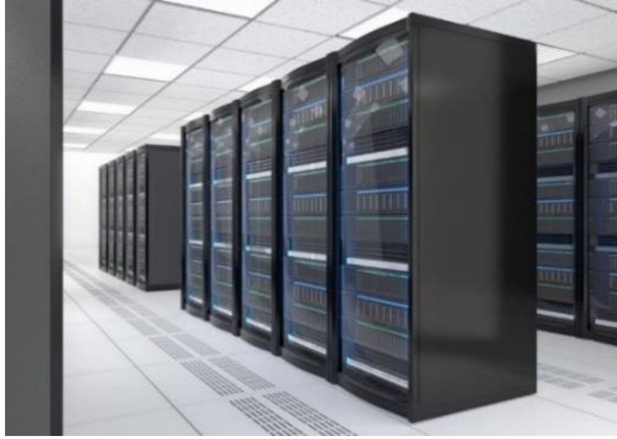

DiRAC

- Model contains ~150 parameters
- Data contains ~9000 observations in total
    - Spread across many time series
- Inference performed using ABC-MBP
    - DiRAC HPC
    - ~5 hours when running on 256 cores

Results: Reproduction number

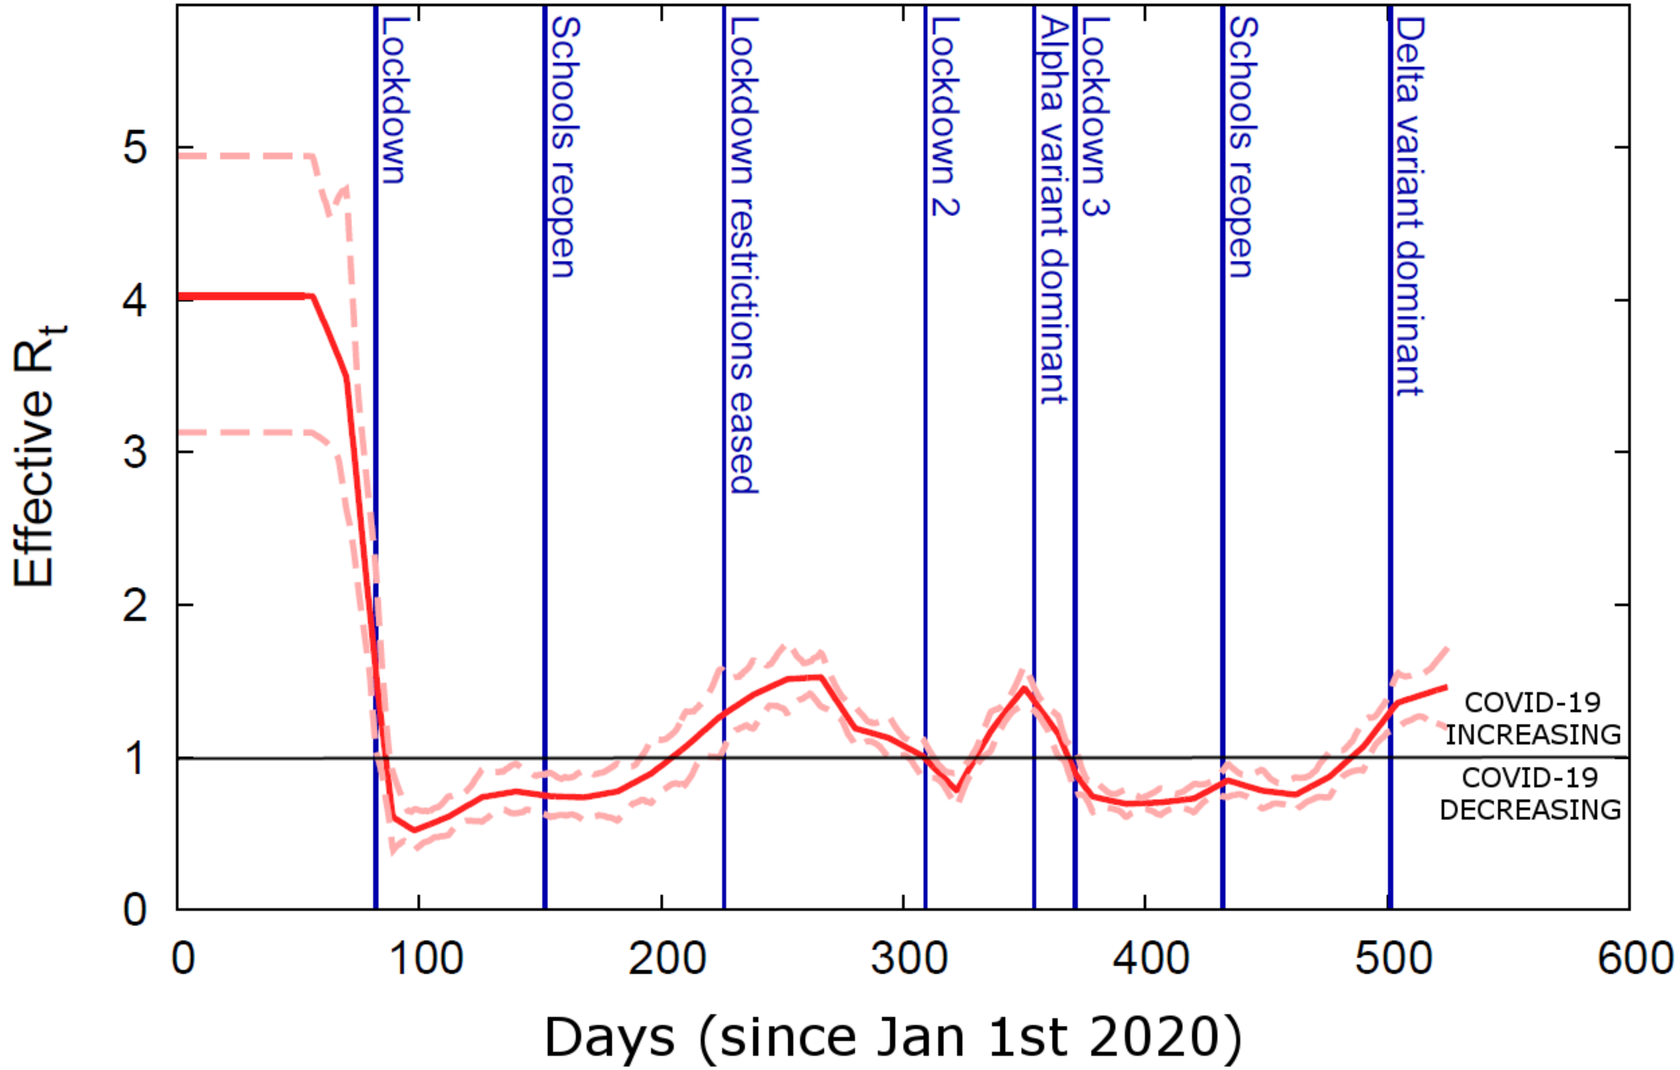

Compartmental parameter estimates

- Obtain a fully parameterised age-structured model

- Improved predictions and simulation studies

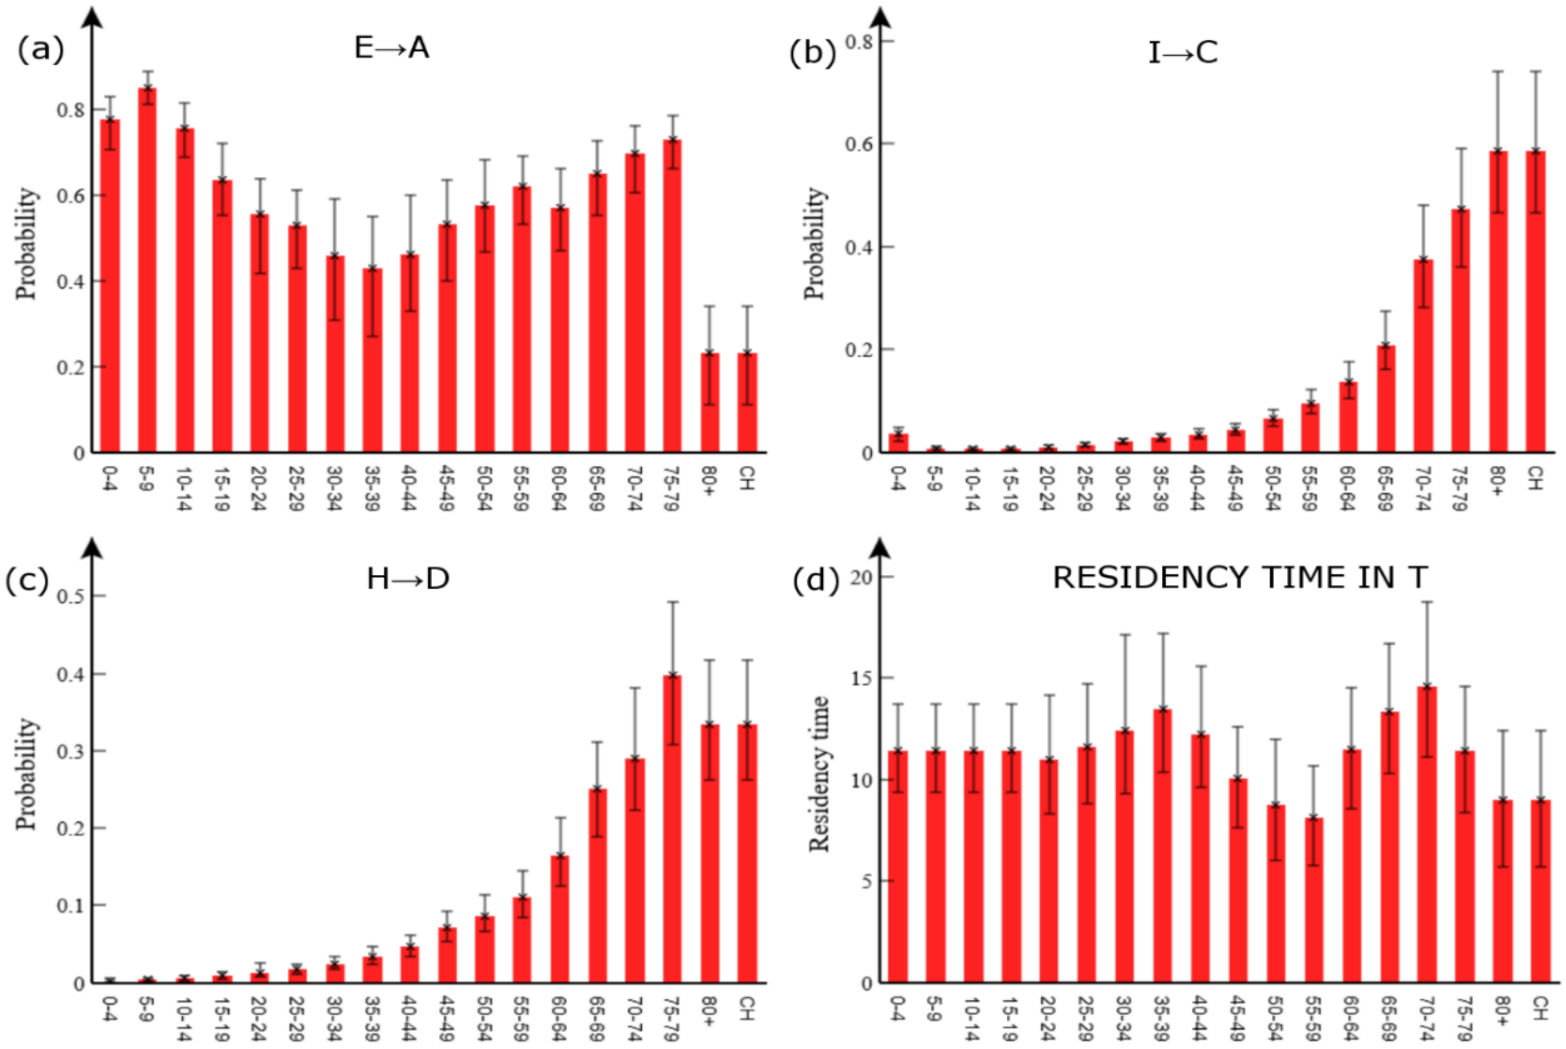

Age contact matrix C

- This shows the relative modification in the contact matrix compared to the expectation obtained from the BBC Pandemic study

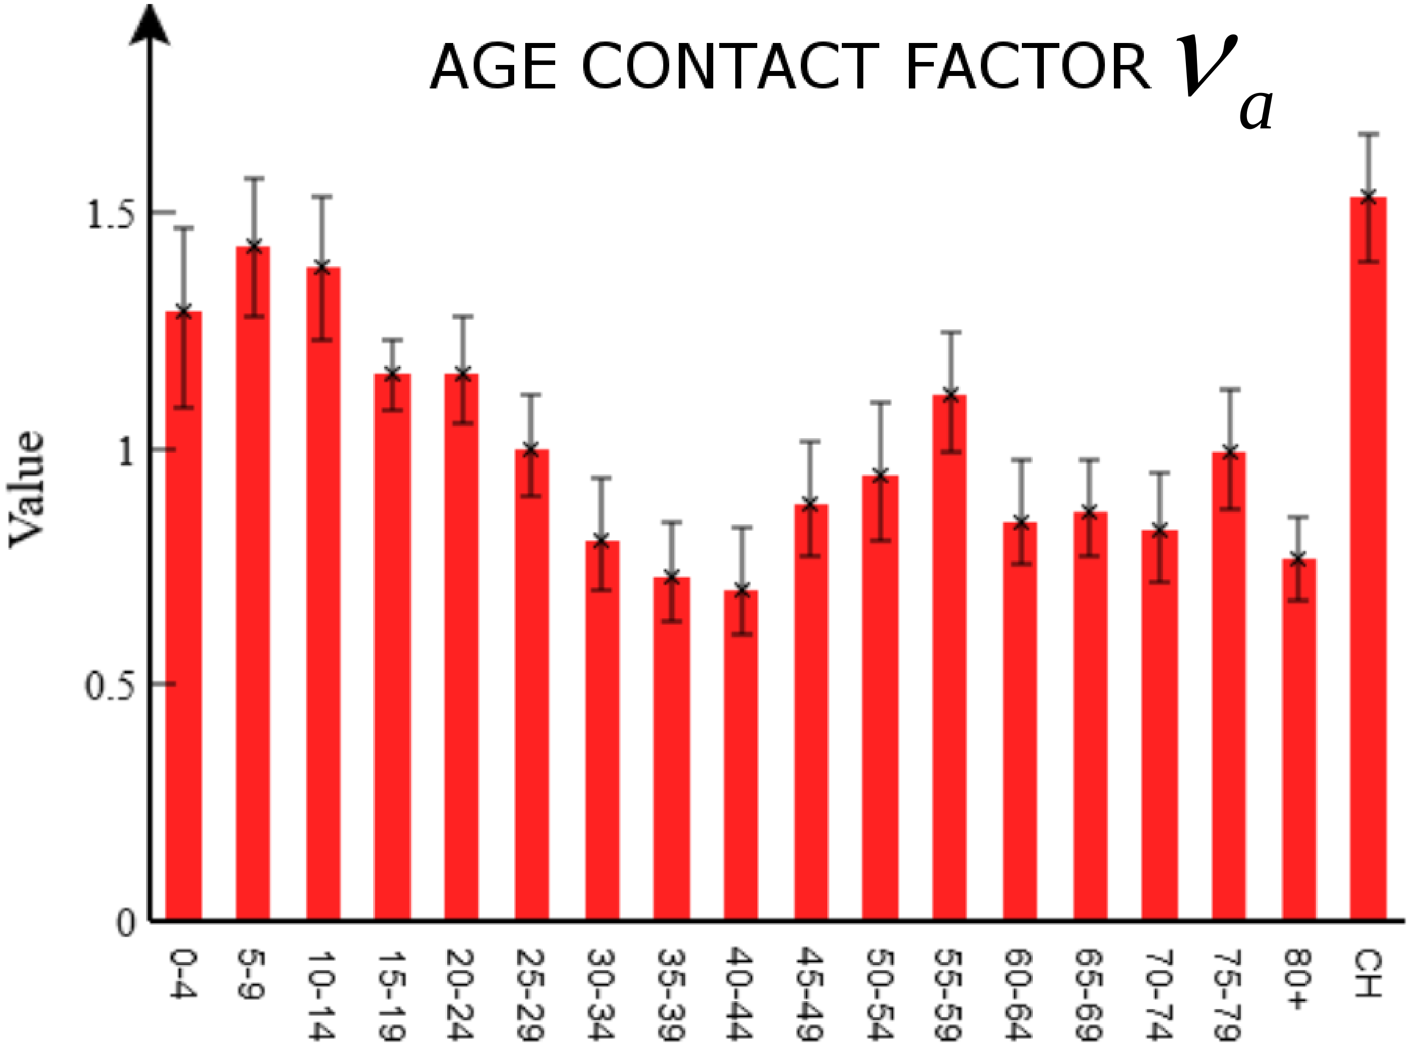

Age contact matrix C

- This shows the relative modification in the contact matrix compared to the expectation obtained from the BBC Pandemic study

- We find higher relative rates of effective contacts in younger individuals compared to older individuals

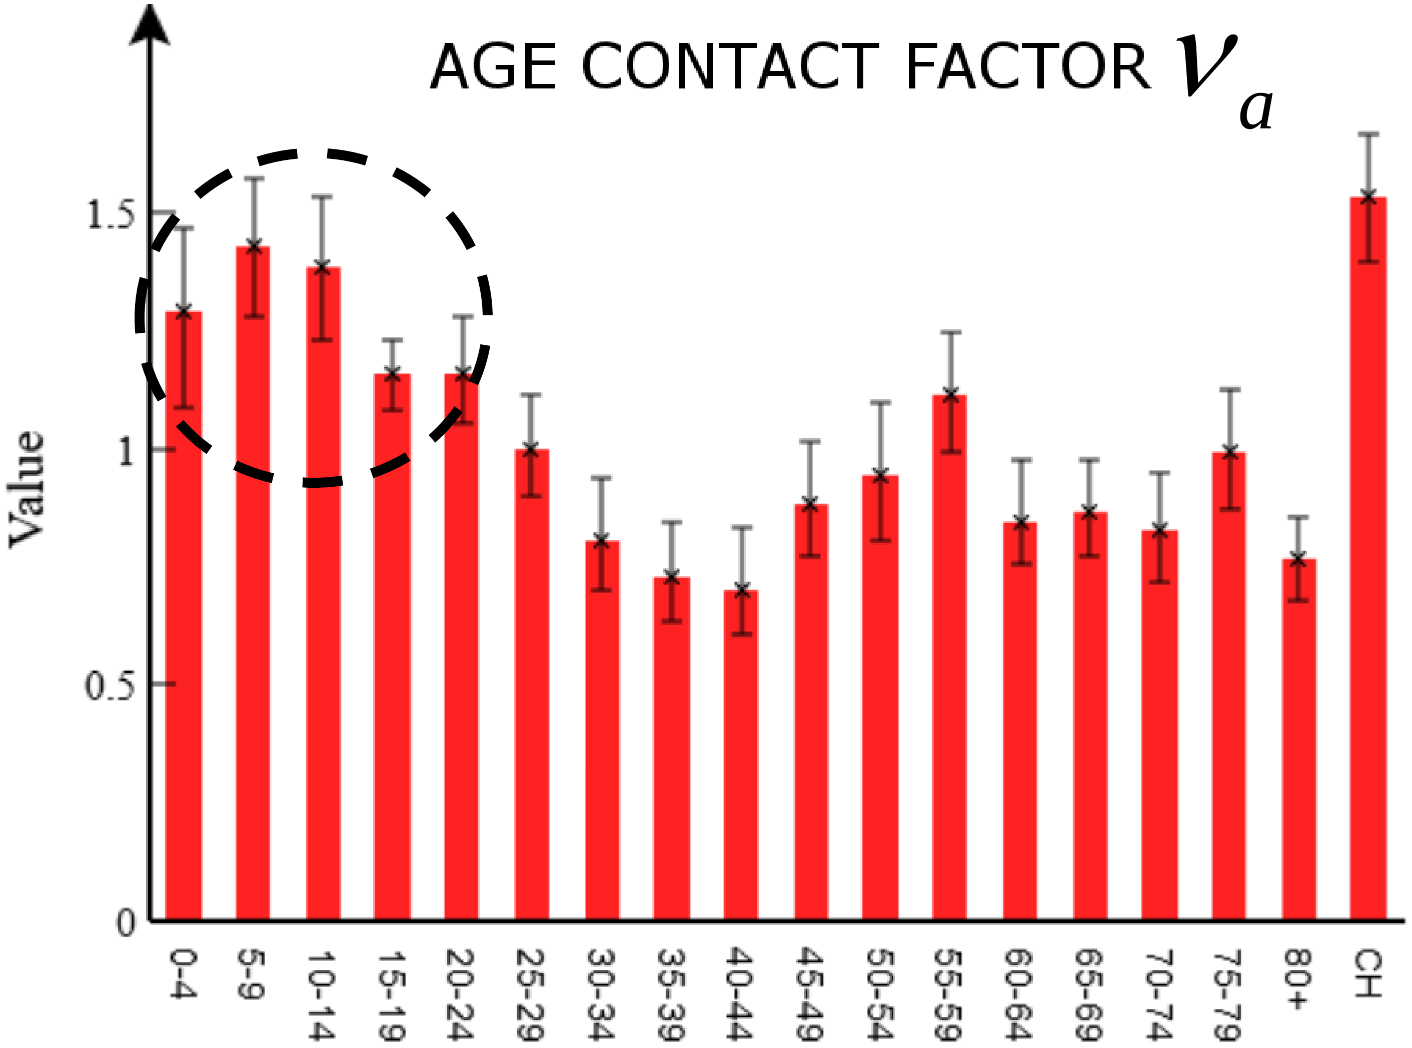

Summary

- Developed a new inference methodology ABC-MBP
  - Orders of magnitude faster than other existing approaches
  - Generic and not specific to Covid-19 or epidemiology
- Created software tool BEEPmbp
  - Incorporate many inference algorithms
  - Flexible to different model specification and data types
- Application to an age-structure model of Covid-19

Future

- Covid-19 provides a huge data source
  - Many opportunities for new ways to analyse
- Current work…
  - Spatial models nationally and internationally
  - Trying to assess the impact of government interventions and travel restrictions
- What can we apply to the next animal disease outbreak?
- What should we do differently next time?

Acknowledgments

Supervisors:

Glenn Marion (BioSS)
Andrea Doeschl-Wilson (The Roslin Institute)

Data curators:

Richard Bailey (The Roslin Institute)
Martin Burke (BioSS)
Helen Brown (The Roslin Institute)

Also:

Scottish Covid-19 Response Consortium (SCRC)
Royal Society RAMP volunteers

RSEs:

Ian Hinder (Manchester)
Robin Williams (Bristol)

Funding:

epic – Centre of Expertise on Animal Disease Outbreaks

Scottish Government
Riaghaltas na h-Alba
gov.scot

RESAS

Data-Driven Innovation
Part of the Edinburgh & South East Scotland City Region Deal