Bayesian calibration of differentiable agent-based models

Arnau Quera-Bofarull, Ayush Chopra, Anisoara Calinescu,
Michael Wooldridge, Joel Dyer

AI4ABM workshop — ICLR 2023

Calibration of ABMs

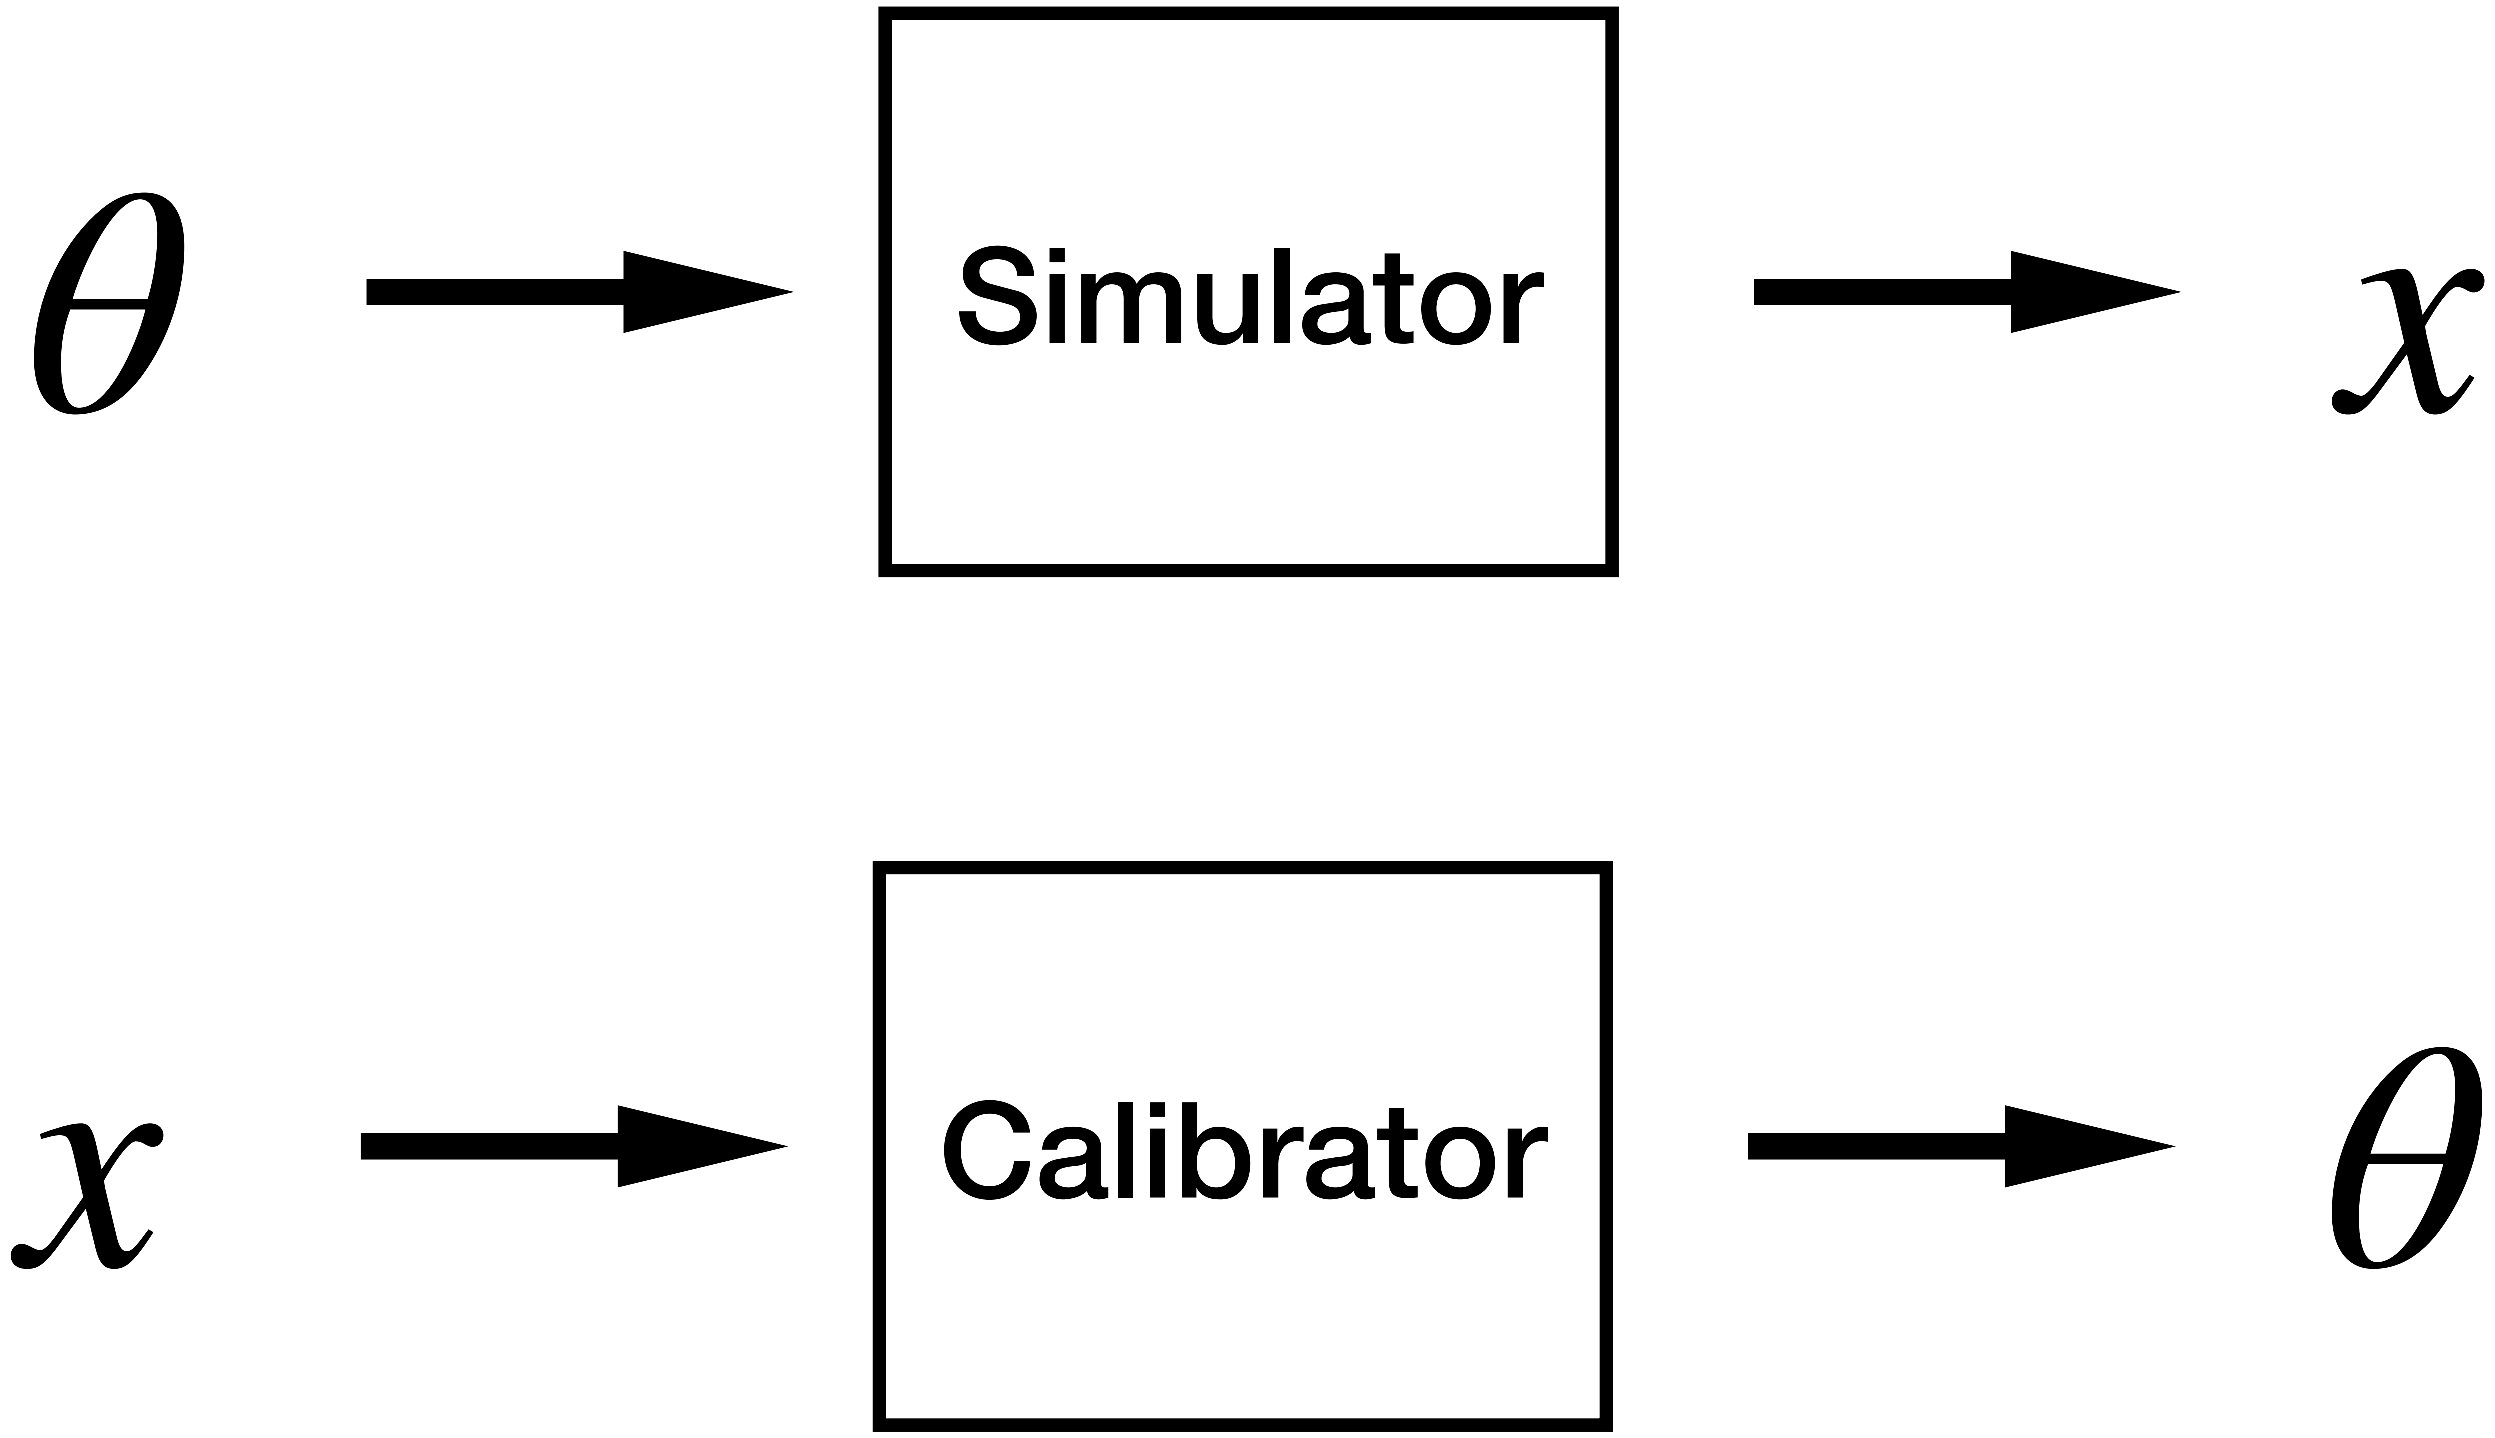

Why is it hard?

- Expensive simulator

+

- Large parameter space

Calibration requirements

1. Uncertainty quantification

Ideally we want to get all θ that can generate *x* with a certain probability

Example

Epidemiological model with 2 parameters:

1. Reproduction number at schools
2. Reproduction number at companies

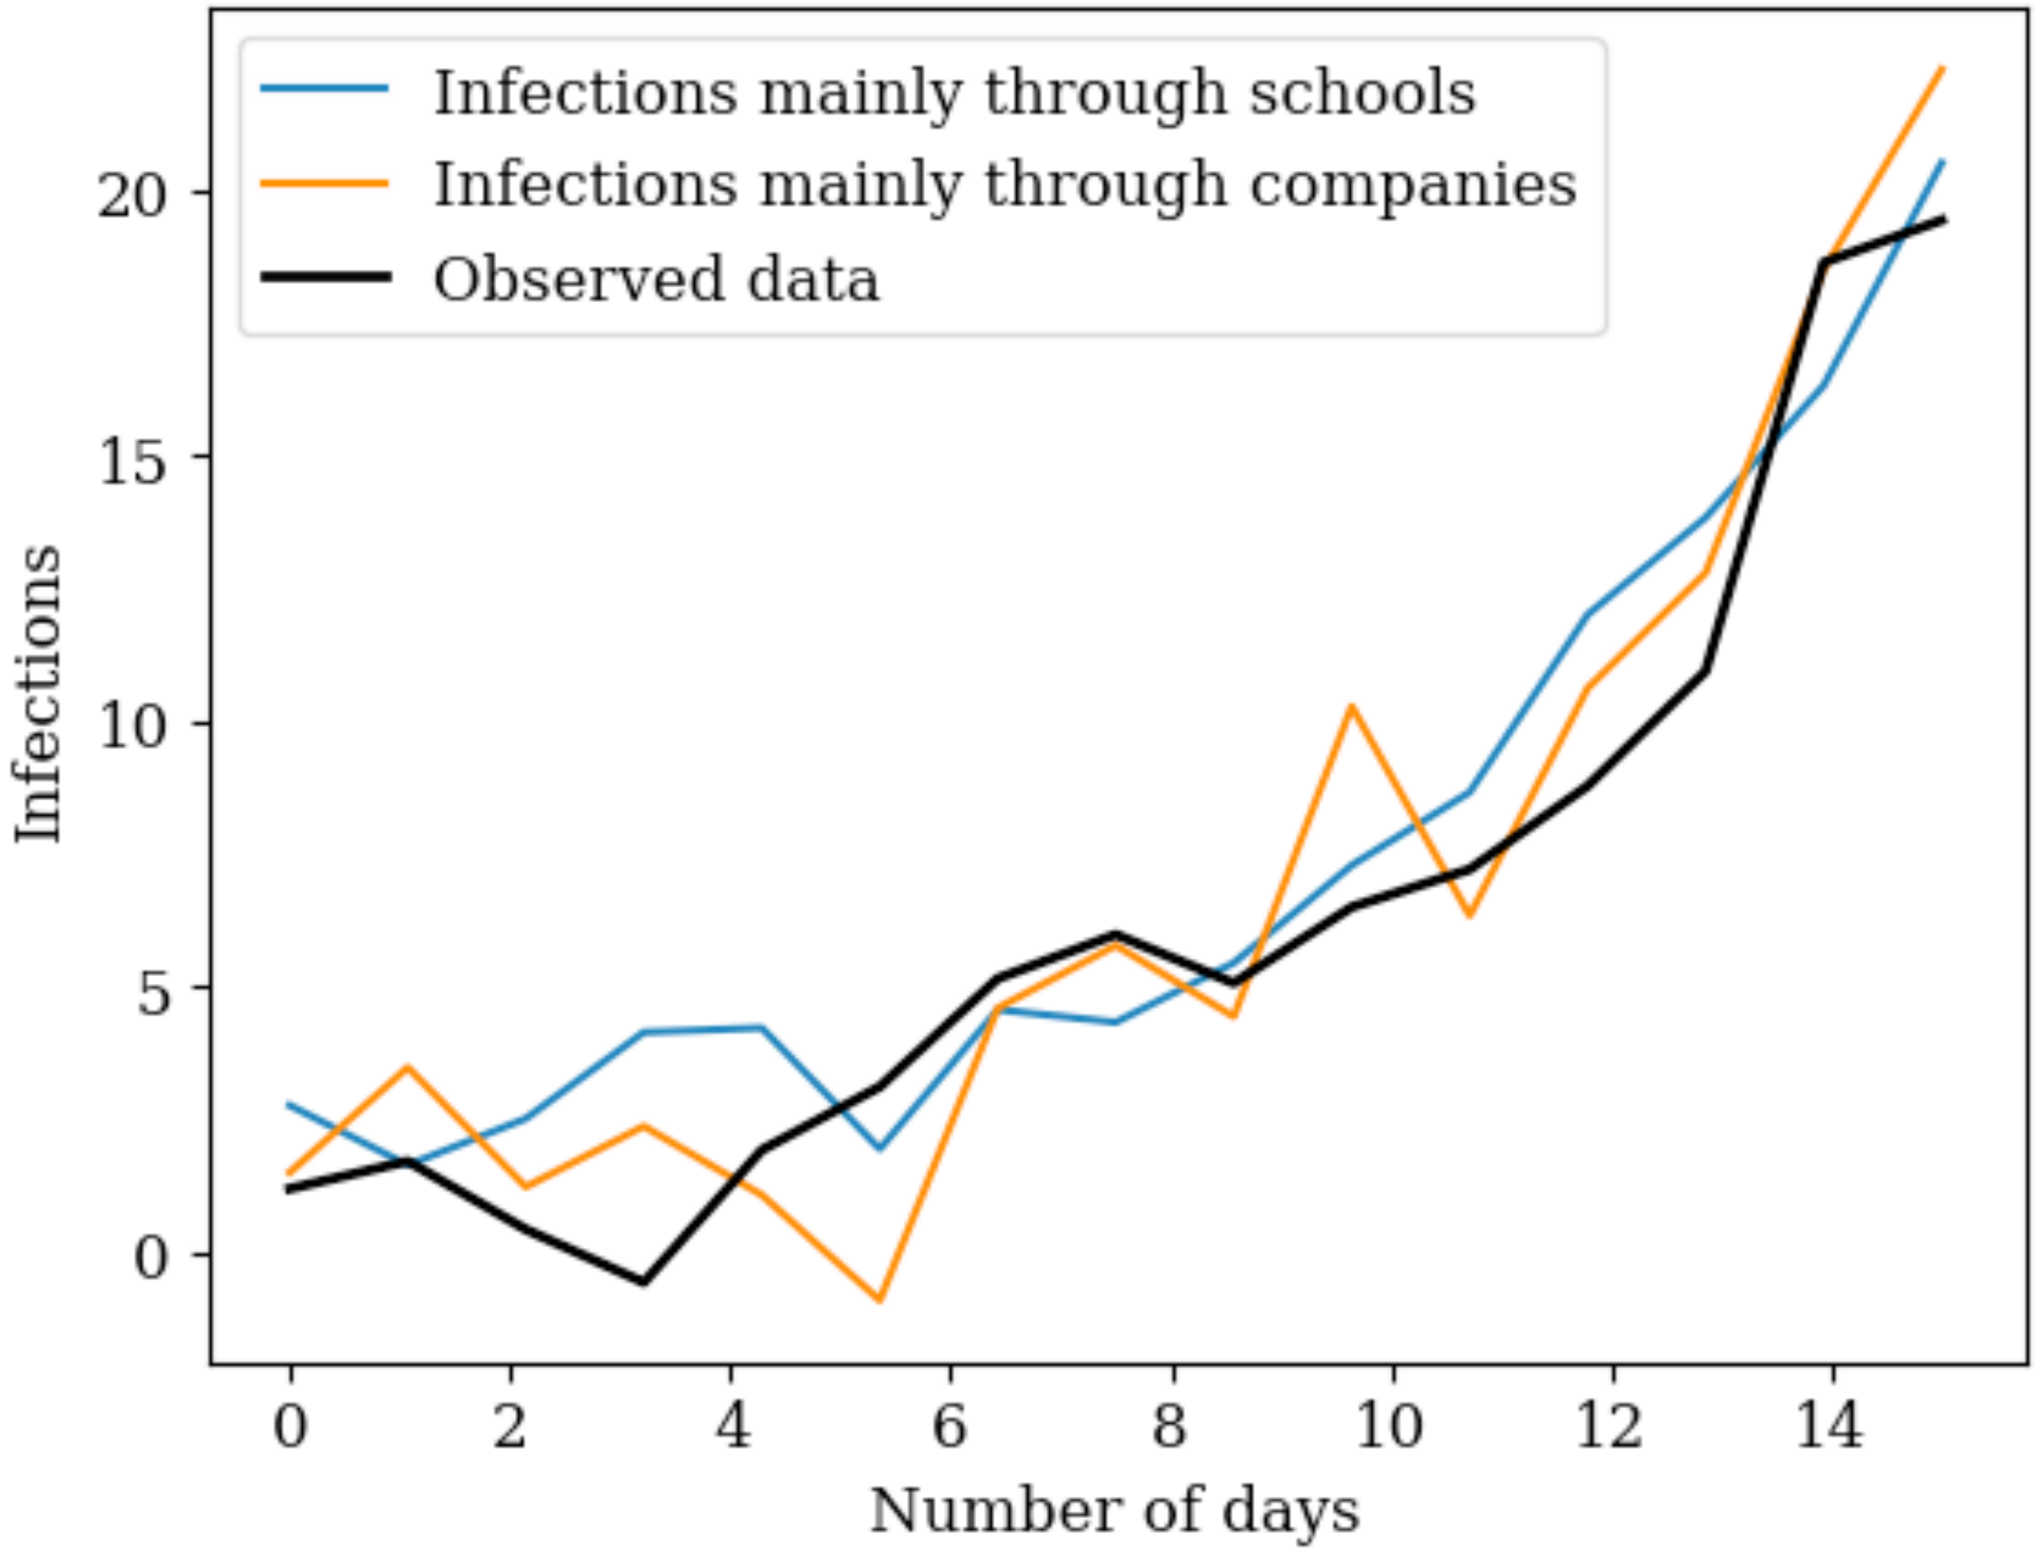

Calibration requirements

1. Uncertainty quantification (UQ)

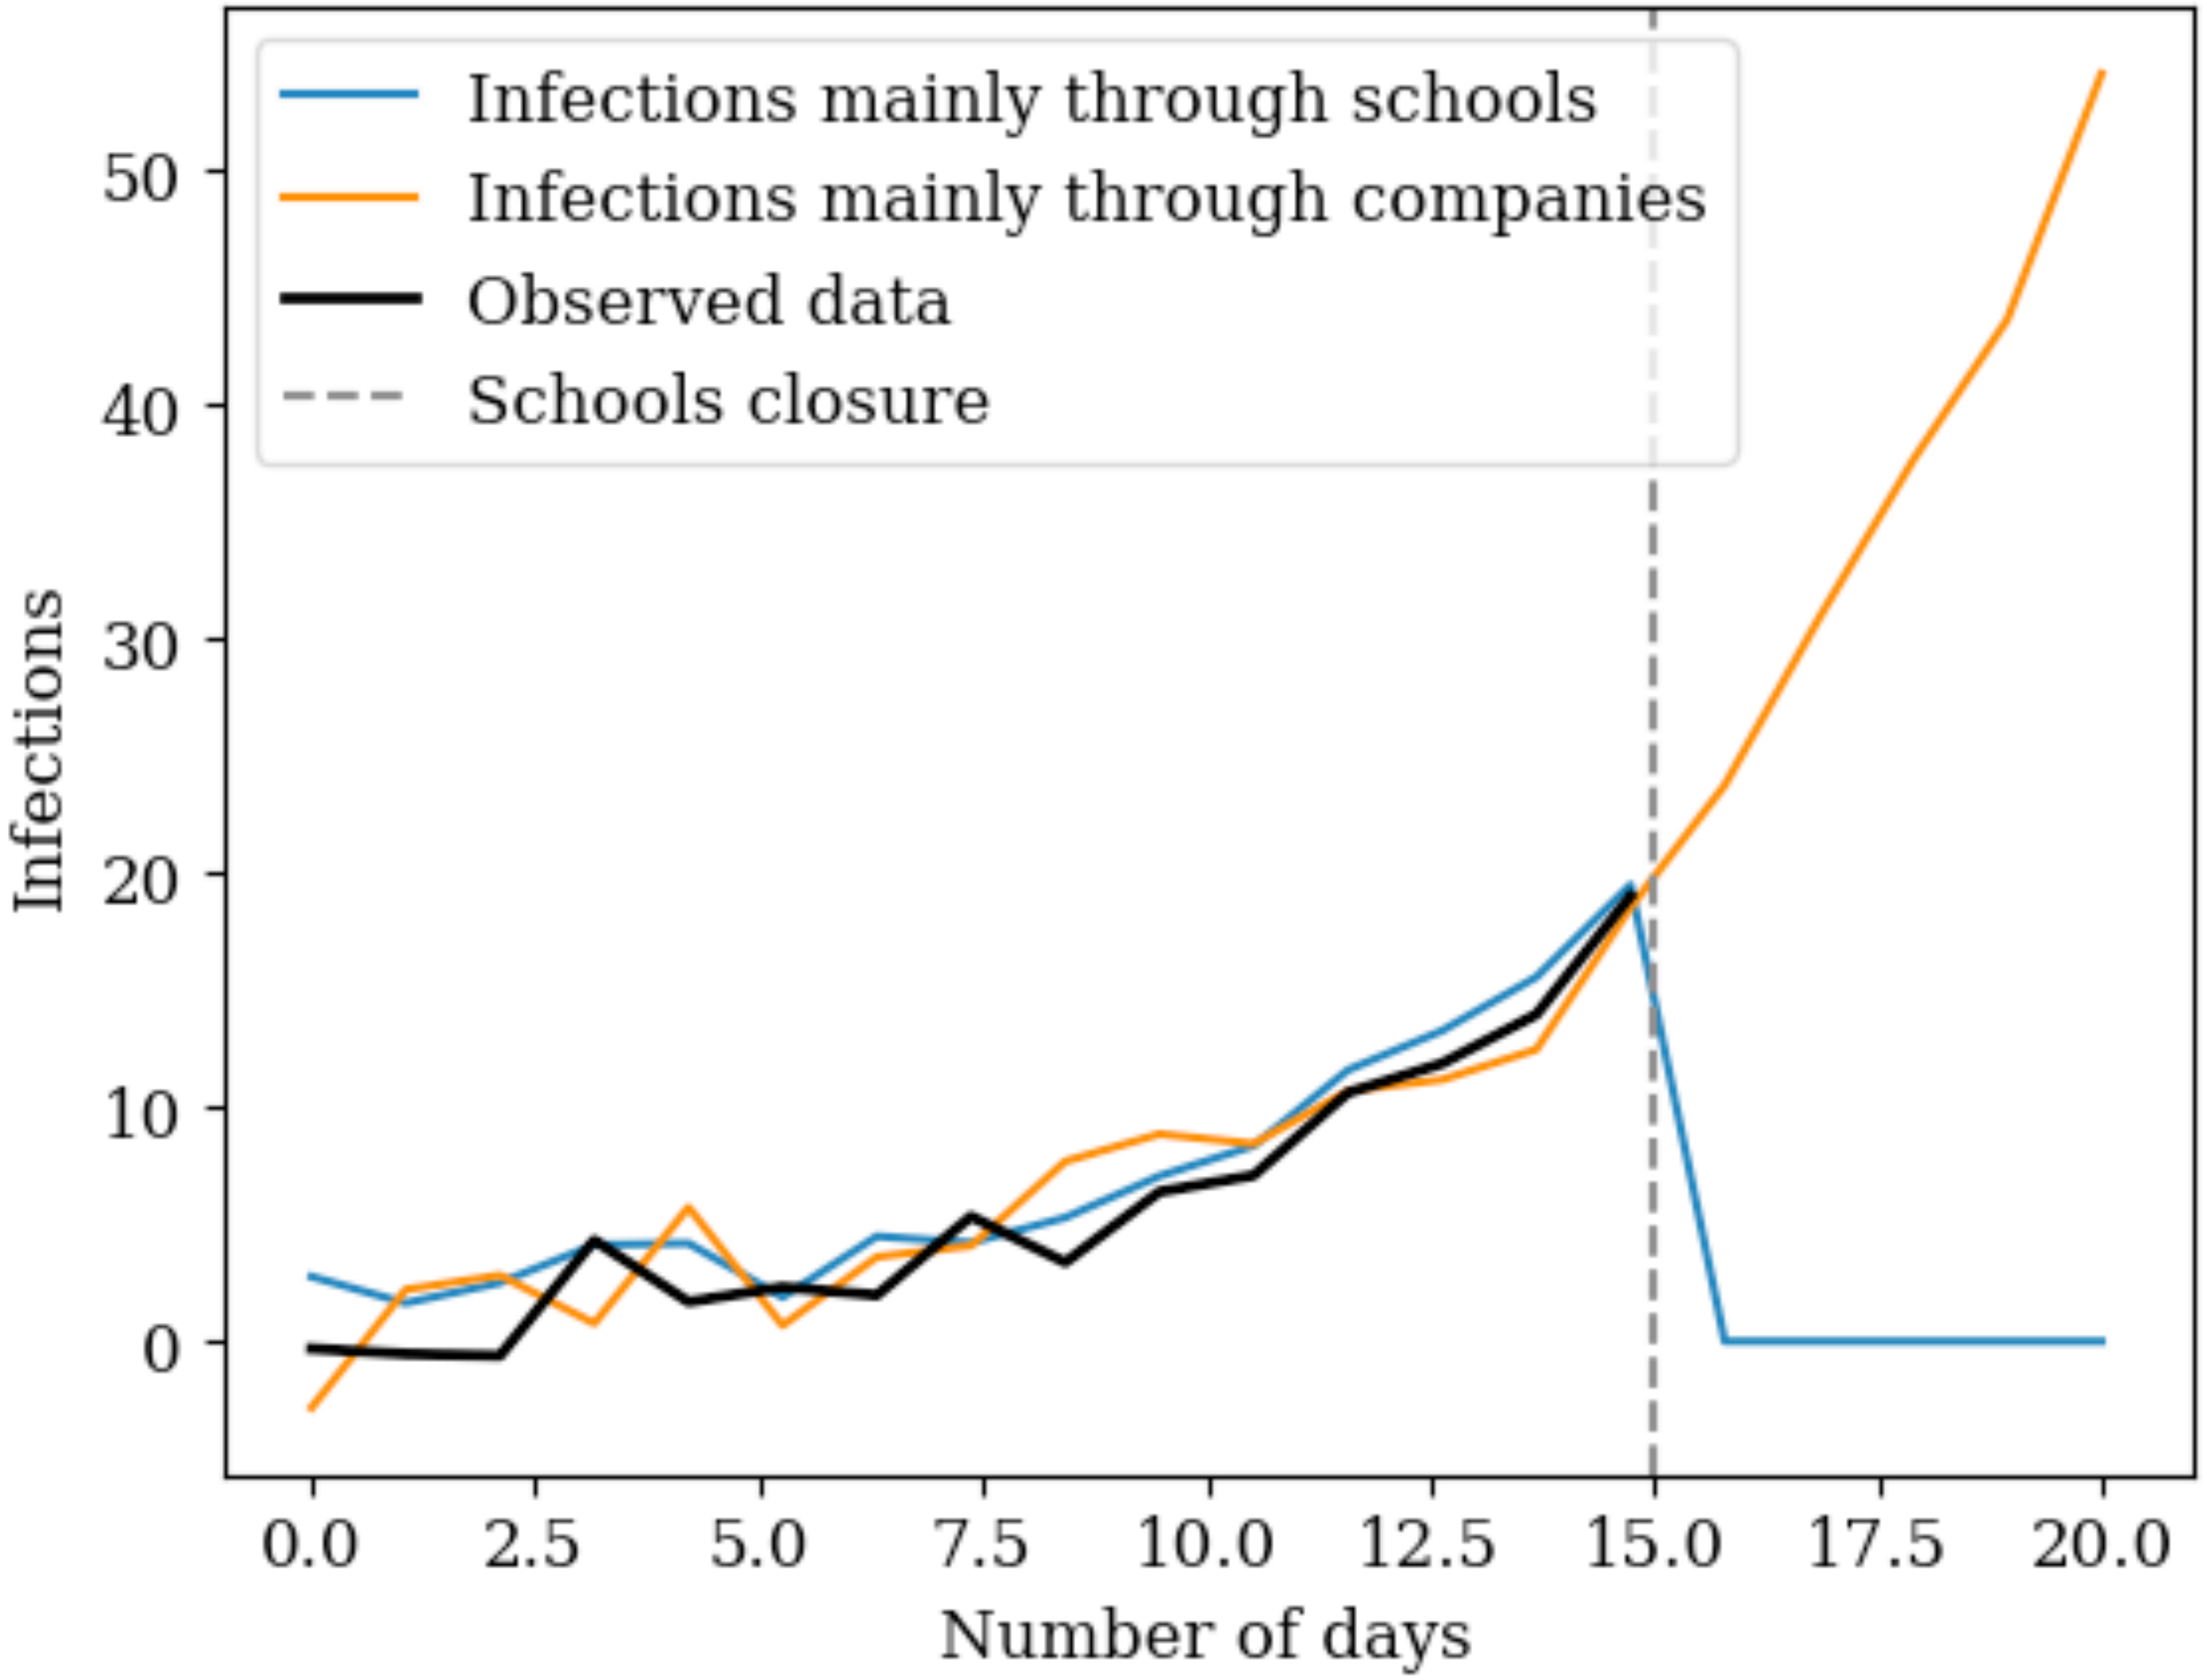

Crucial for policy analysis

Calibration requirements

2. Expert (prior) knowledge

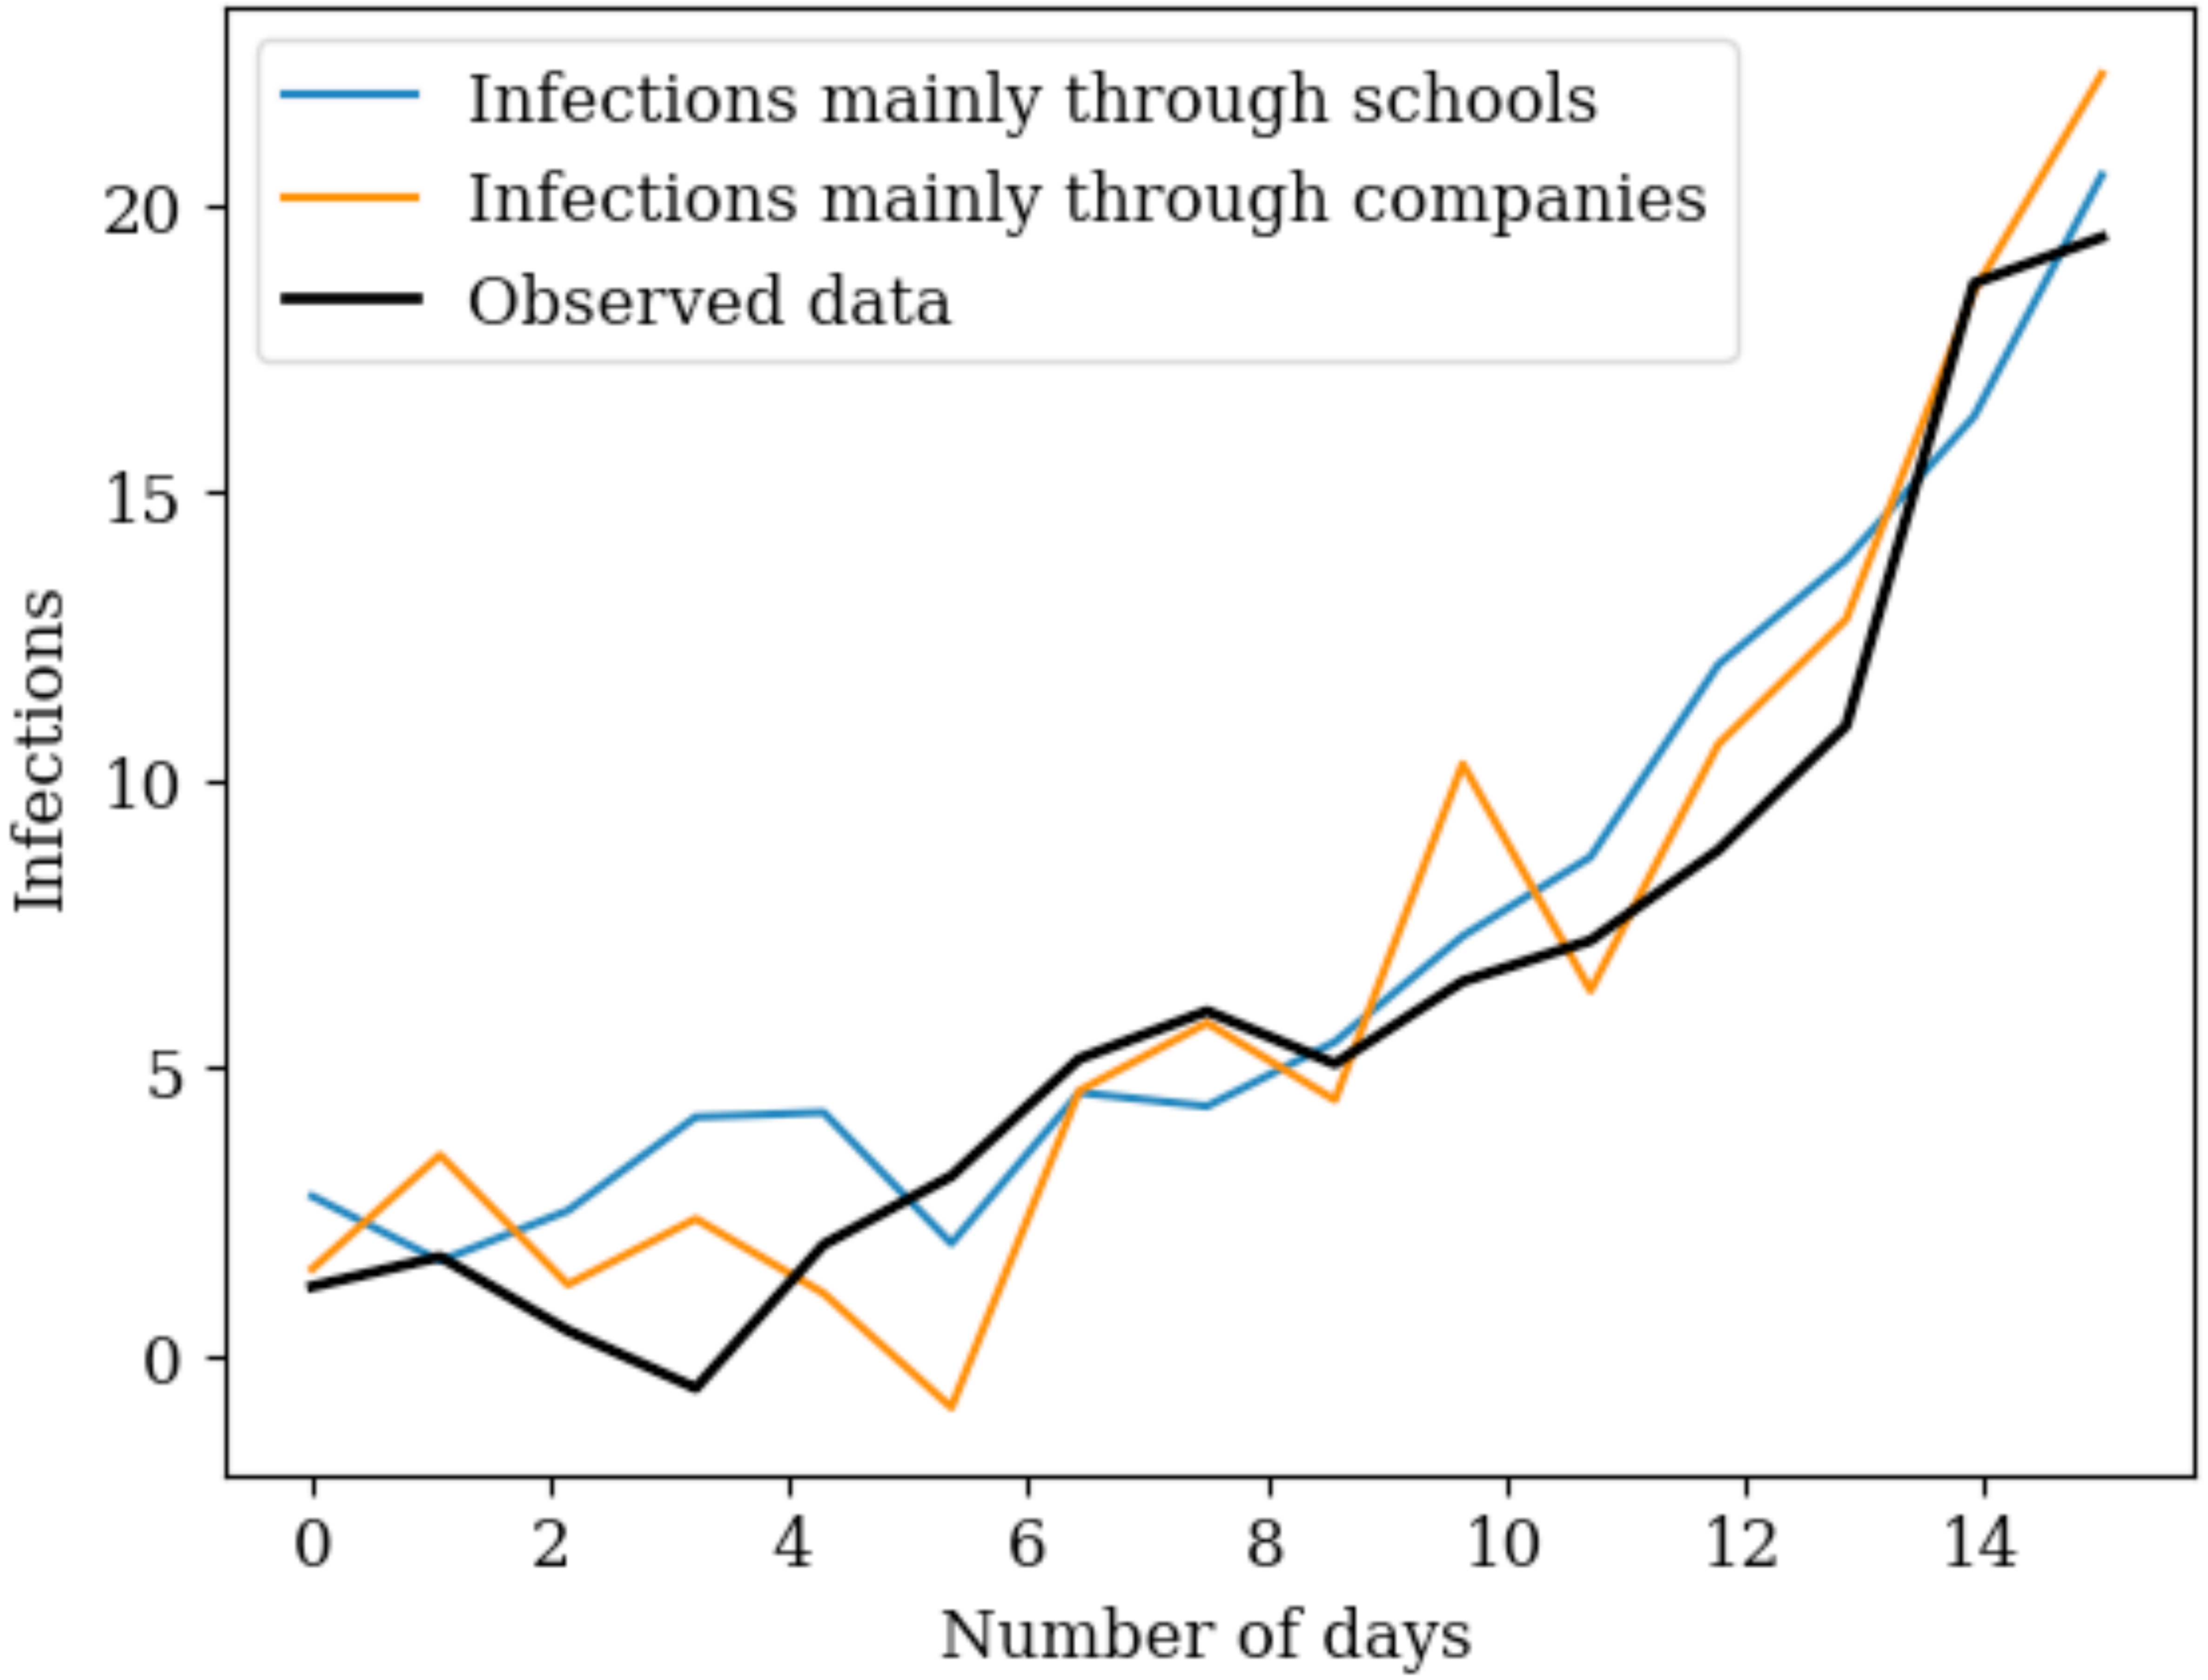

Need to include prior information in our calibration process

Bayesian inference

Allows to tackle both problems

$$\underbrace{p(\theta \mid x)}_{\text{posterior}} \propto \underbrace{p(x \mid \theta)}_{\text{likelihood}} \; \underbrace{p(\theta)}_{\text{prior}}$$

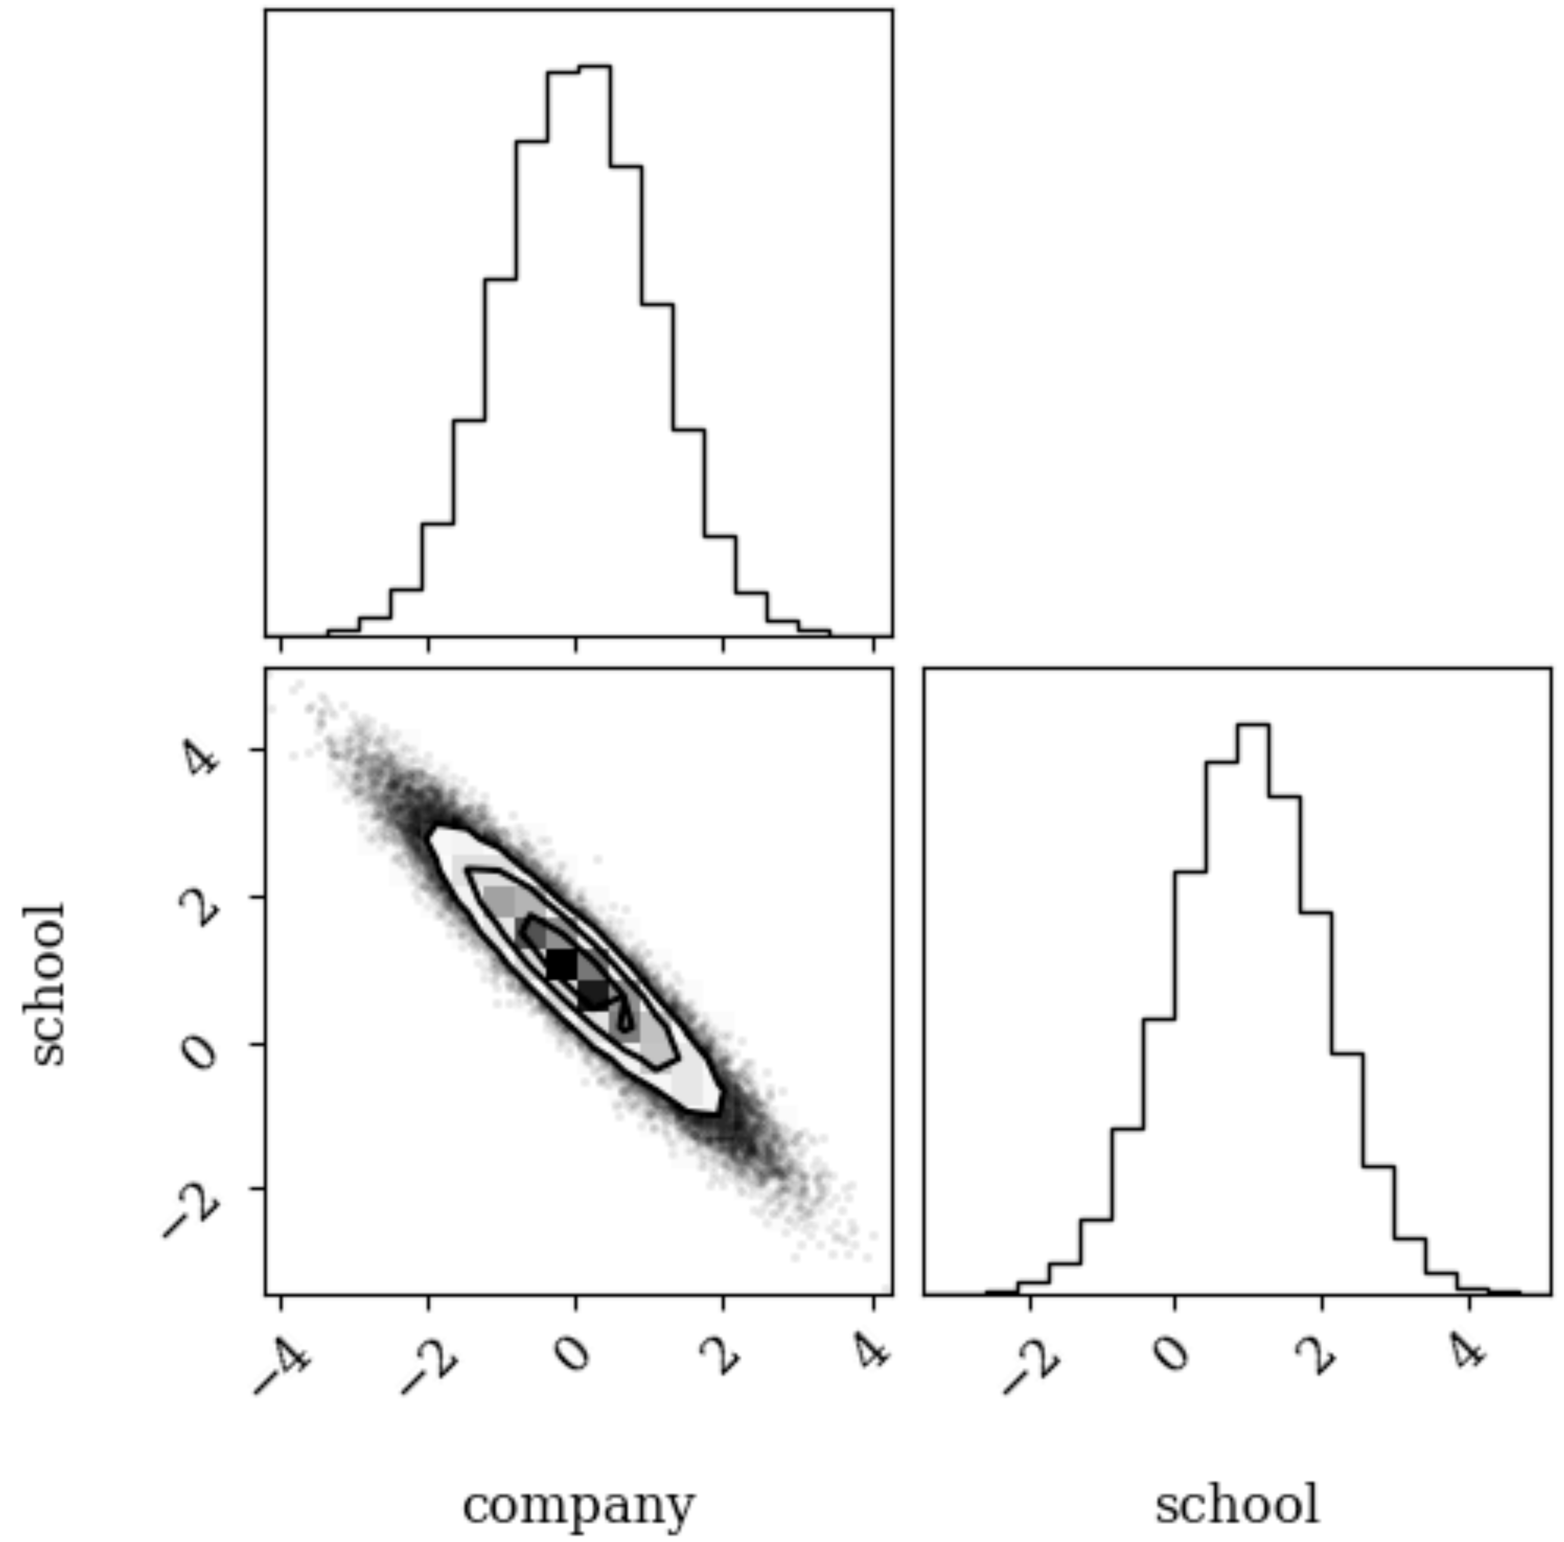

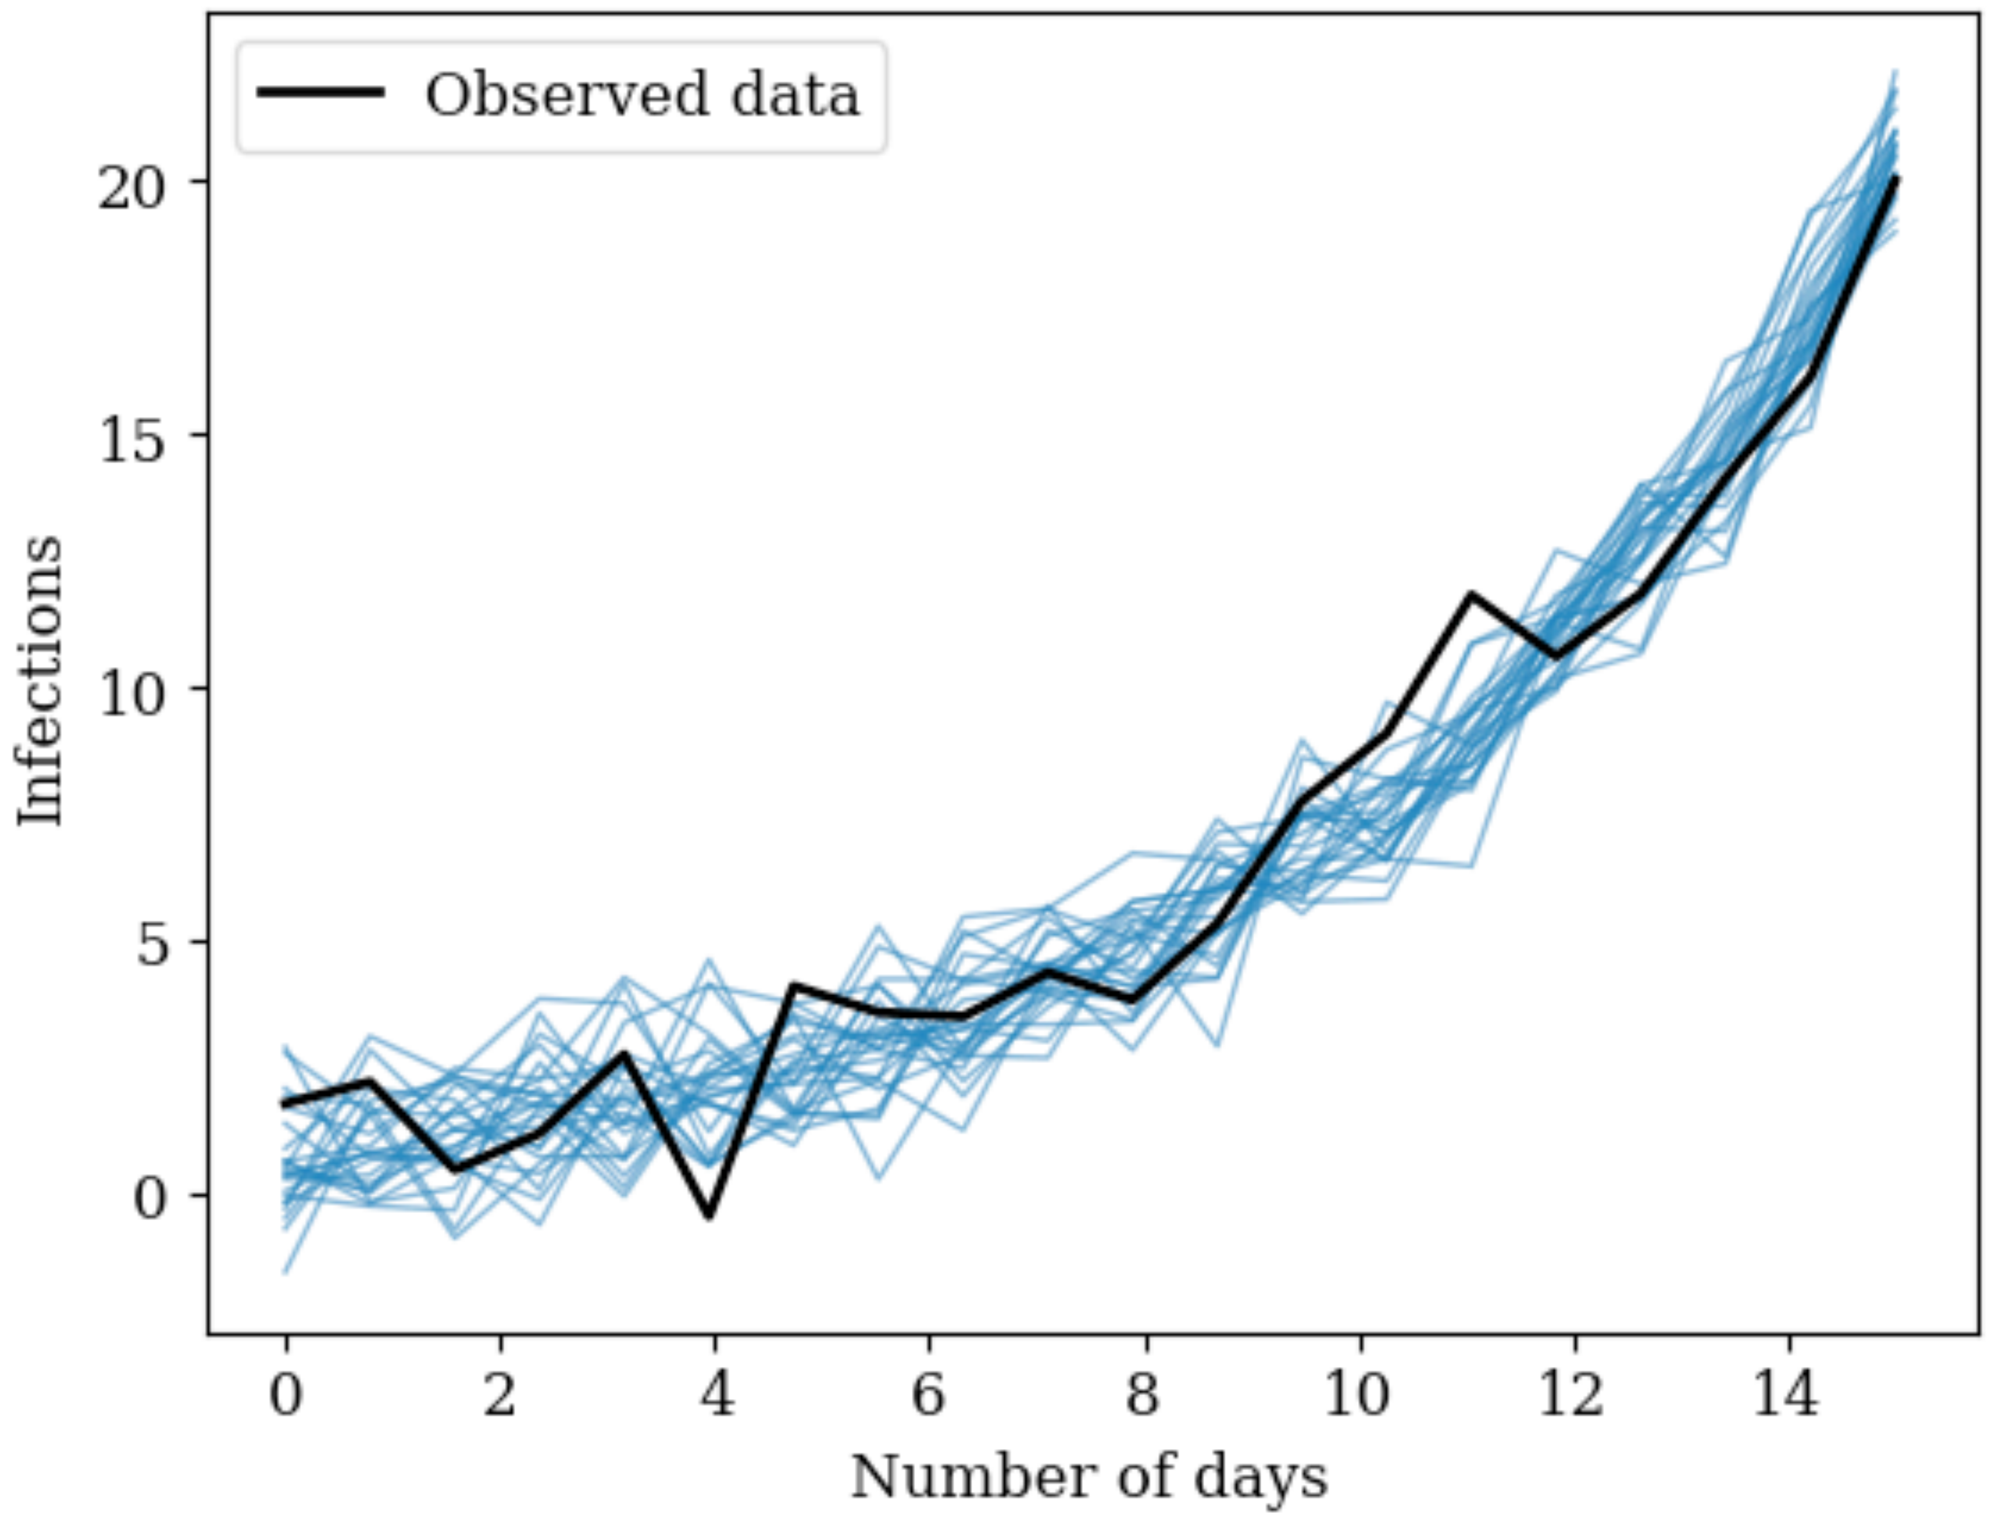

Likelihood p(x | θ) is intractable for ABMs

Proposed solutions include

Emulation

Approximate Bayesian Computation

Neural density ratio estimation

Variational Inference: Bayesian inference as an optimisation problem

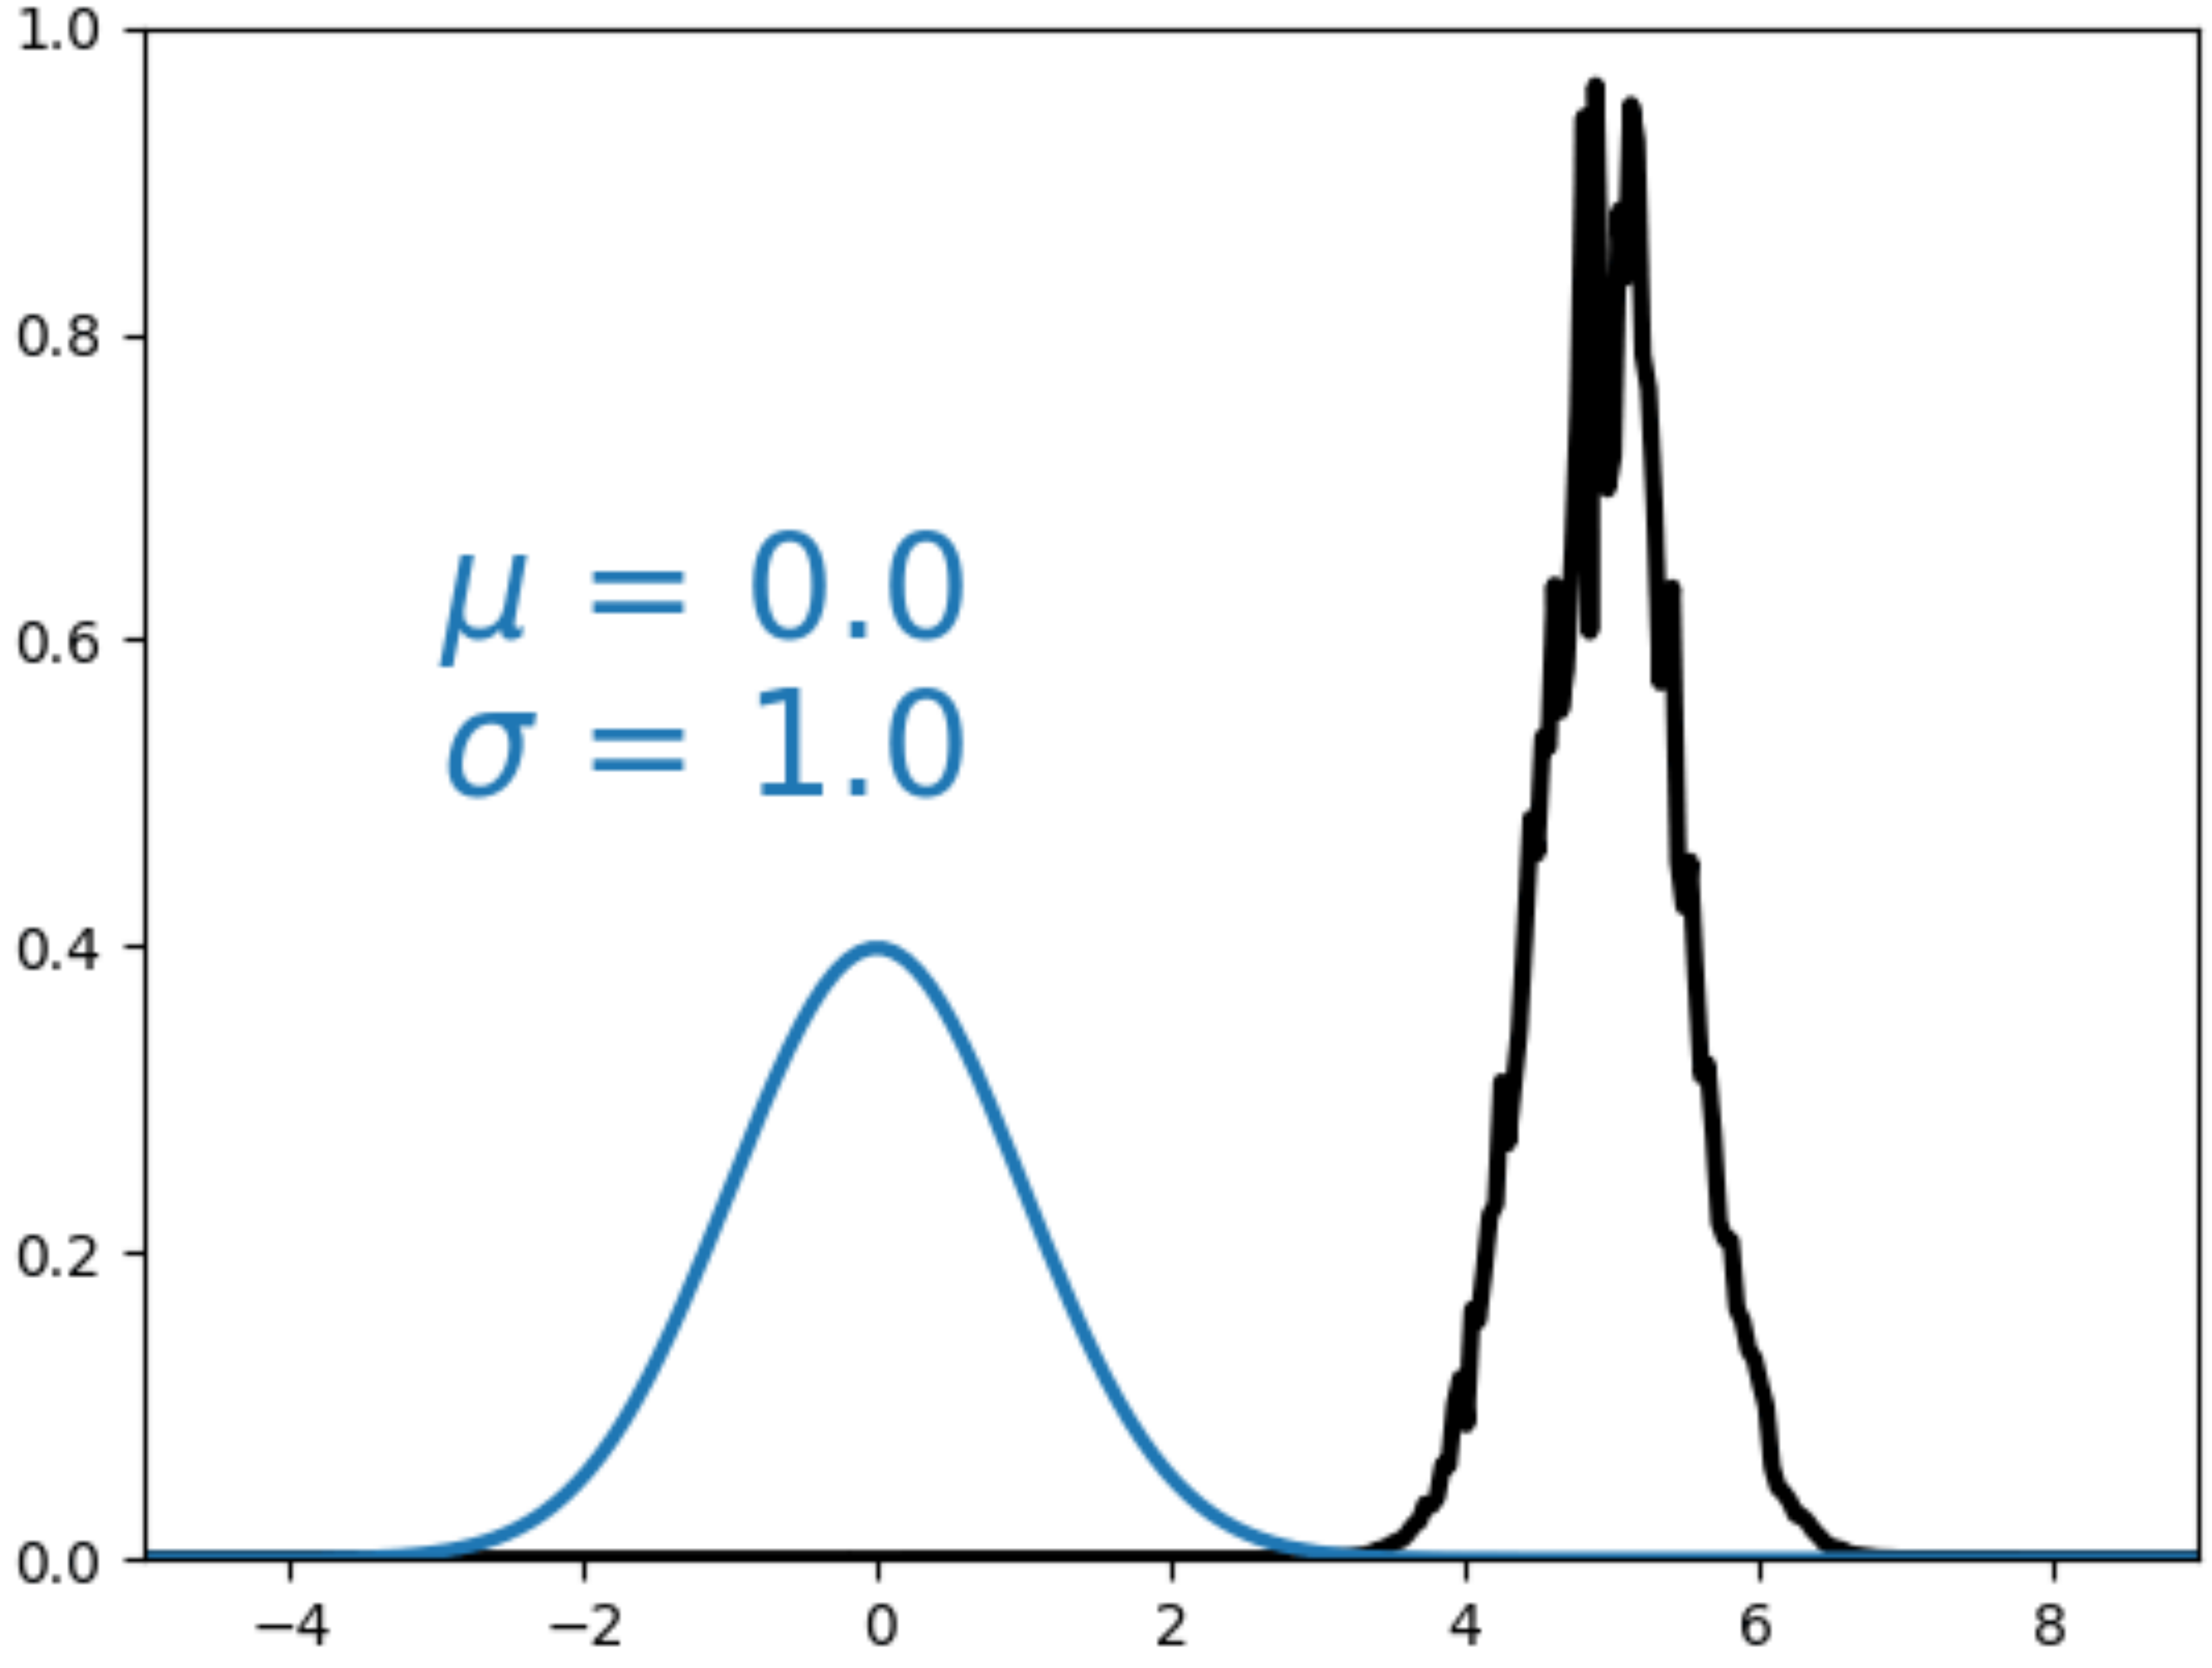

1. Assume posterior can be approximated by a family of distributions

2. Optimise for optimal parameters

Generalised Variational Inference

Knoblauch et al., (2022)

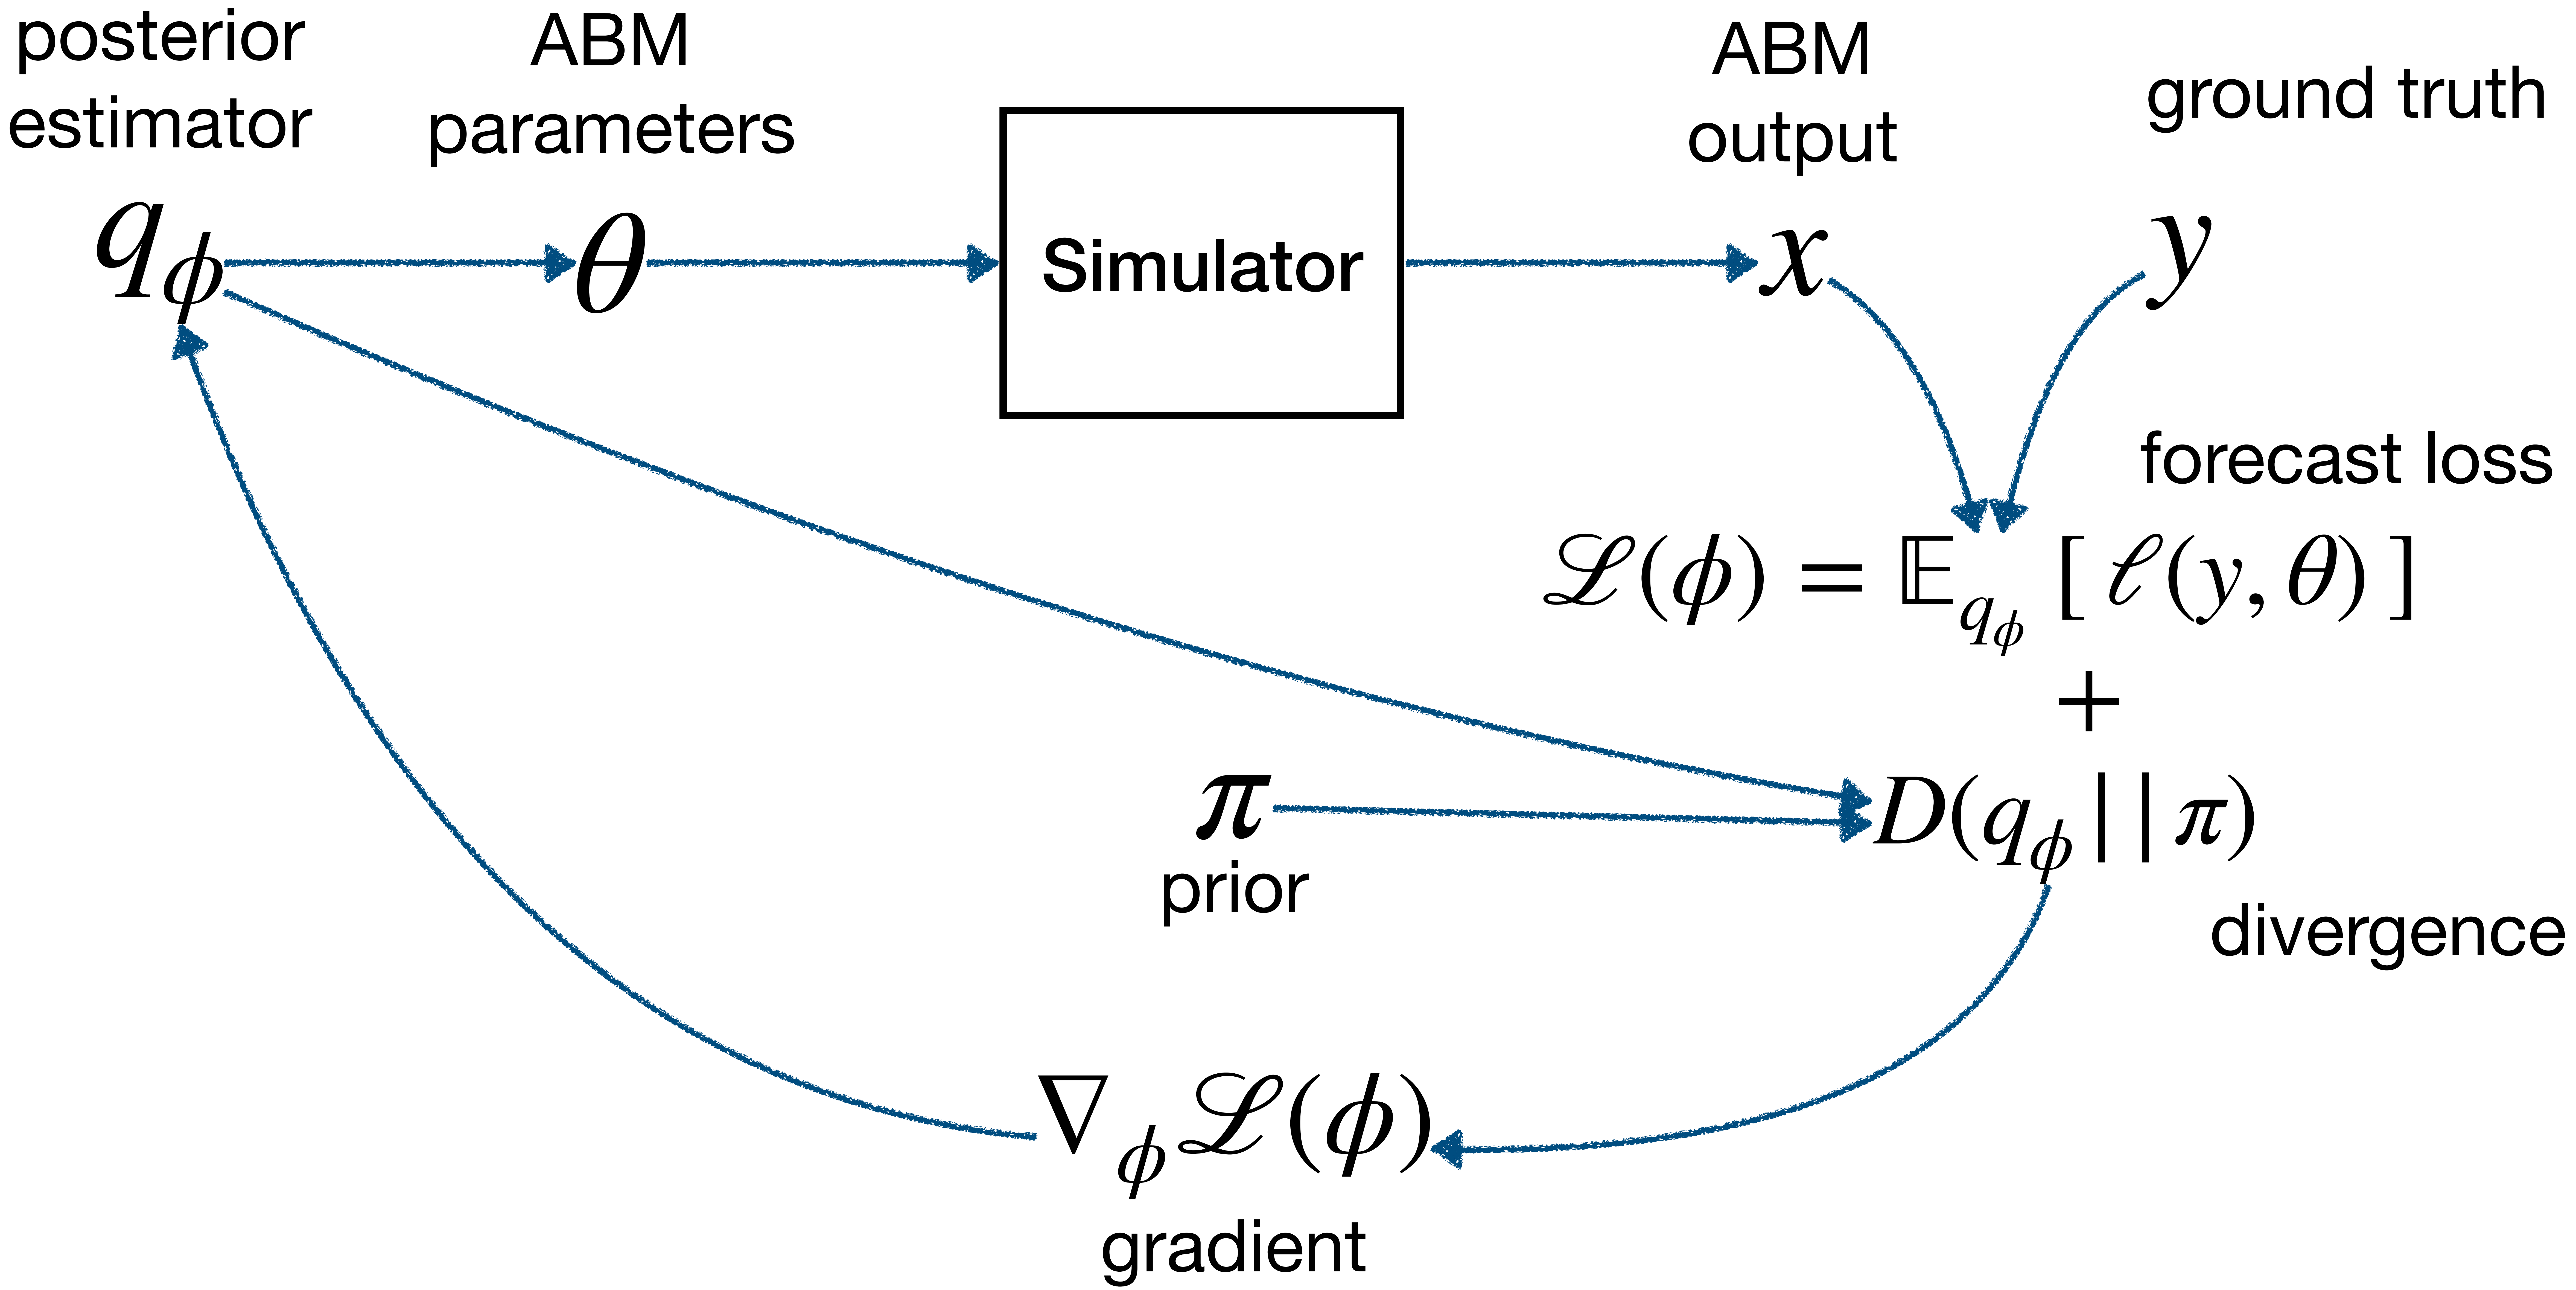

Gradients: path-wise vs score

- Gradient-assisted calibration algorithms need

$$\underbrace{\nabla_\phi \mathbb{E}_{p(\theta\,;\,\phi)}}_{\text{posterior estimator}} \underbrace{\left[f(\theta)\right]}_{\text{simulator}}$$

- Two ways of obtaining the gradient:

1. Differentiating the measure (**score-based gradient**)
2. Differentiating the simulator (**path-wise gradient**)

Typically **path-wise gradient has (much) lower variance** (see Mohamed (2019))

Differentiable simulators

- Leverage Automatic Differentiation to build simulators

- Use “reparameterisation” techniques to differentiate through randomness.

$$x \sim \mathcal{N}(\mu, \sigma) \quad \iff \quad x = \mu + \sigma\, r \quad r \sim \mathcal{N}(0,1)$$

$$\frac{\mathrm{d}x}{\mathrm{d}\mu} = 1 \quad \frac{\mathrm{d}x}{\mathrm{d}\sigma} = r$$

Differentiable ABMs

The problem of discrete randomness

- Discrete sampling + flow control = no differentiability?
  - Gumbel-Softmax

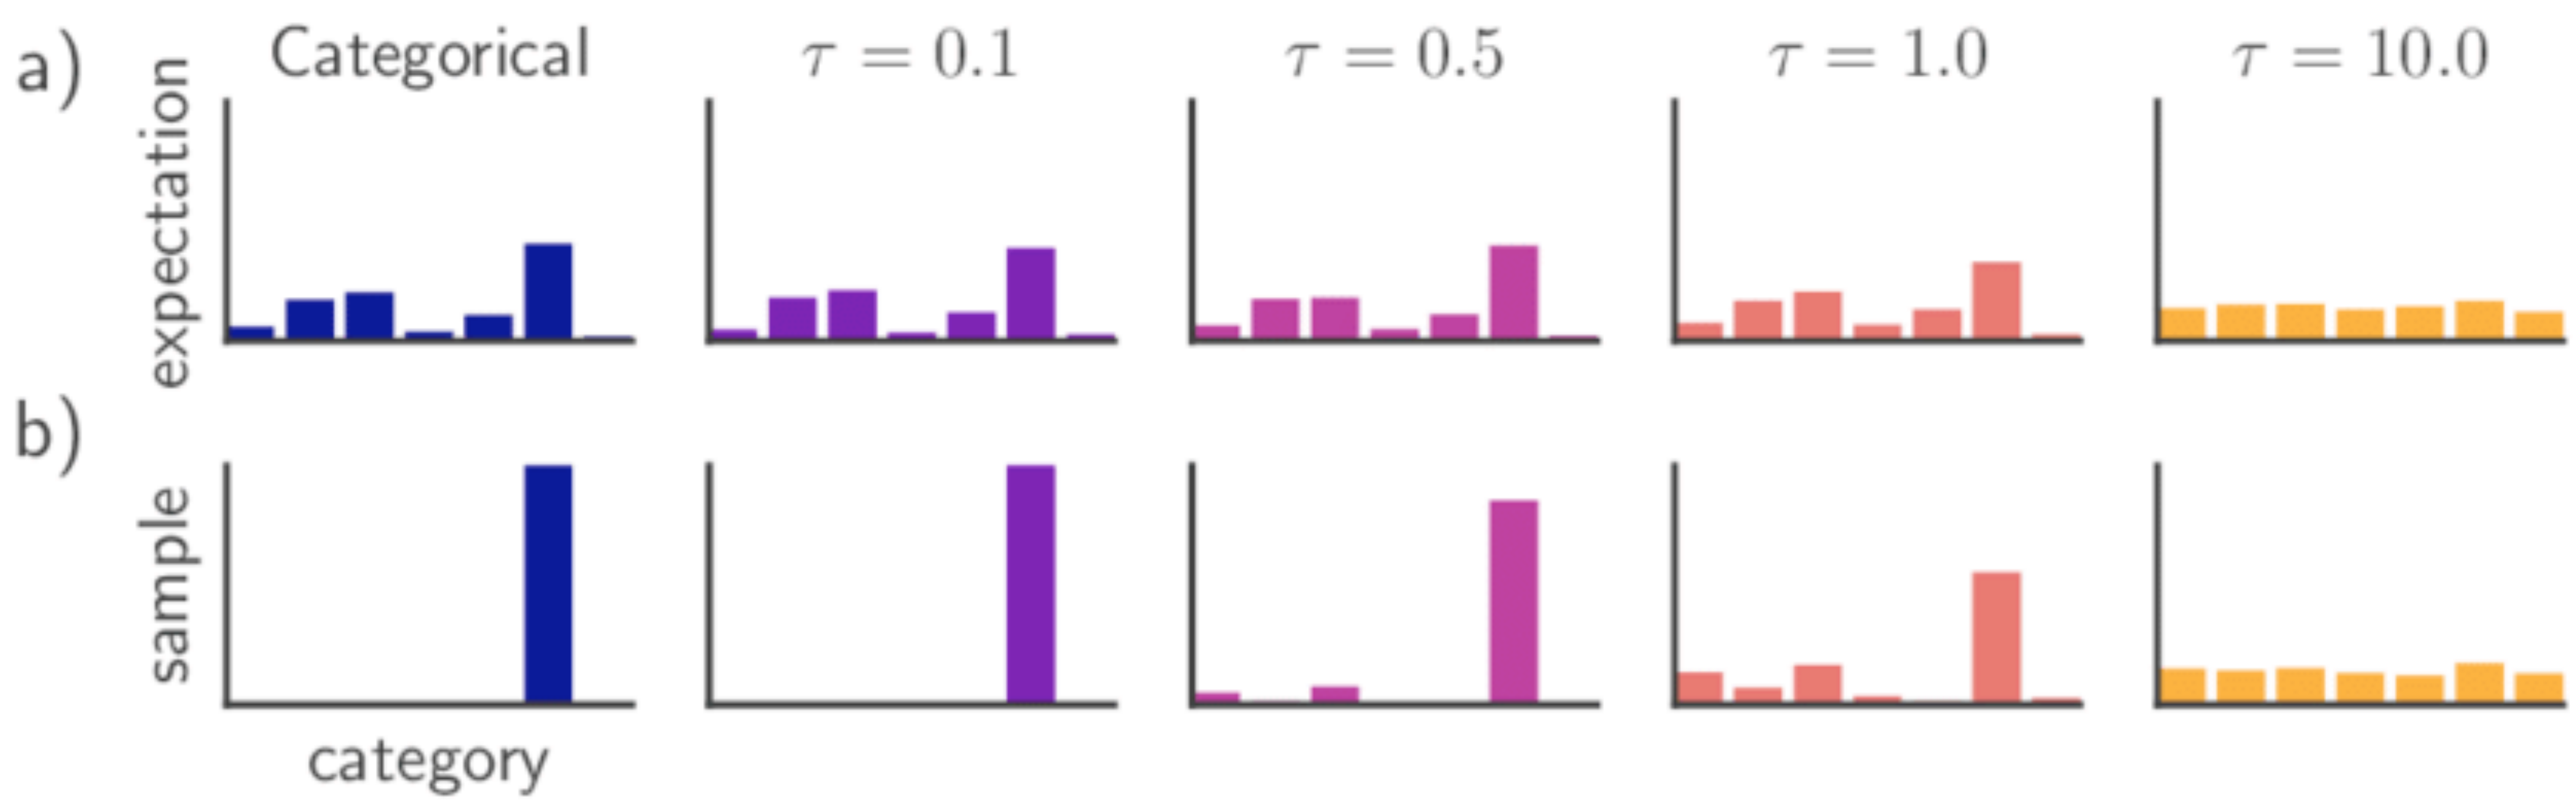

Jang et al. (2016)

Differentiable Agent-Based Epidemiology

Chopra et al. (2023), Quera-Bofarull et al. (2023)

JUNE model
8 M agents (London)

	Simulation
JUNE	50 hours
GradABM-June (CPU)	5 minutes
GradABM-June (GPU)	5 seconds

x40,000 speed-up !

Bayesian Inference for Differentiable Simulators (BIRDS)

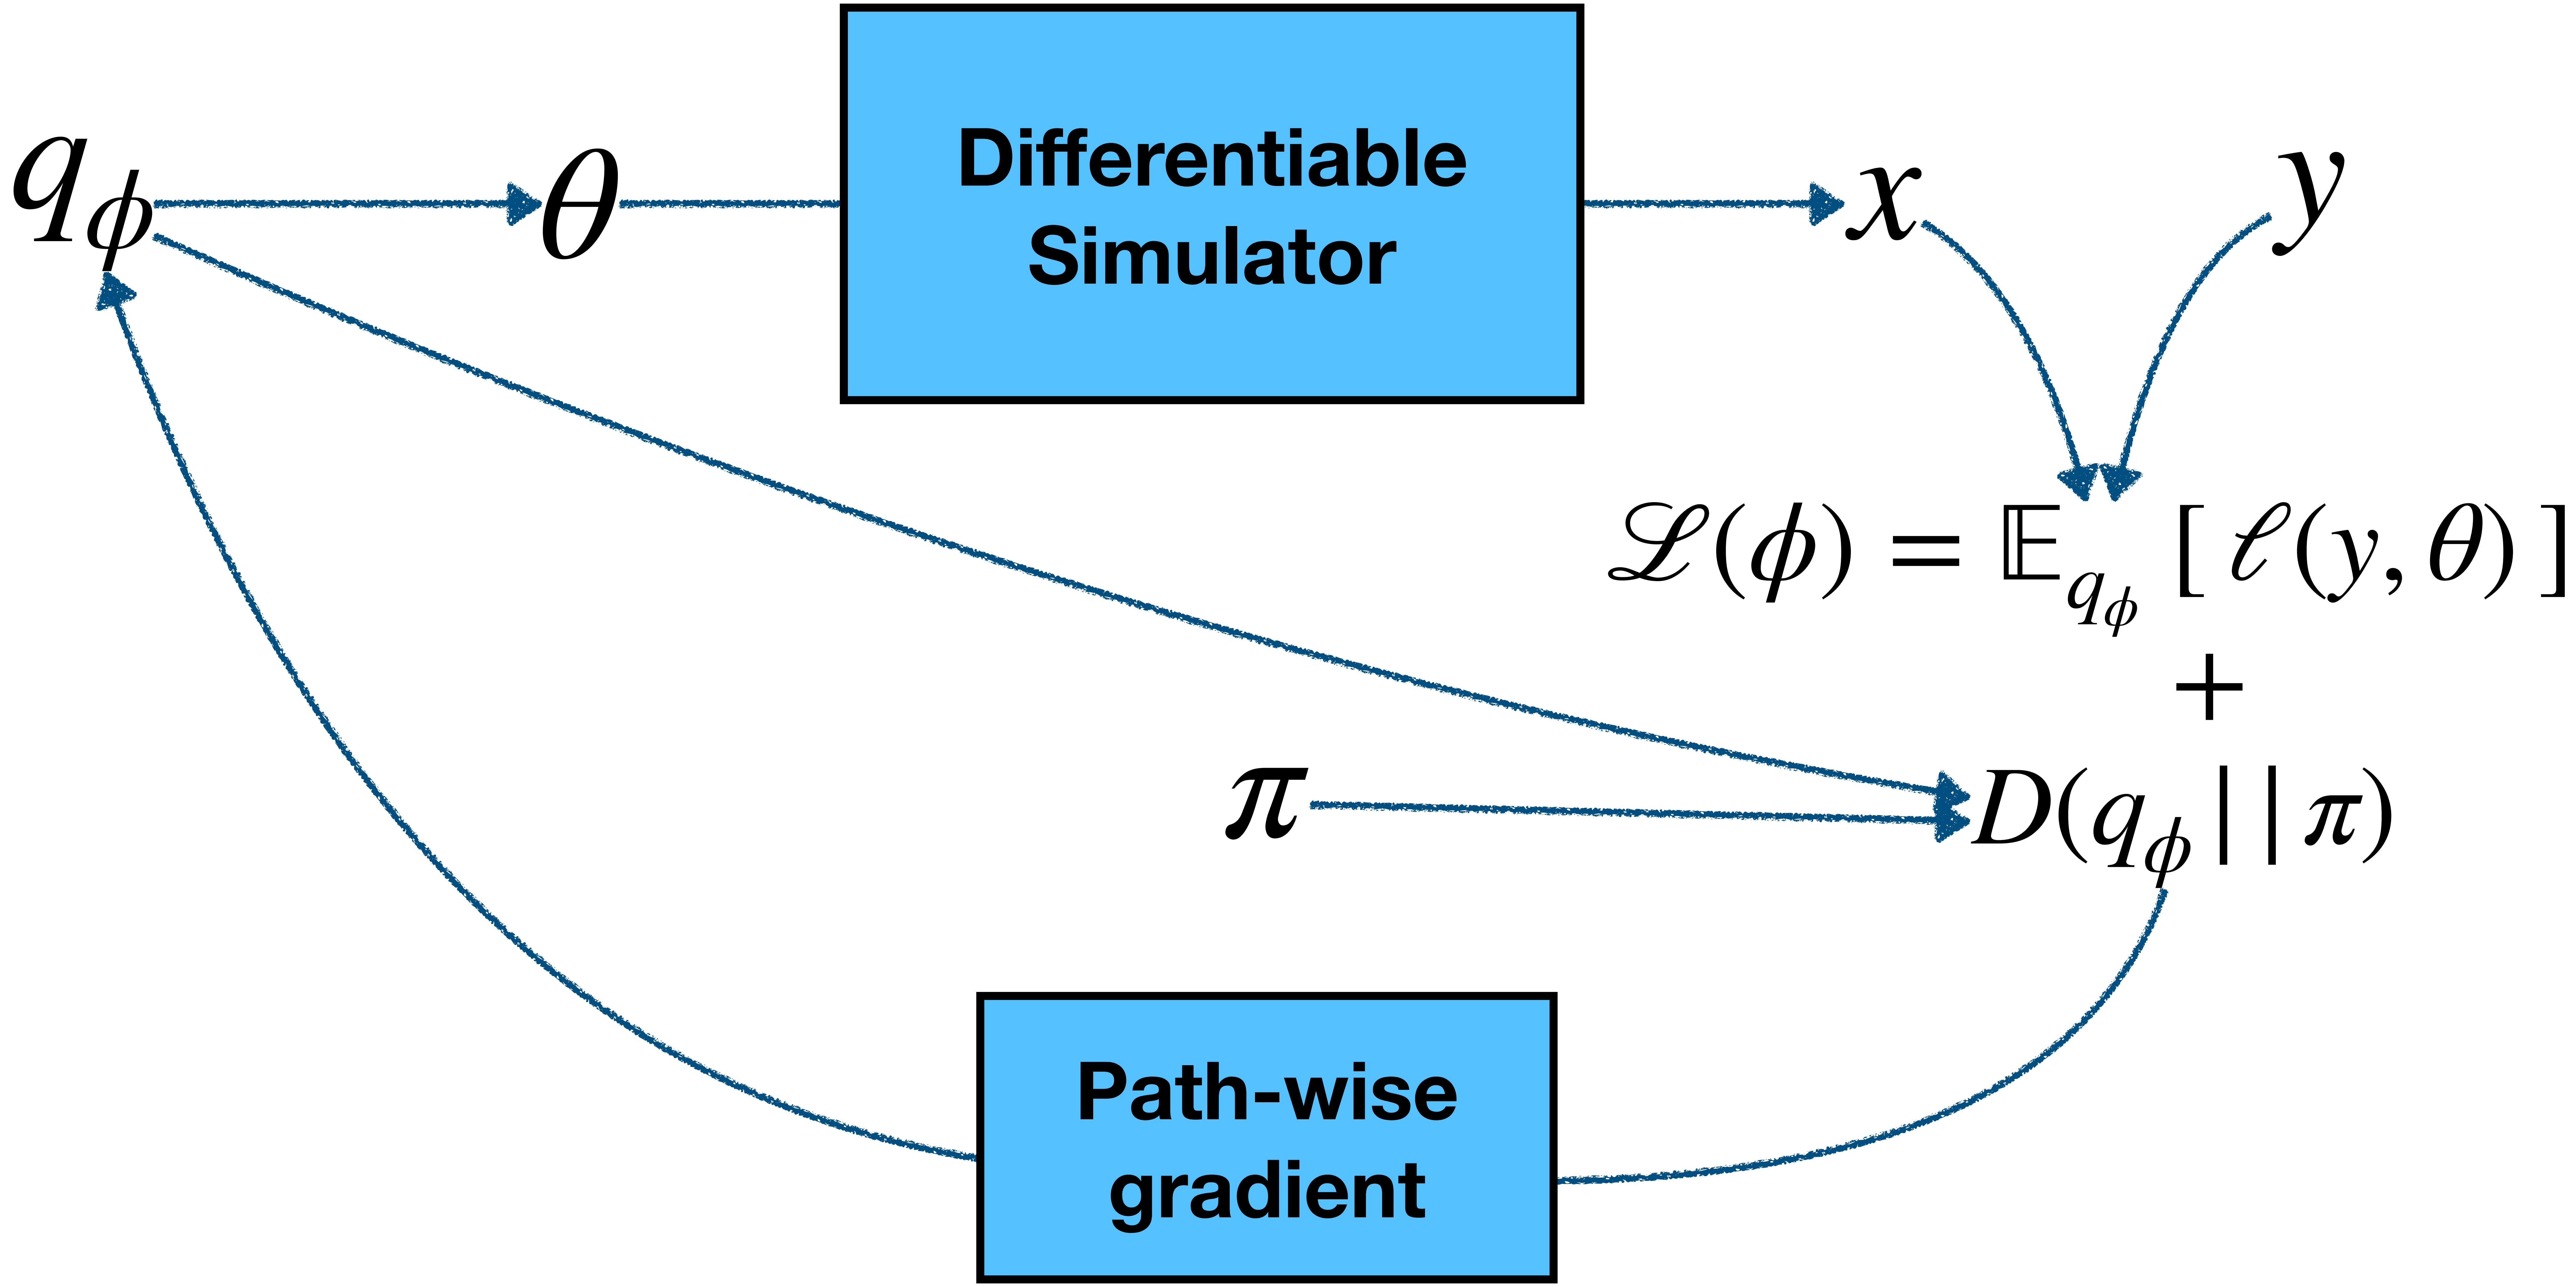

Normalizing Flows

What do we choose for q?

Image credit: Lilian Weng

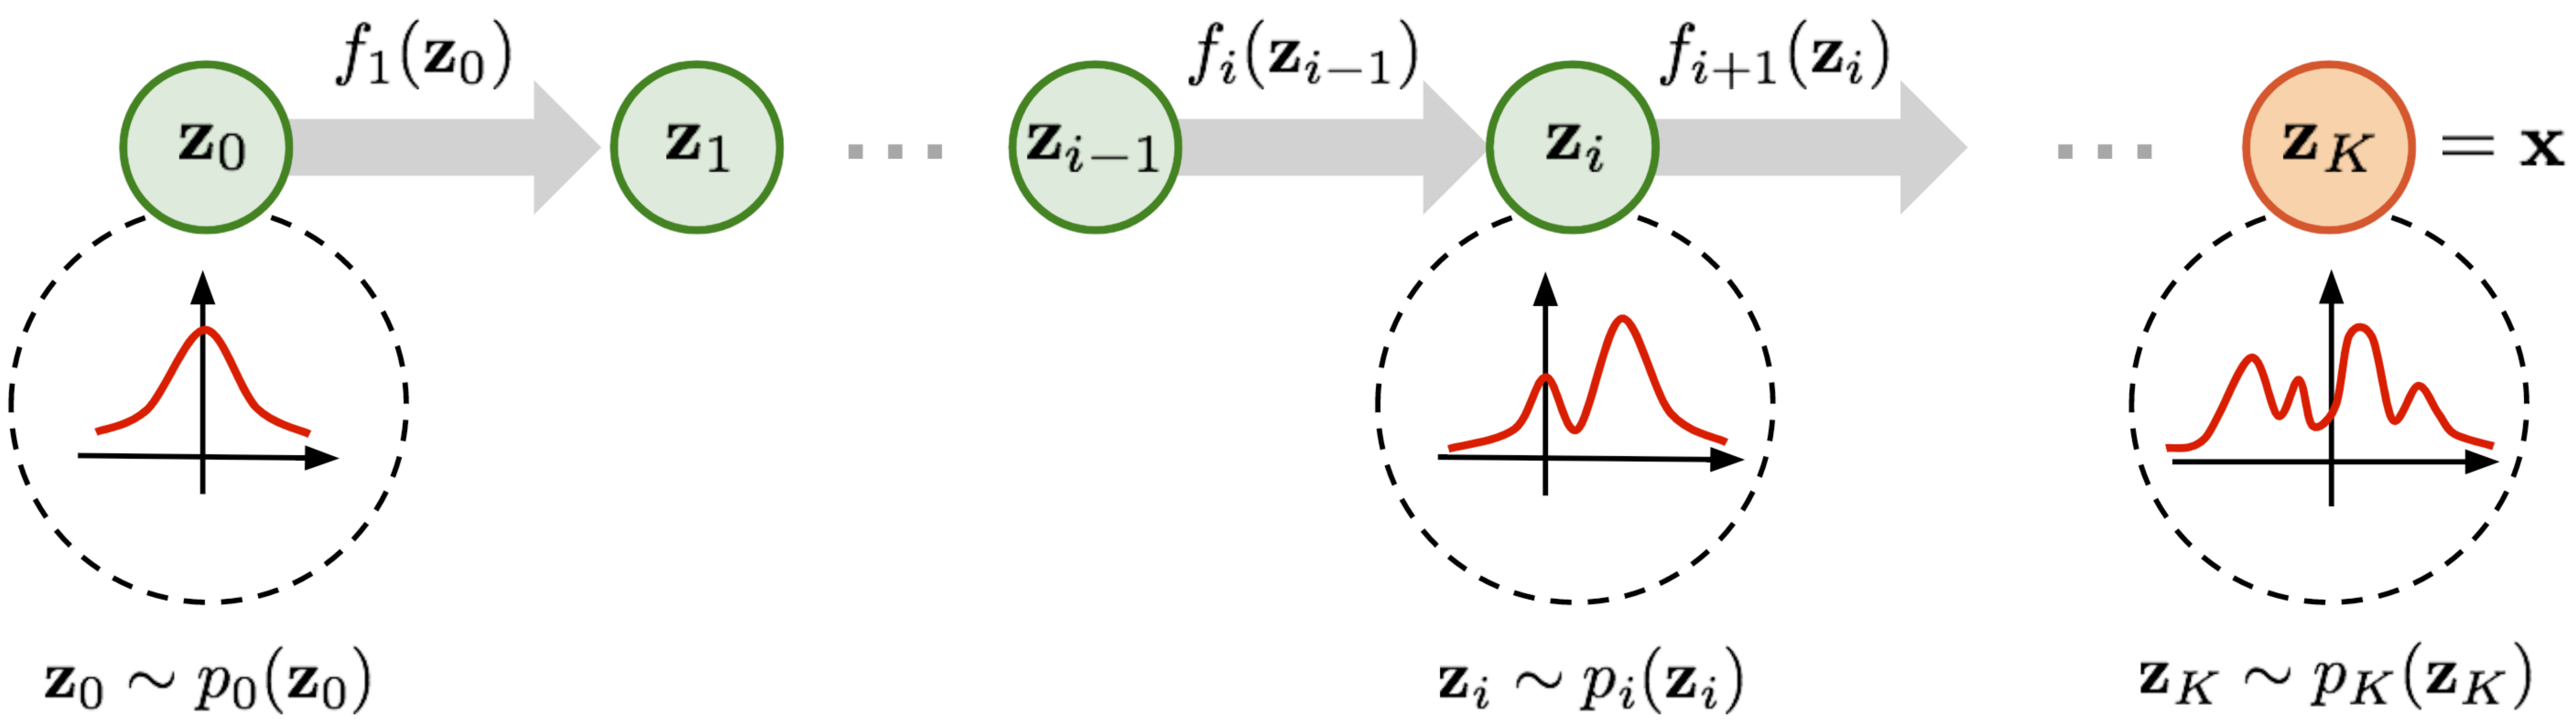

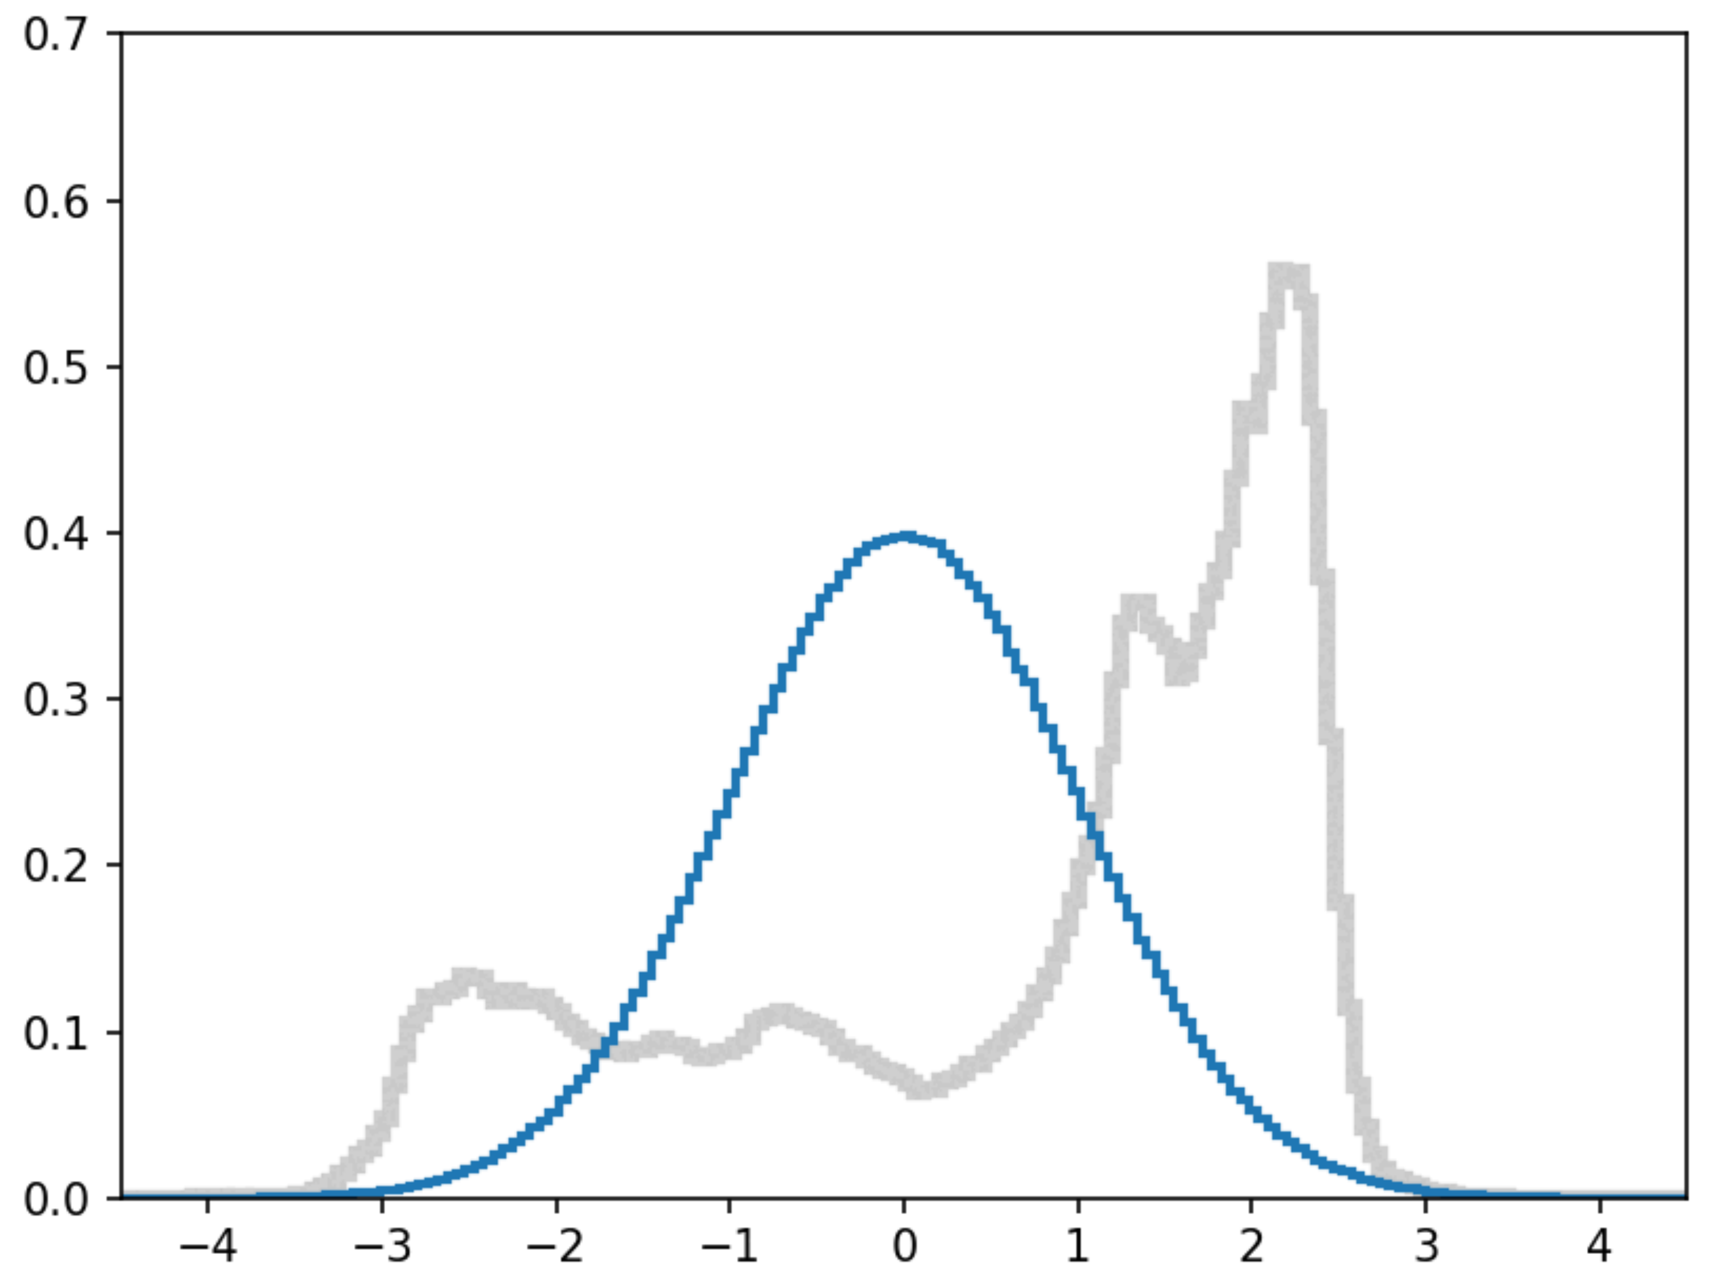

Bayesian Inference for Differentiable Simulators (BIRDS)

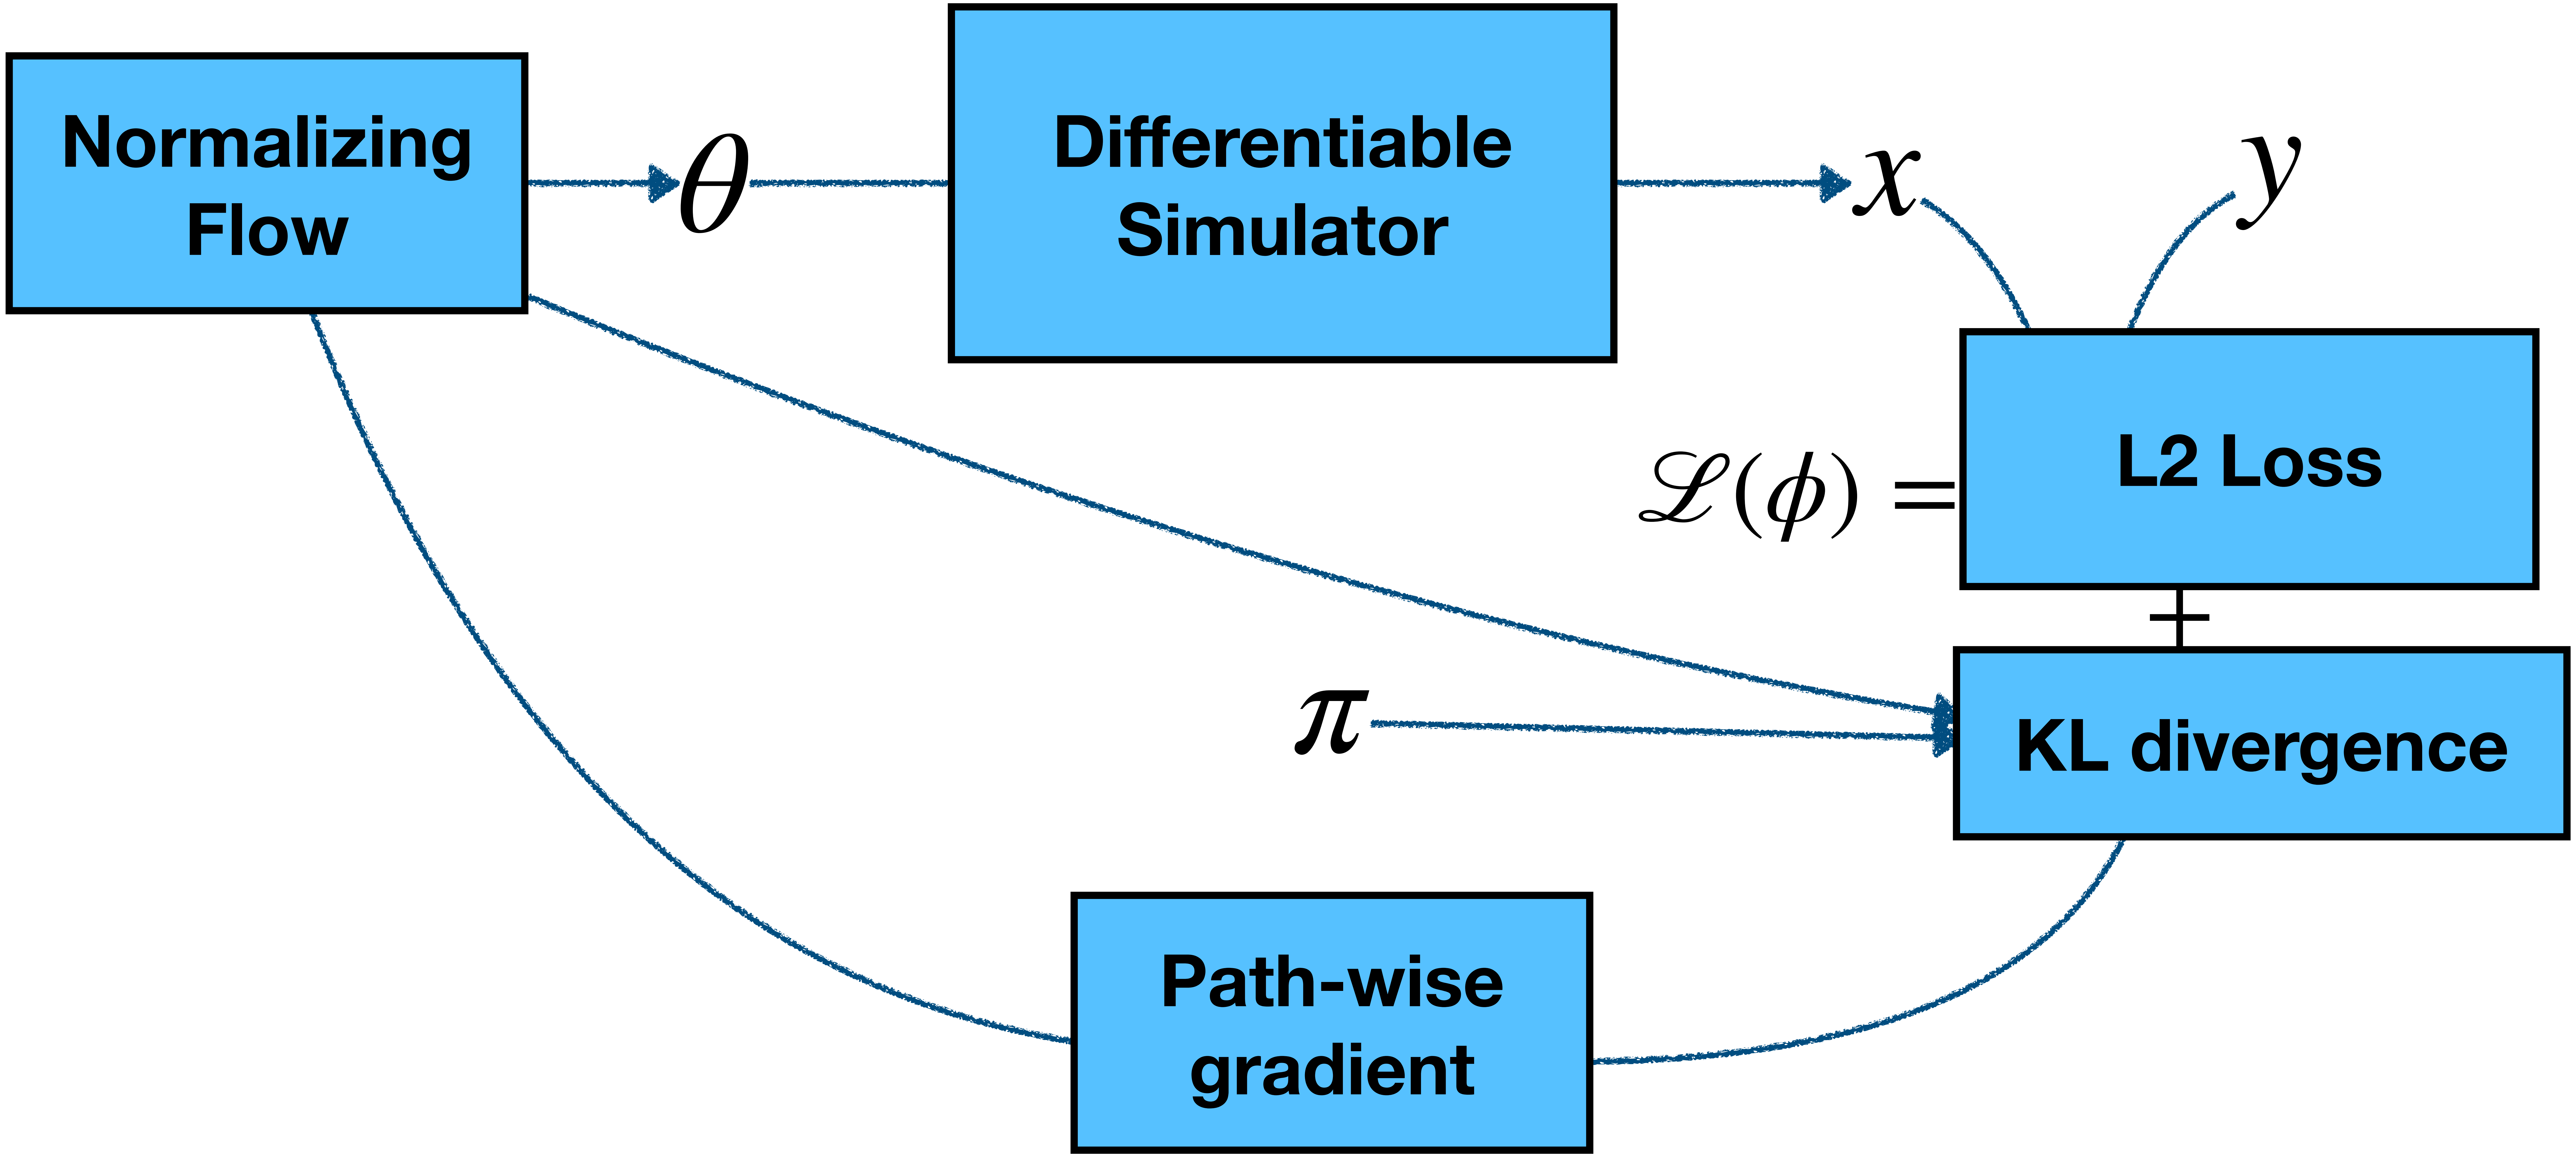

Experiment with JUNE

- ABM model of Covid19
- Model
  - ~200k agents
  - 3 layers of interactions (household, company, school)
  - Calibrate to synthetic data

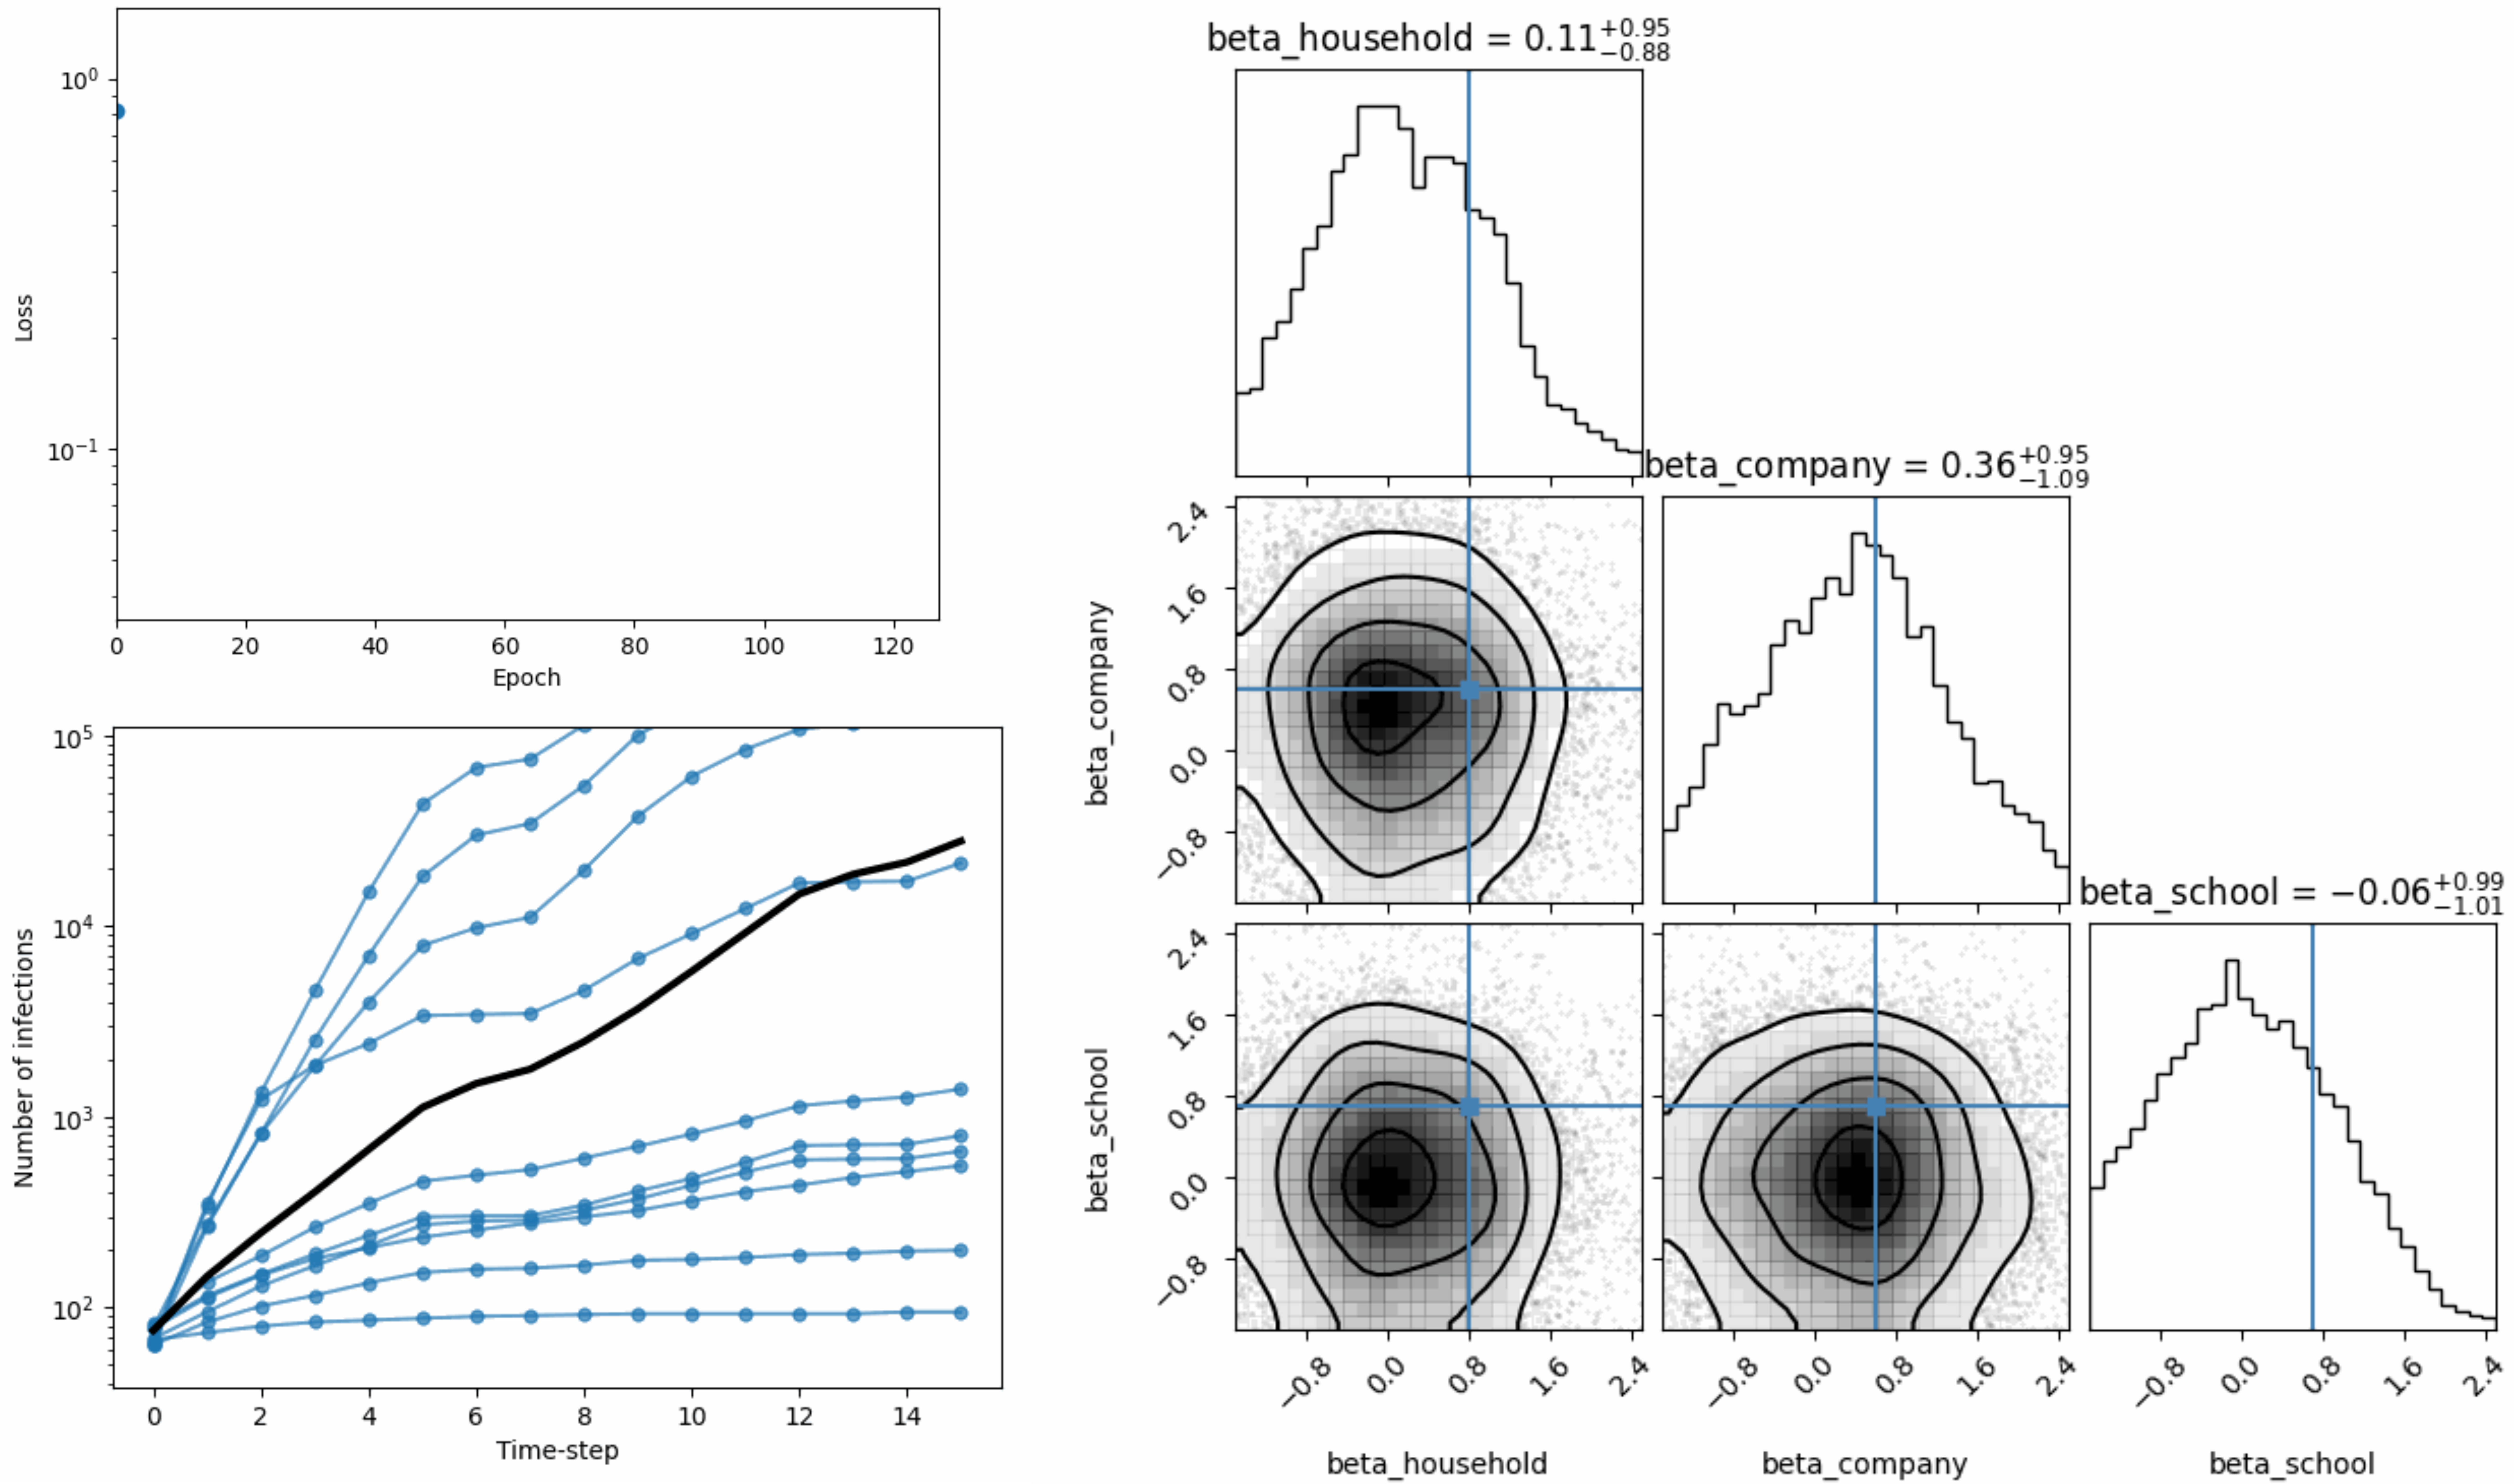

10^0
10^{-1}
Loss
0
20
40
60
80
100
120
Epoch
10^5
10^4
10^3
10^2
Number of infections
0
2
4
6
8
10
12
14
Time-step
beta_household = $0.11^{+0.95}_{-0.88}$
beta_company = $0.36^{+0.95}_{-1.09}$
beta_school = $-0.06^{+0.99}_{-1.01}$
beta_company
beta_school
beta_household
beta_company
beta_school
−0.8
0.0
0.8
1.6
2.4

Conclusions

1. Bayesian approaches to calibrating ABMs have numerous benefits

2. ABMs can be made differentiable even with discrete randomness and control flow

3. Diff simulators + Bayesian inference (via Normalizing Flows) promising route to calibrate large-scale ABMs efficiently

Paper + slides: www.arnau.ai/iclr

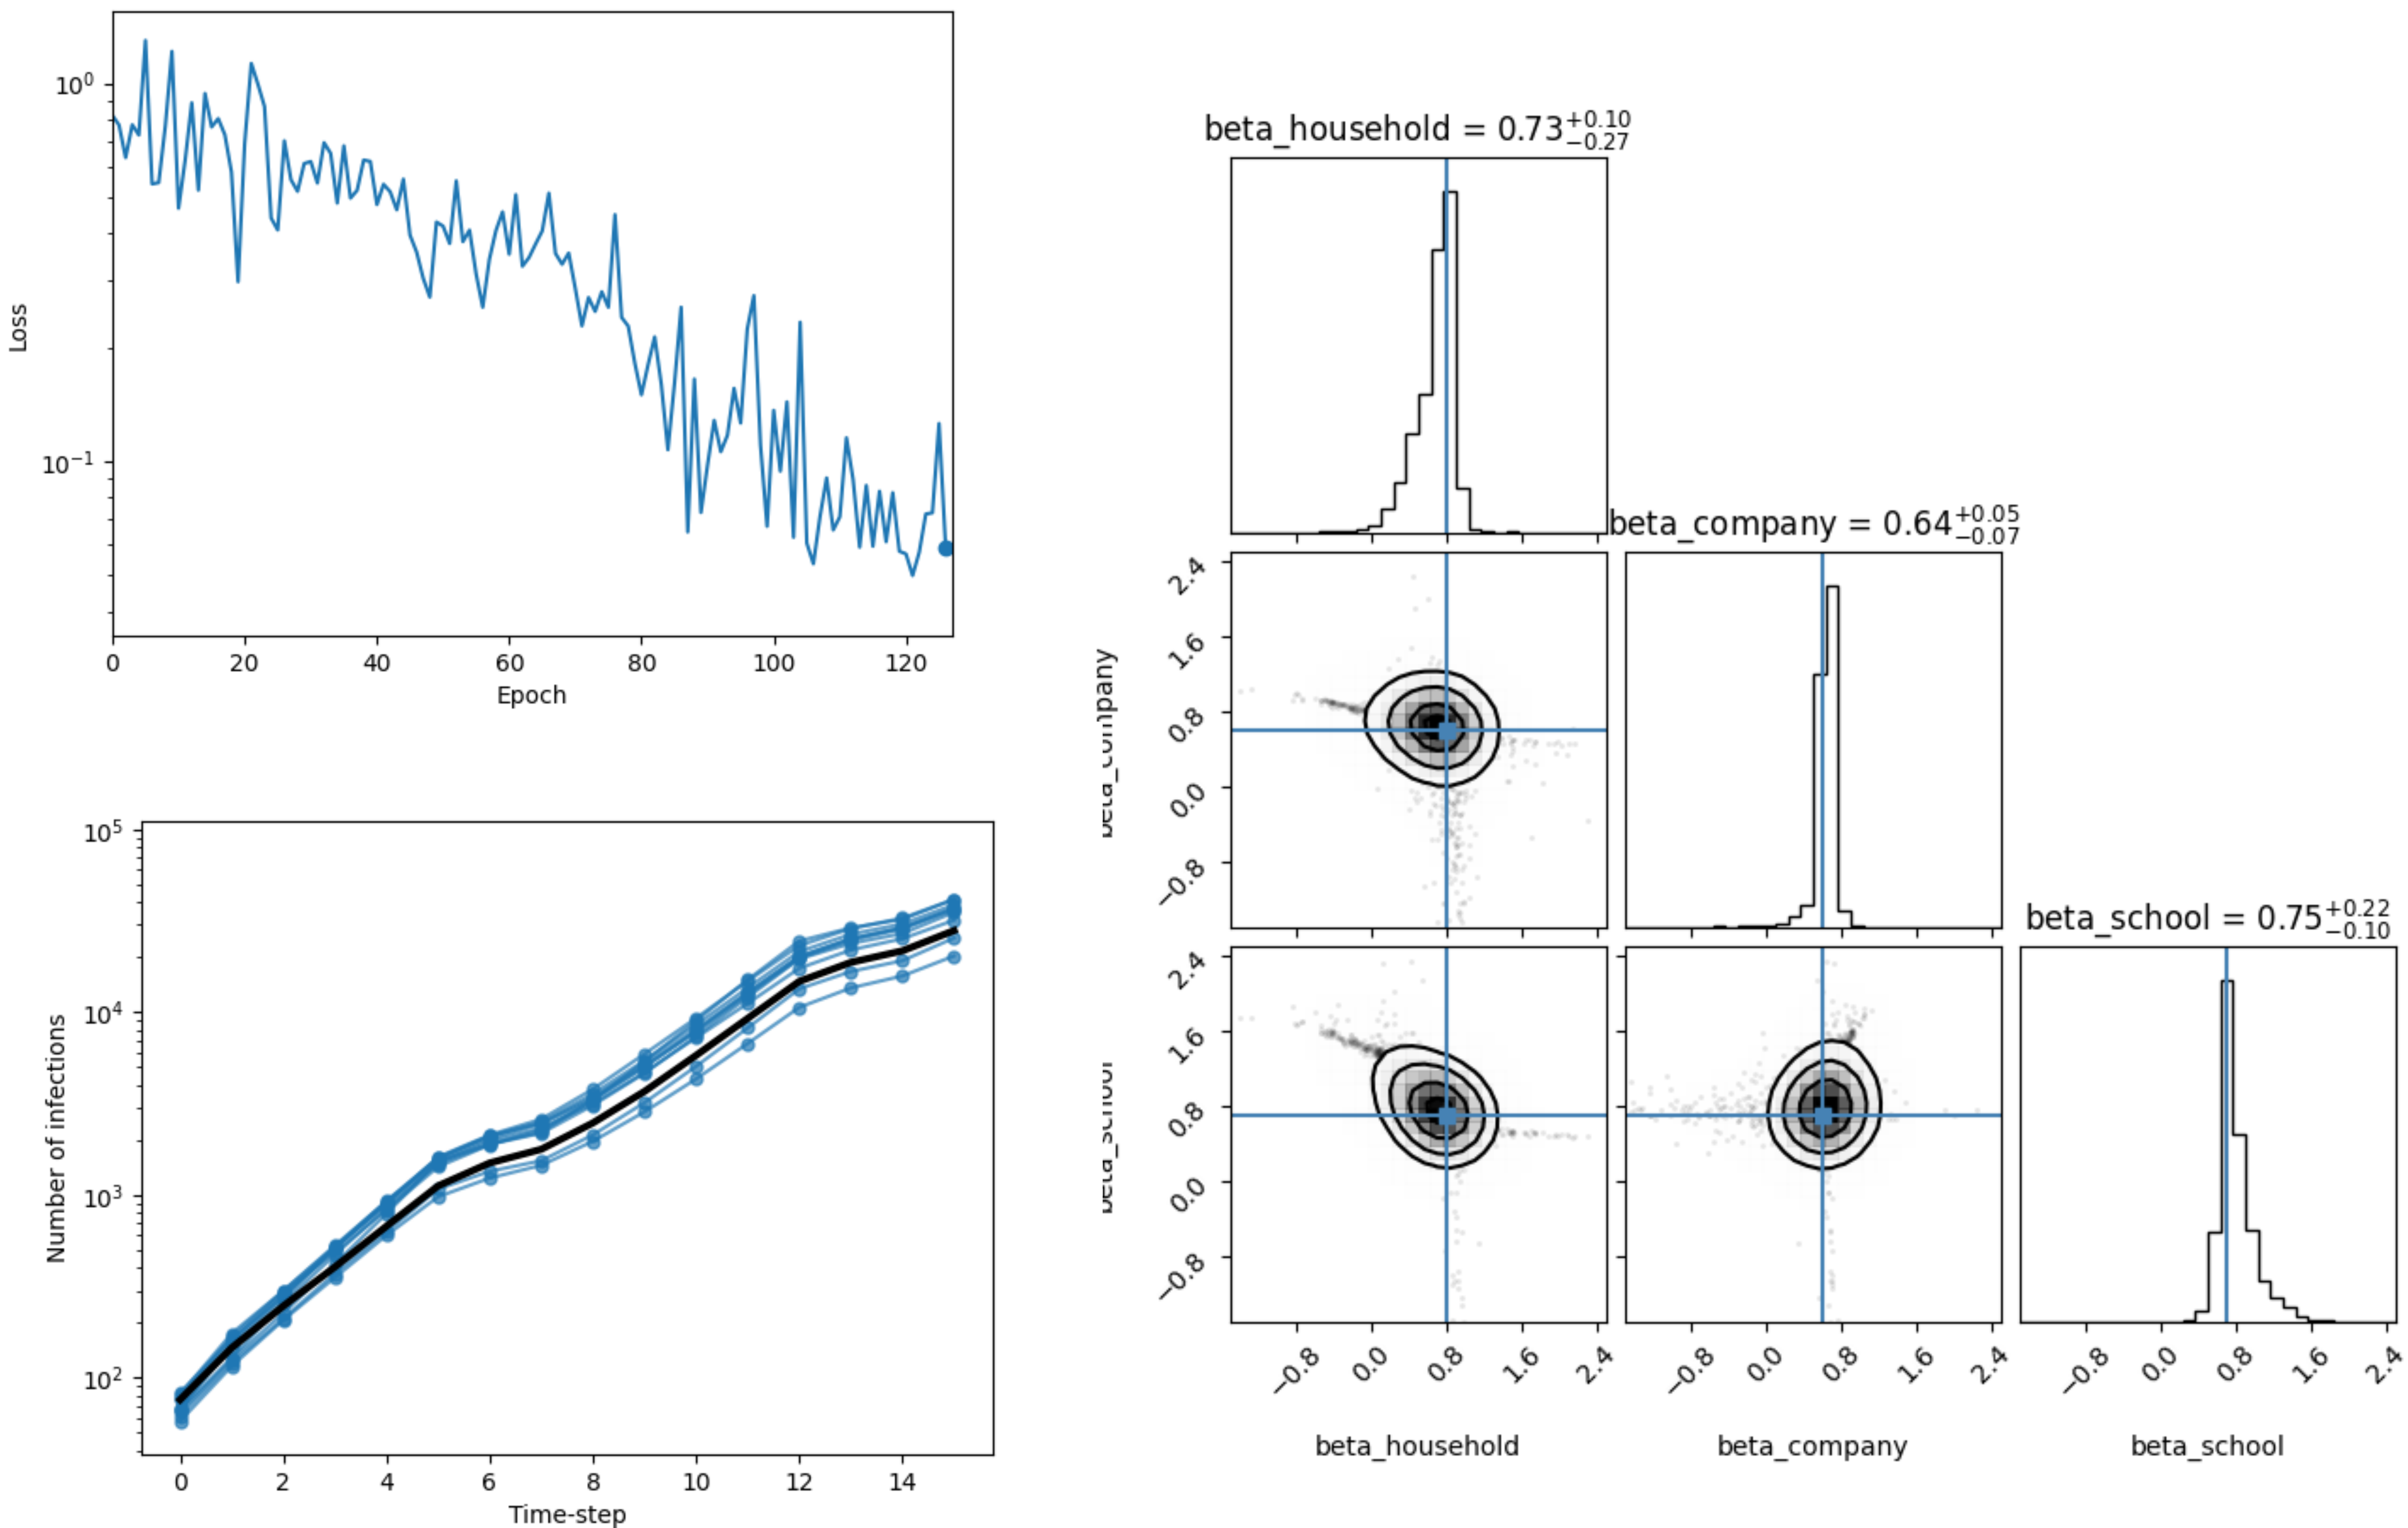

Loss
Epoch
Number of infections
Time-step
beta_household = 0.73$^{+0.10}_{-0.27}$
beta_company = 0.64$^{+0.05}_{-0.07}$
beta_school = 0.75$^{+0.22}_{-0.10}$
beta_household
beta_company
beta_school