



UNIVERSITY OF
OXFORD

Bayesian calibration of differentiable simulators

Arnau Quera-Bofarull, Joel Dyer

Anisoara Calinescu, Michael Wooldridge

BiInference conference – Oxford, June 2023

www.arnau.ai

Differentiable Simulators

$$f : \mathbb{R}^m \longrightarrow \mathbb{R}^n$$

$$(J_f)_{ij} = \frac{\partial f_i}{\partial x_j}$$

Differentiation of Computer Programs

Numerical differentiation

$$\frac{\partial f}{\partial x} = \lim_{h \rightarrow 0} \frac{f(x+h) - f(x)}{h}$$

Inaccurate

Expensive $\mathcal{O}(m \times n)$

$$f : \mathbb{R}^m \longrightarrow \mathbb{R}^n$$

Differentiation of Computer Programs

Symbolic differentiation

$$f(x) = (2x + \sin(x))x^2(x + 3)(x + 5)$$

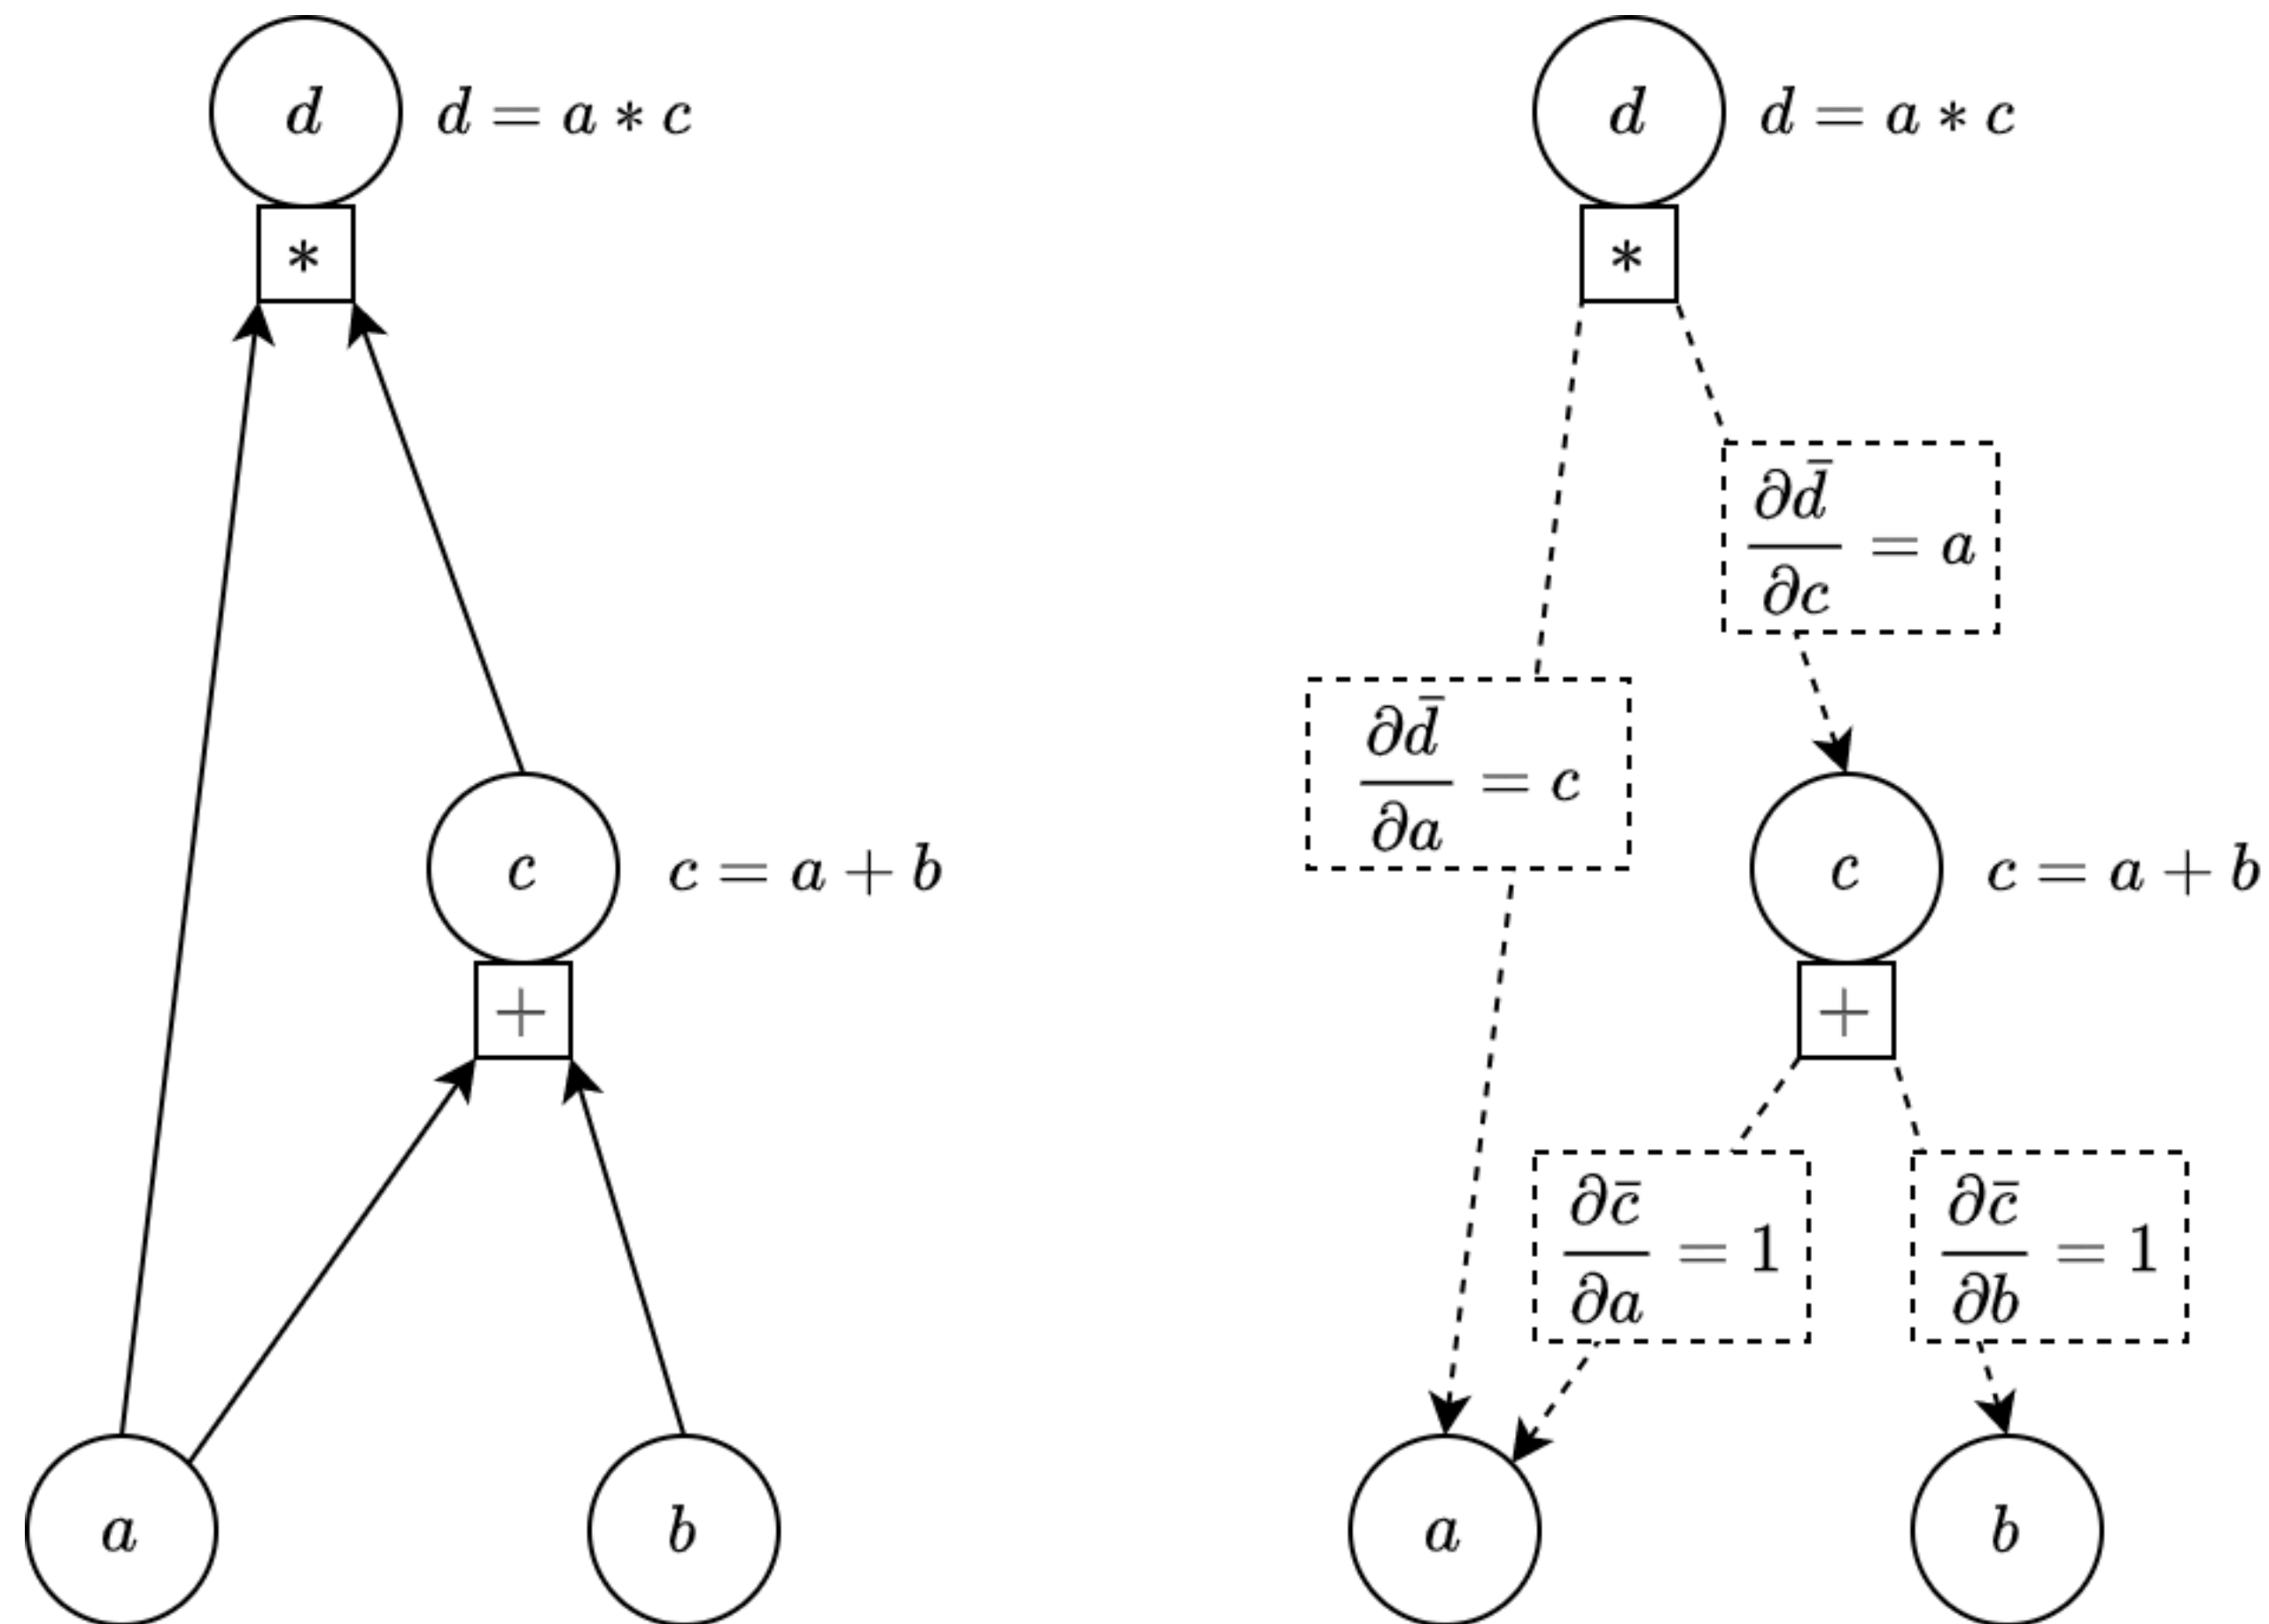
$$f'(x) = x(2x(45 + 32x + 5x^2) + x(15 + 8x + x^2)\cos(x) + (30 + 24x + 4x^2)\sin(x))$$

Exact

Expensive

Differentiation of Computer Programs

Automatic differentiation (AD)



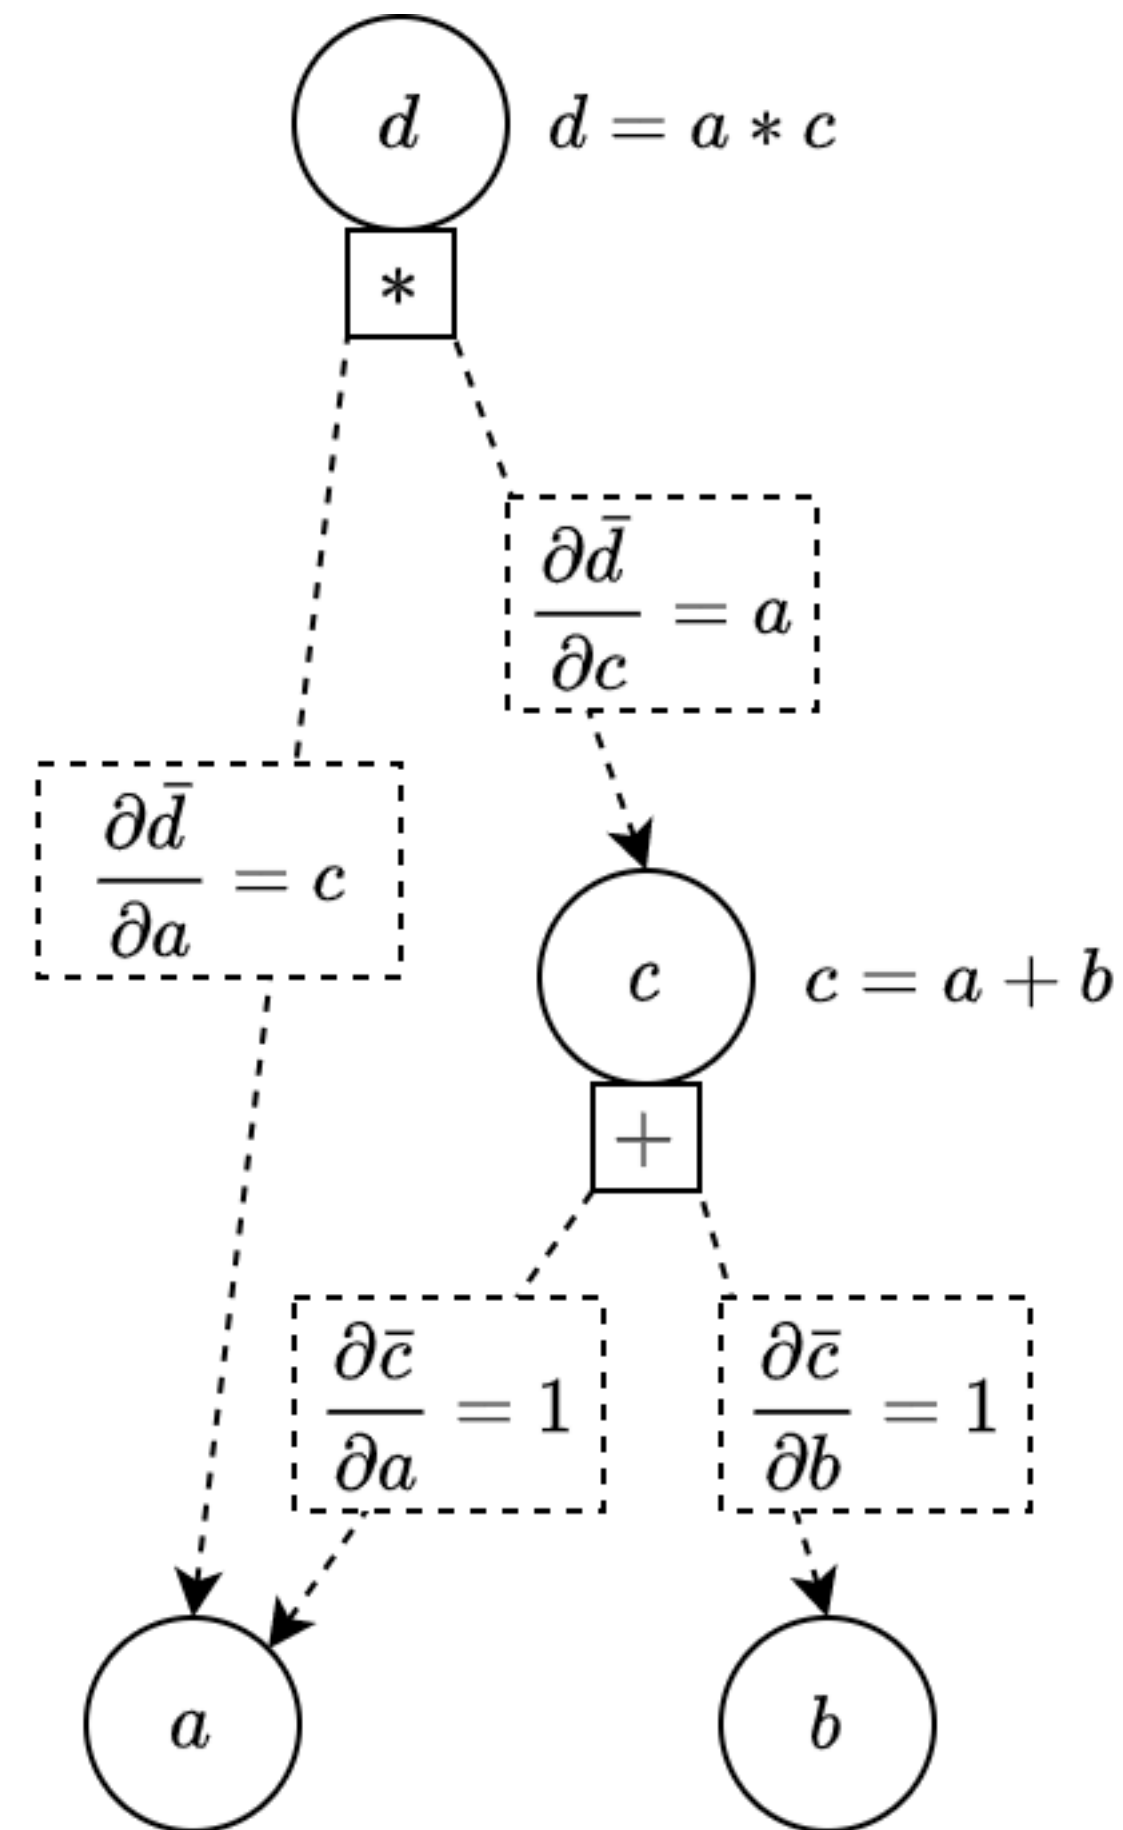

Exact

Fast $\mathcal{O}(\min(m, n))$

$$f : \mathbb{R}^m \longrightarrow \mathbb{R}^n$$


Forward vs Reverse mode AD

$$f : \mathbb{R}^m \longrightarrow \mathbb{R}^n$$



Forward $\sim m$
Reverse $\sim n$

Agent-Based Models



- Expensive to calibrate
- Difficult to interpret / validate

Differentiable Agent-Based Models

Can ABMs be made differentiable?

Stochasticity?

Com

Yes, they can

Memory?

Meaningful gradients?

Case study: the JUNE epidemiological model

- JUNE is a 1:1 epi model of England (56 million agents)
- GradABM-JUNE is its differentiable implementation (PyTorch).

	Simulation
JUNE	50 hours
GRADABM-JUNE (GPU)	5 seconds

Tensorisation enables scalability to millions (billions?) of agents

References:
arnau.ai/talks


Case study: the JUNE epidemiological model

- We can use generalized variational inference for calibration

	Simulation	Calibration (No UQ)	Bayesian Calibration
JUNE	50 hours	-	100k hours
GRADABM-JUNE (GPU)	5 seconds	20 minutes	8 hours

**Differentiability enables fast
and accurate model calibration**

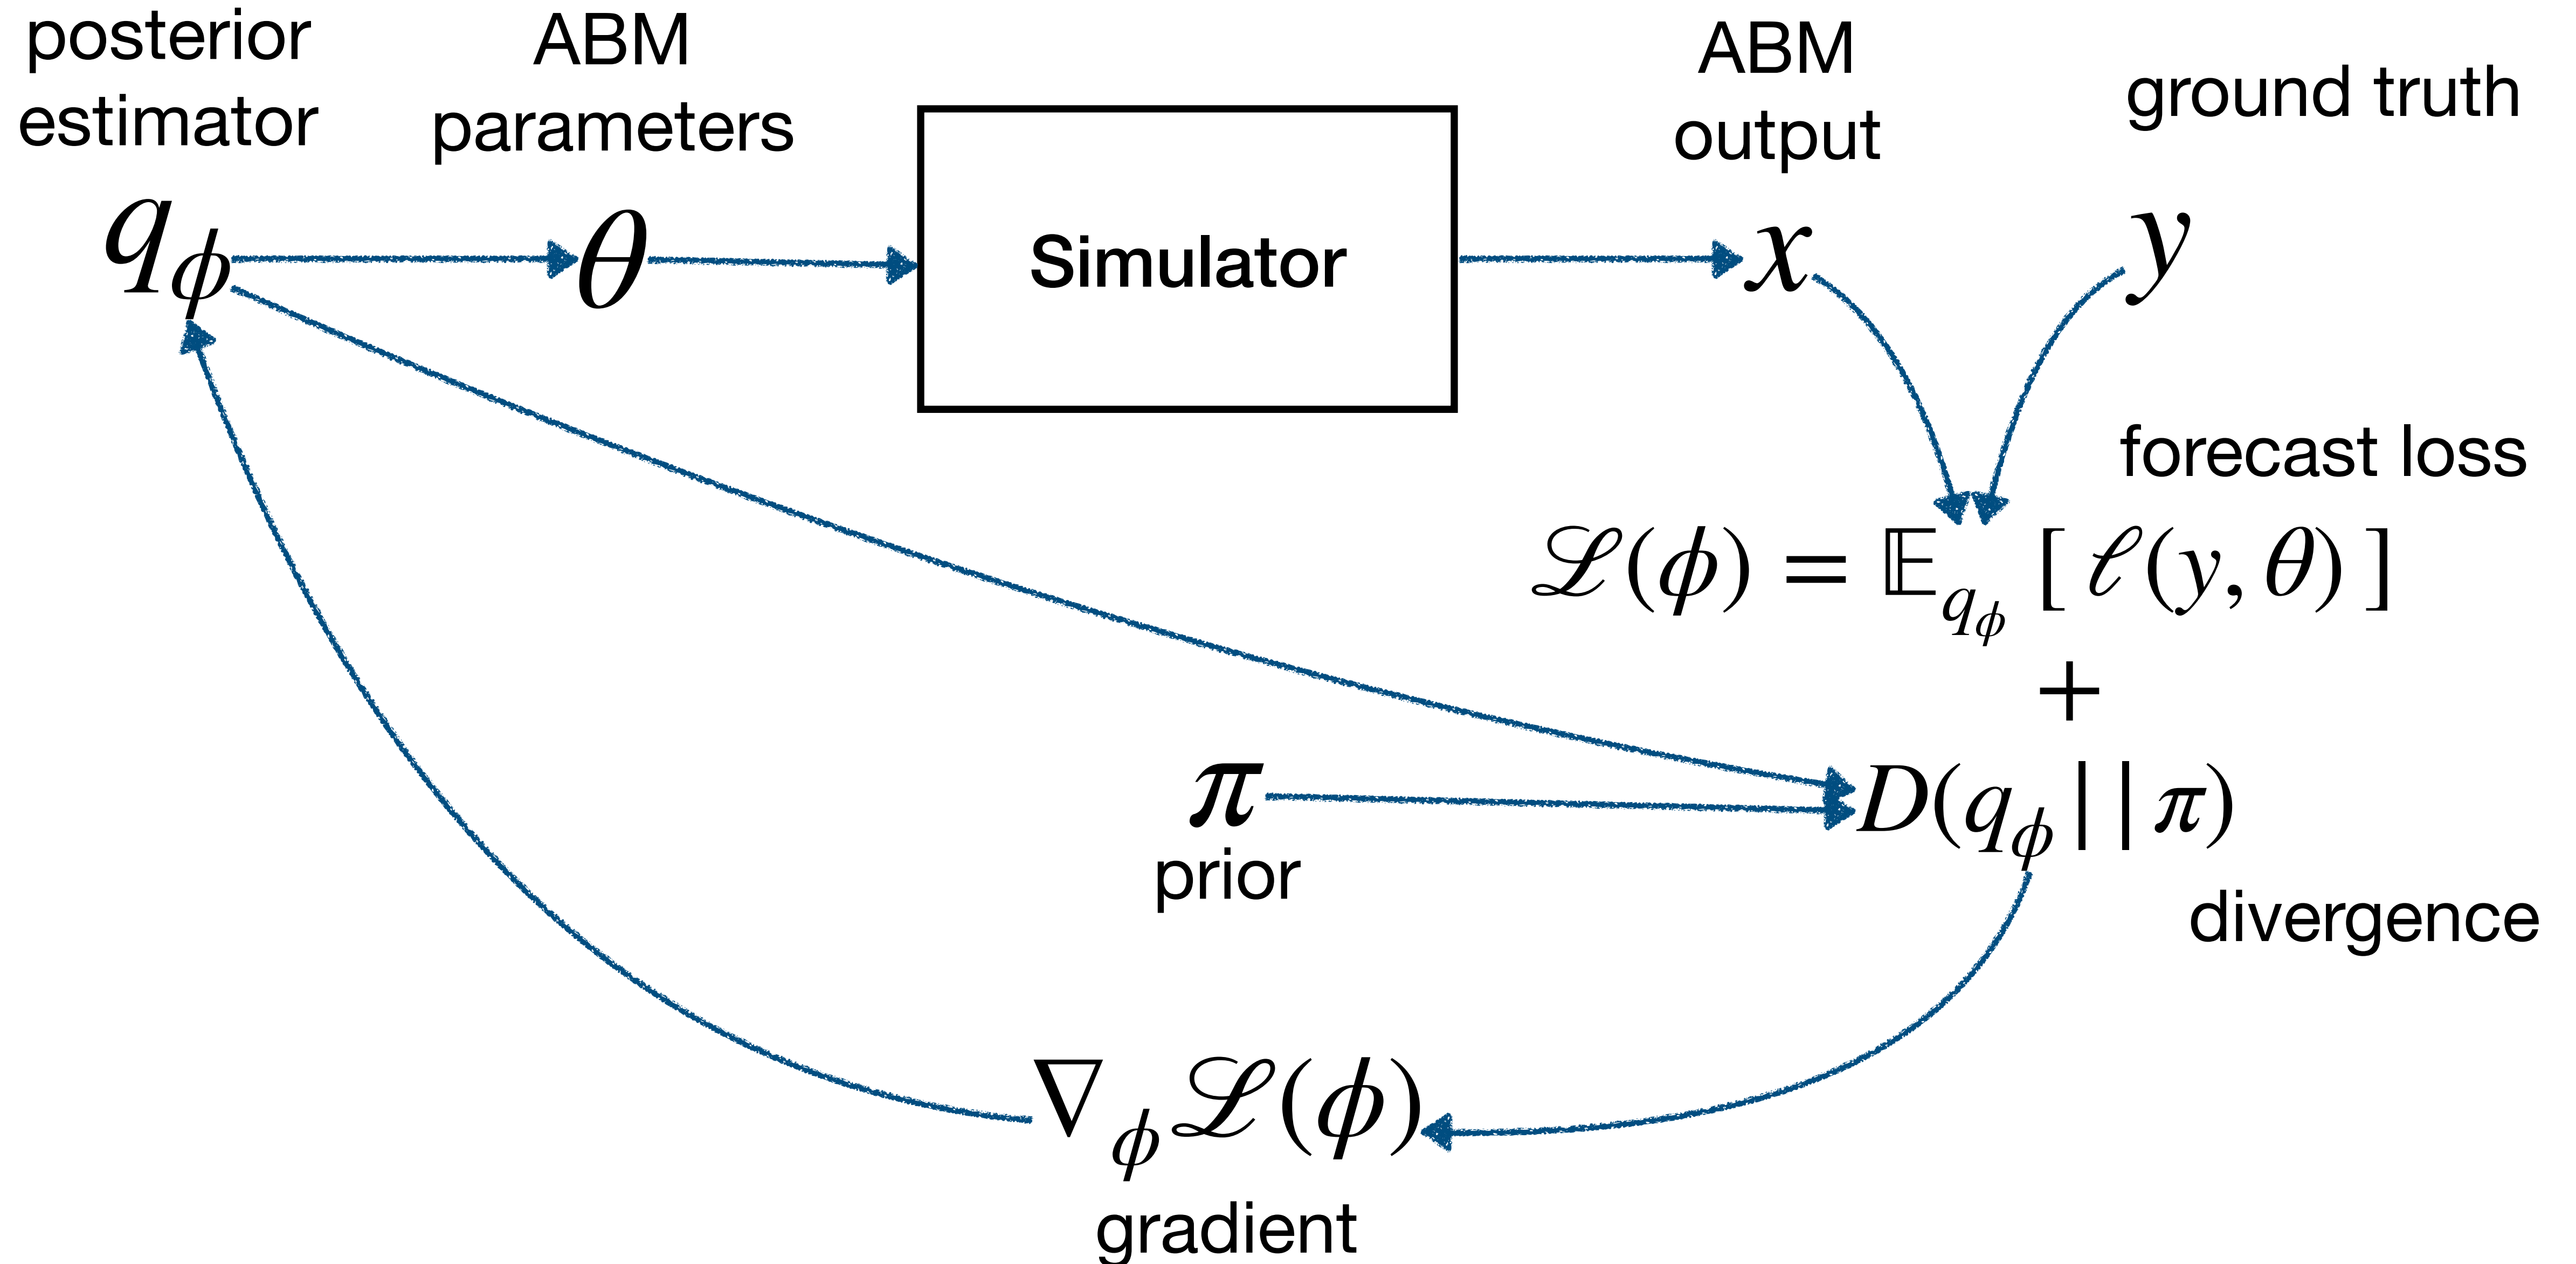
Variational Inference: Bayesian inference as an optimisation problem



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

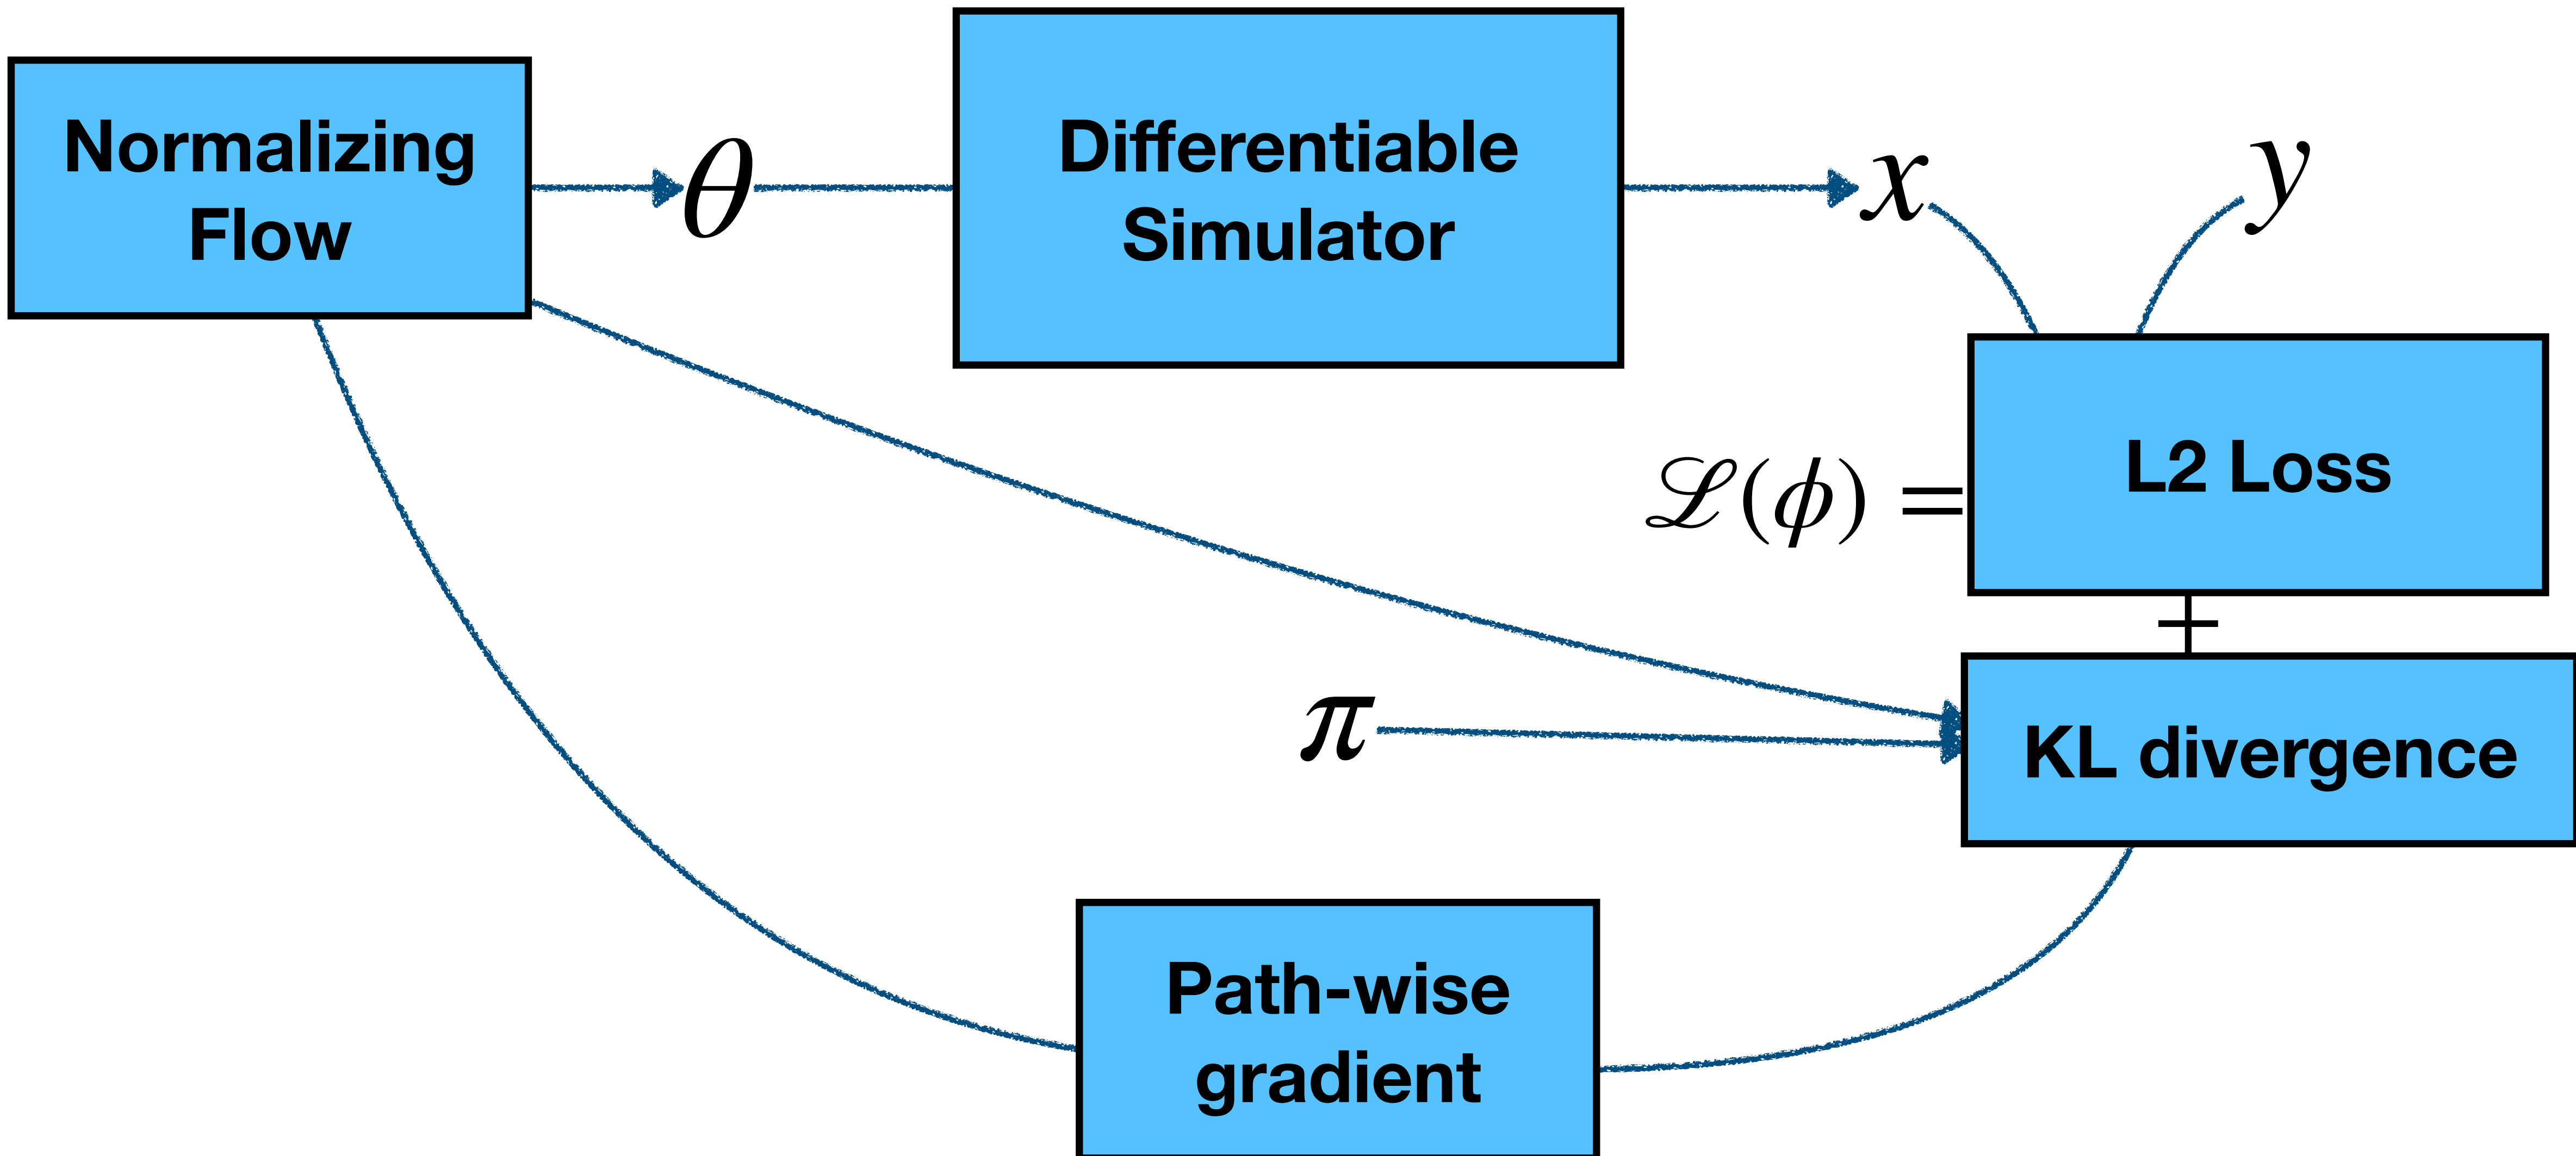
- Gradient-assisted calibration algorithms need



- 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))



Generalised Variational Inference



Normalizing Flows

What do we choose for q ?


Image credit: Lilian Weng



Experiment with JUNE

- > 8M agents (London)
- 10 layers of interactions (household, company, school, pub, ...)
- 1 Seed parameter
- Calibrate to **synthetic** data

Gradient Horizon Problem



Model does not train!

Gradient Horizon Problem

Suppose we run an ABM with parameters θ for 4 time-steps

x_1	x_2	x_3	x_4
-------	-------	-------	-------

$$x_1 = f(\theta)$$

$$x_2 = f(x_1, \theta)$$

$$x_3 = f(x_1, x_2, \theta)$$

$$x_4 = \dots$$


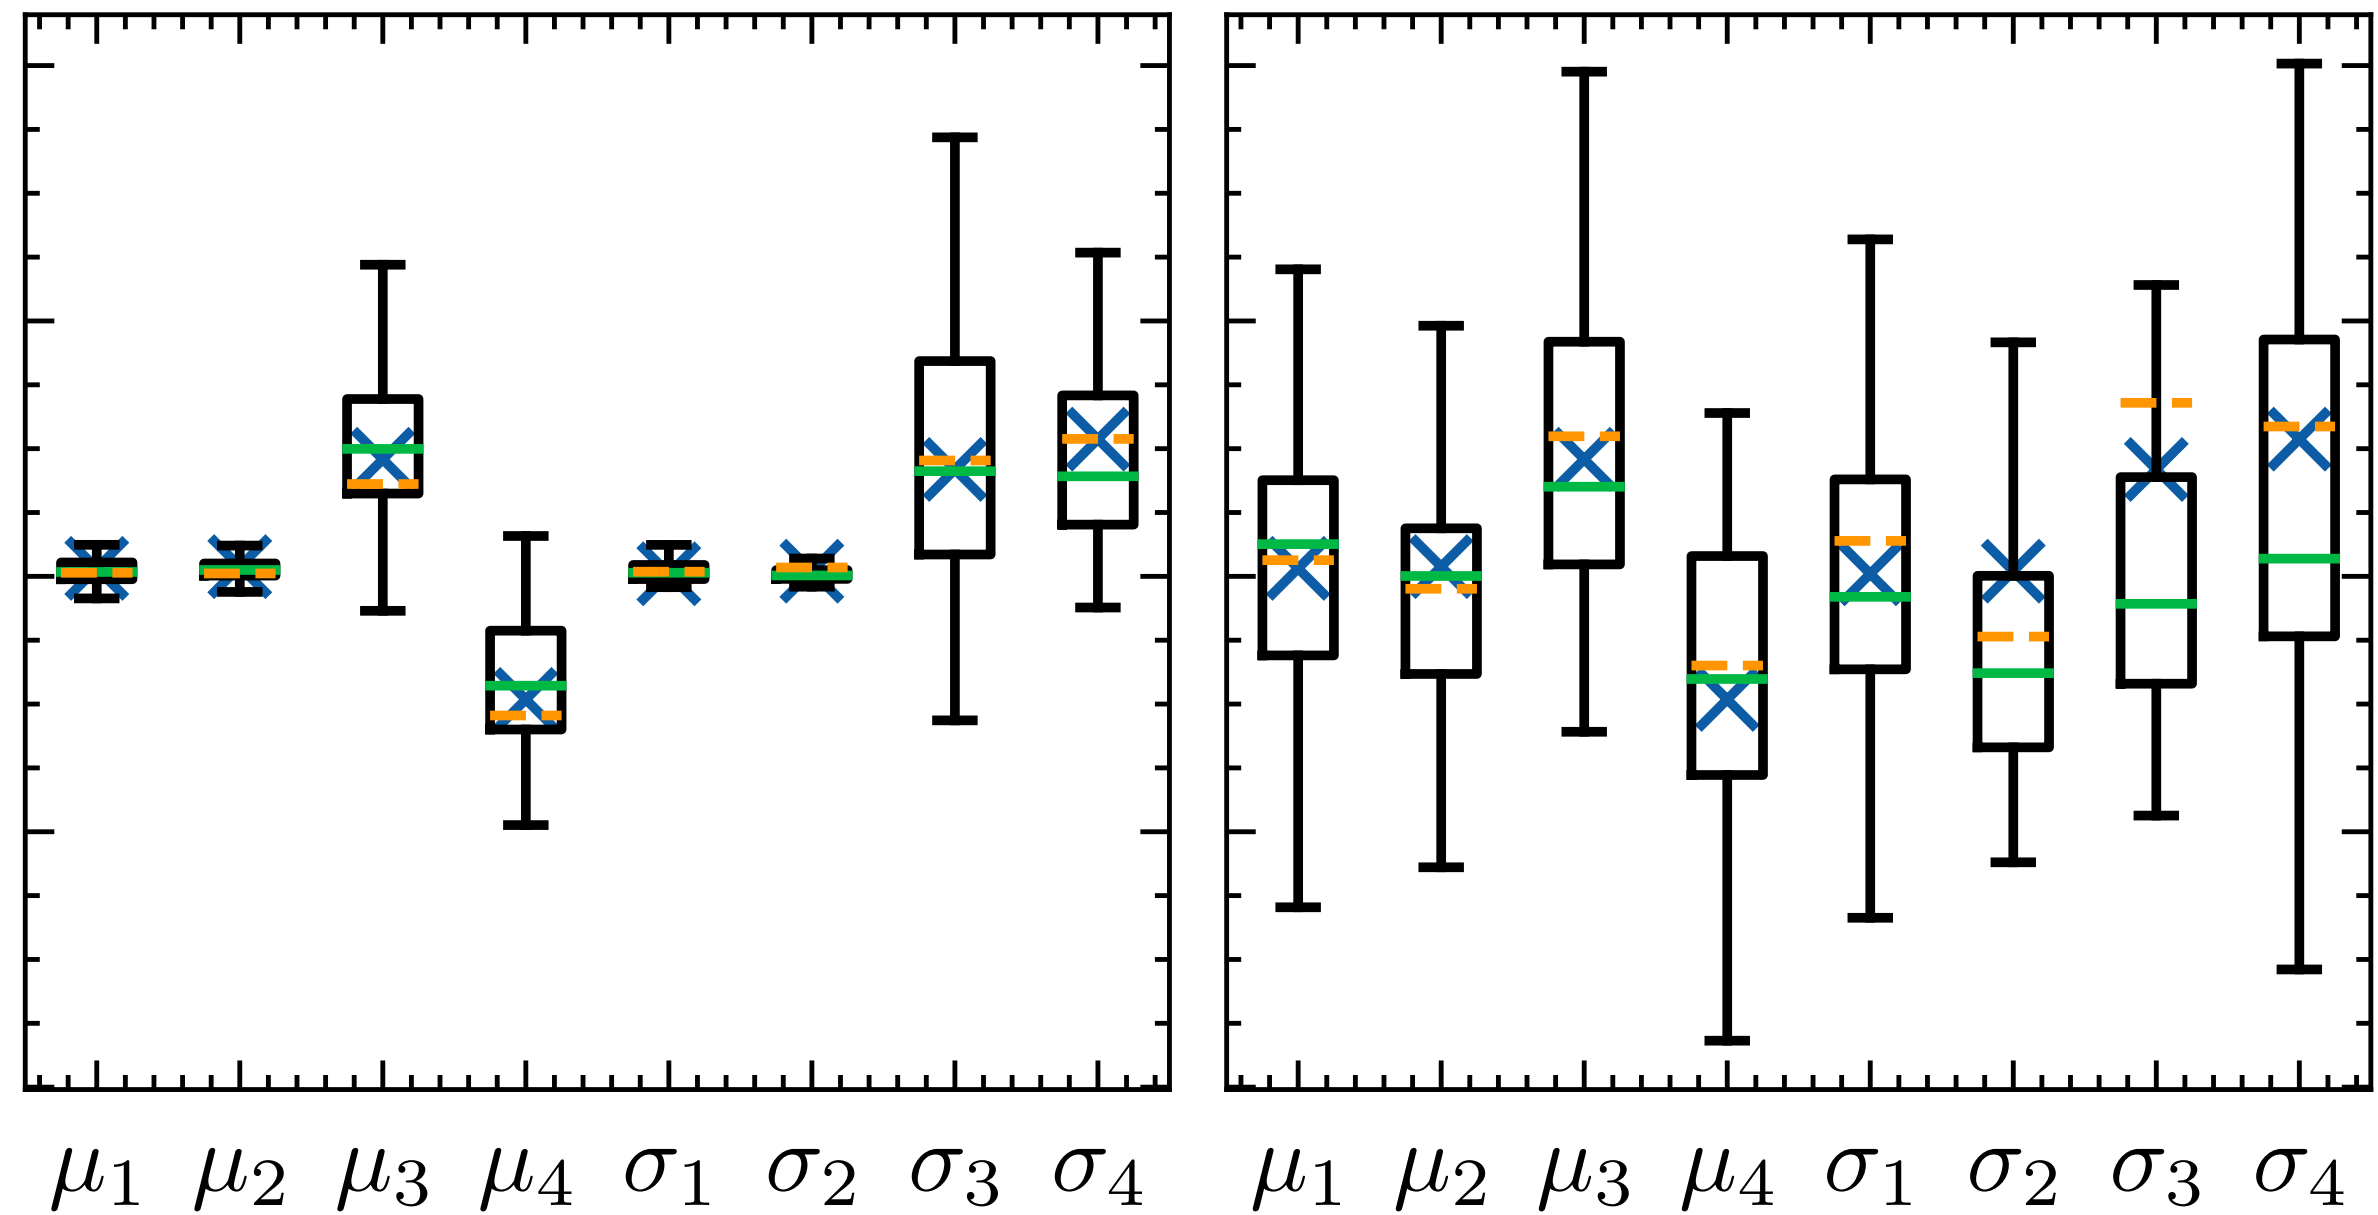
$$\begin{aligned} \frac{dx_3}{d\theta} &\equiv \frac{\partial f}{\partial x_1} \frac{\partial x_1}{\partial \theta} + \frac{\partial f}{\partial x_2} \frac{\partial x_2}{\partial \theta} \\ &+ \frac{\partial f}{\partial x_2} \frac{\partial x_2}{\partial x_1} \frac{\partial x_1}{\partial \theta} \\ &+ \frac{\partial f}{\partial \theta} \end{aligned}$$

Gradient Horizon Problem


$$\frac{dx_3}{d\theta} \underset{H=100}{=} \frac{\cancel{\partial f}}{\cancel{\partial x_1}} \frac{\partial x_1}{\partial \theta} + \frac{\cancel{\partial f}}{\cancel{\partial x_2}} \frac{\partial x_2}{\partial \theta}$$




Score function

Truncating the gradient reduces variance!


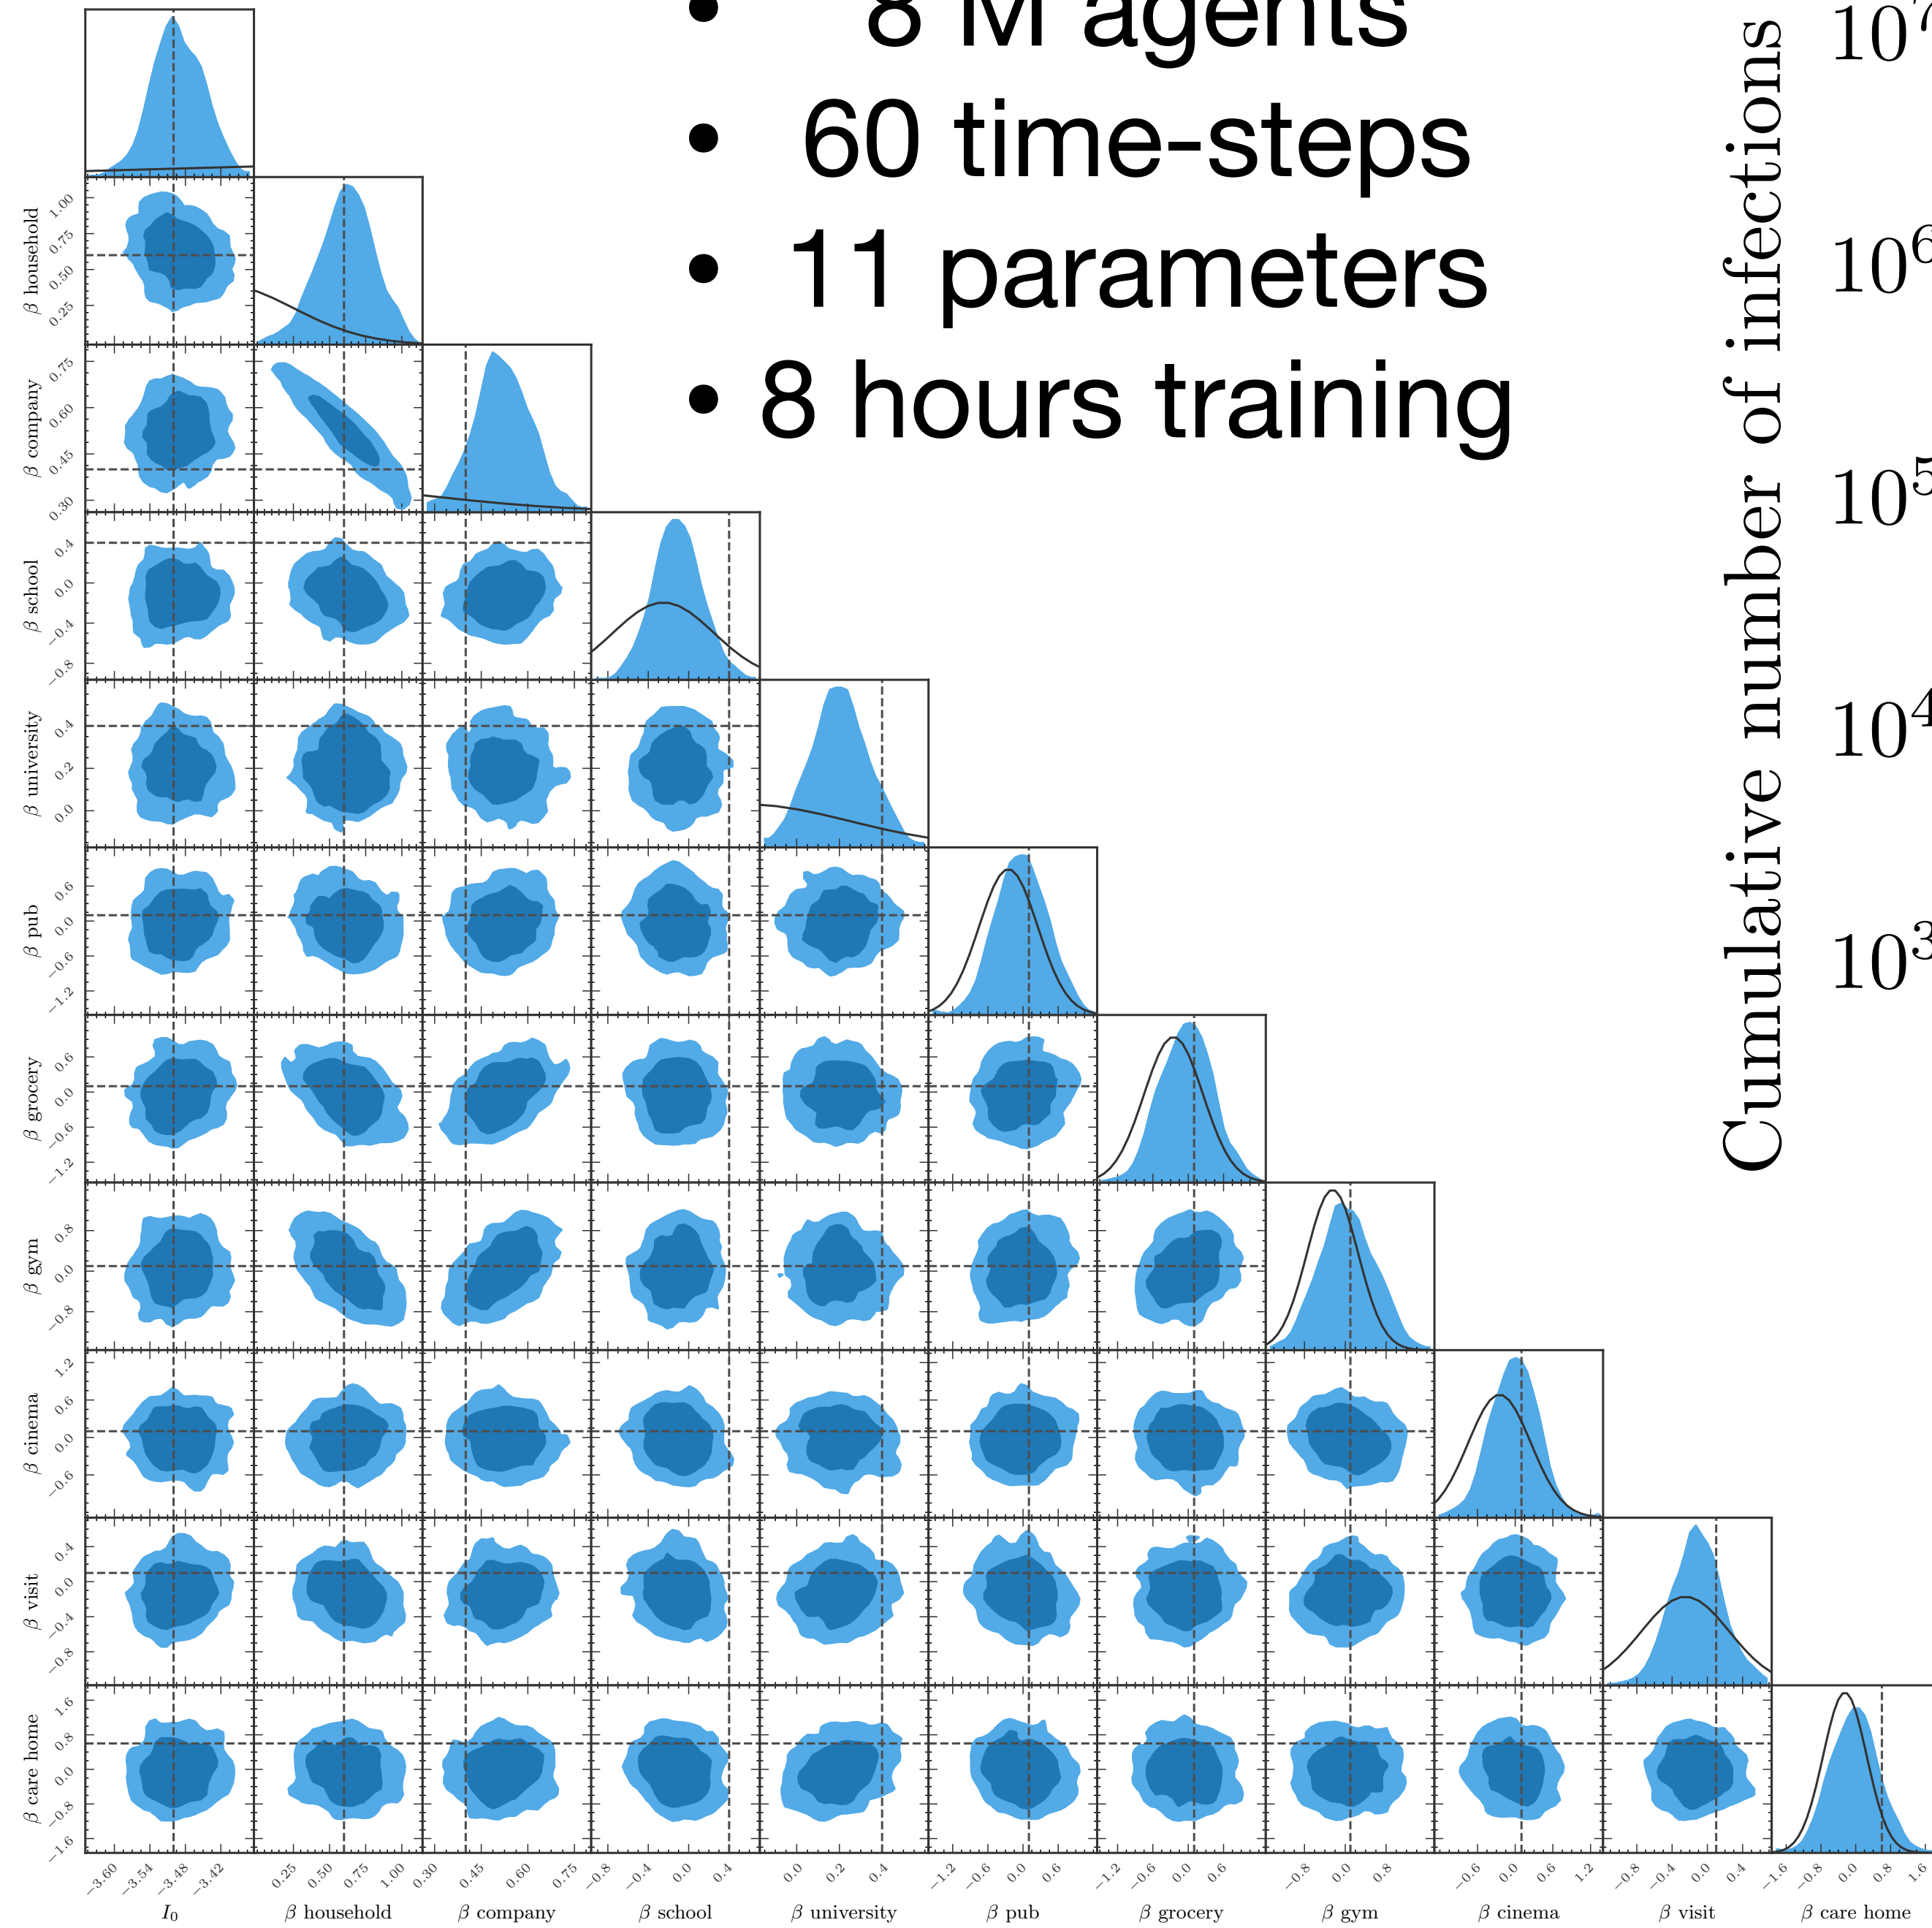


Gradient Horizon Problem

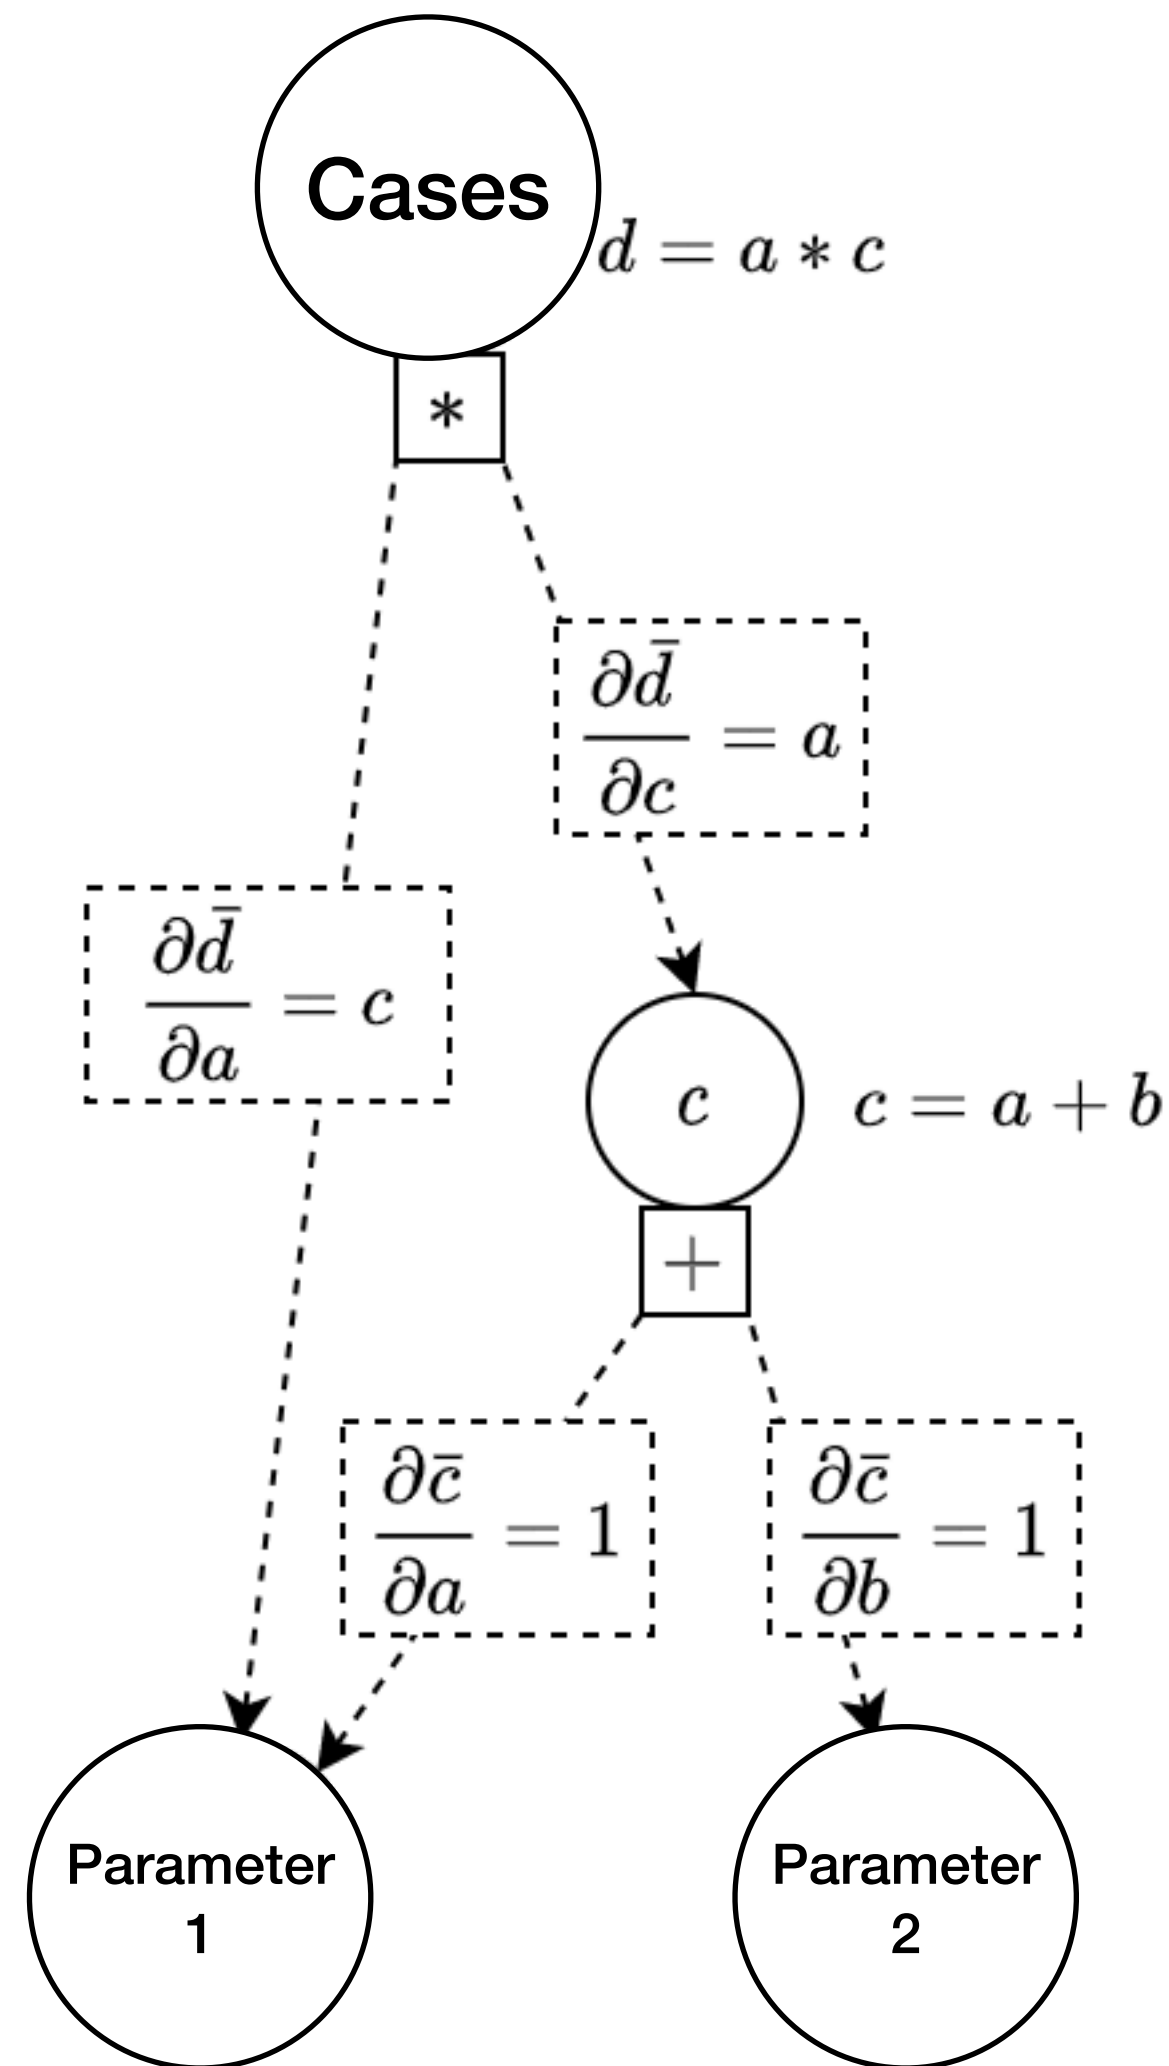




- 8 M agents
- 60 time-steps
- 11 parameters
- 8 hours training



Sensitivity Analyses (SA)



AD performs SA
with a single simulation run,
independent of # of parameters!

The impact of uncertainty on predictions of the CovidSim epidemiological code



Wouter Edeling¹, Hamid Arabnejad², Robbie Sinclair³, Diana Suleimenova²,
Krishnakumar Gopalakrishnan³, Bartosz Bosak⁴, Derek Groen², Imran Mahmood²,
Daan Crommelin^{1,5} and Peter V. Coveney^{3,6} ✉

Epidemiological modelling has assisted in identifying interventions that reduce the impact of COVID-19. The UK government relied, in part, on the CovidSim model to guide its policy to contain the rapid spread of the COVID-19 pandemic during March and April 2020; however, CovidSim contains several sources of uncertainty that affect the quality of its predictions: parametric uncertainty, model structure uncertainty and scenario uncertainty. **Here we report on parametric sensitivity analysis and uncertainty quantification of the code. From the 940 parameters used as input into CovidSim, we find a subset of 19 to which the code output is most sensitive**—imperfect knowledge of these inputs is magnified in the outputs by up to 300%. The model displays substantial bias with respect to observed data, failing to describe validation data well. Quantifying parametric input uncertainty is therefore not sufficient: the effect of model structure and scenario uncertainty must also be properly understood.

Ensemble execution. Consequently, through the use of adaptive methods we make the uncertainty analysis of CovidSim tractable, but our analysis nevertheless required us to perform thousands of runs, each with its own unique set of input parameters. Specifically, we used the Eagle supercomputer at the Posnan

**Reverse-mode AD
independent of number of
parameters!**

Sensitivity Analyses



Conclusions

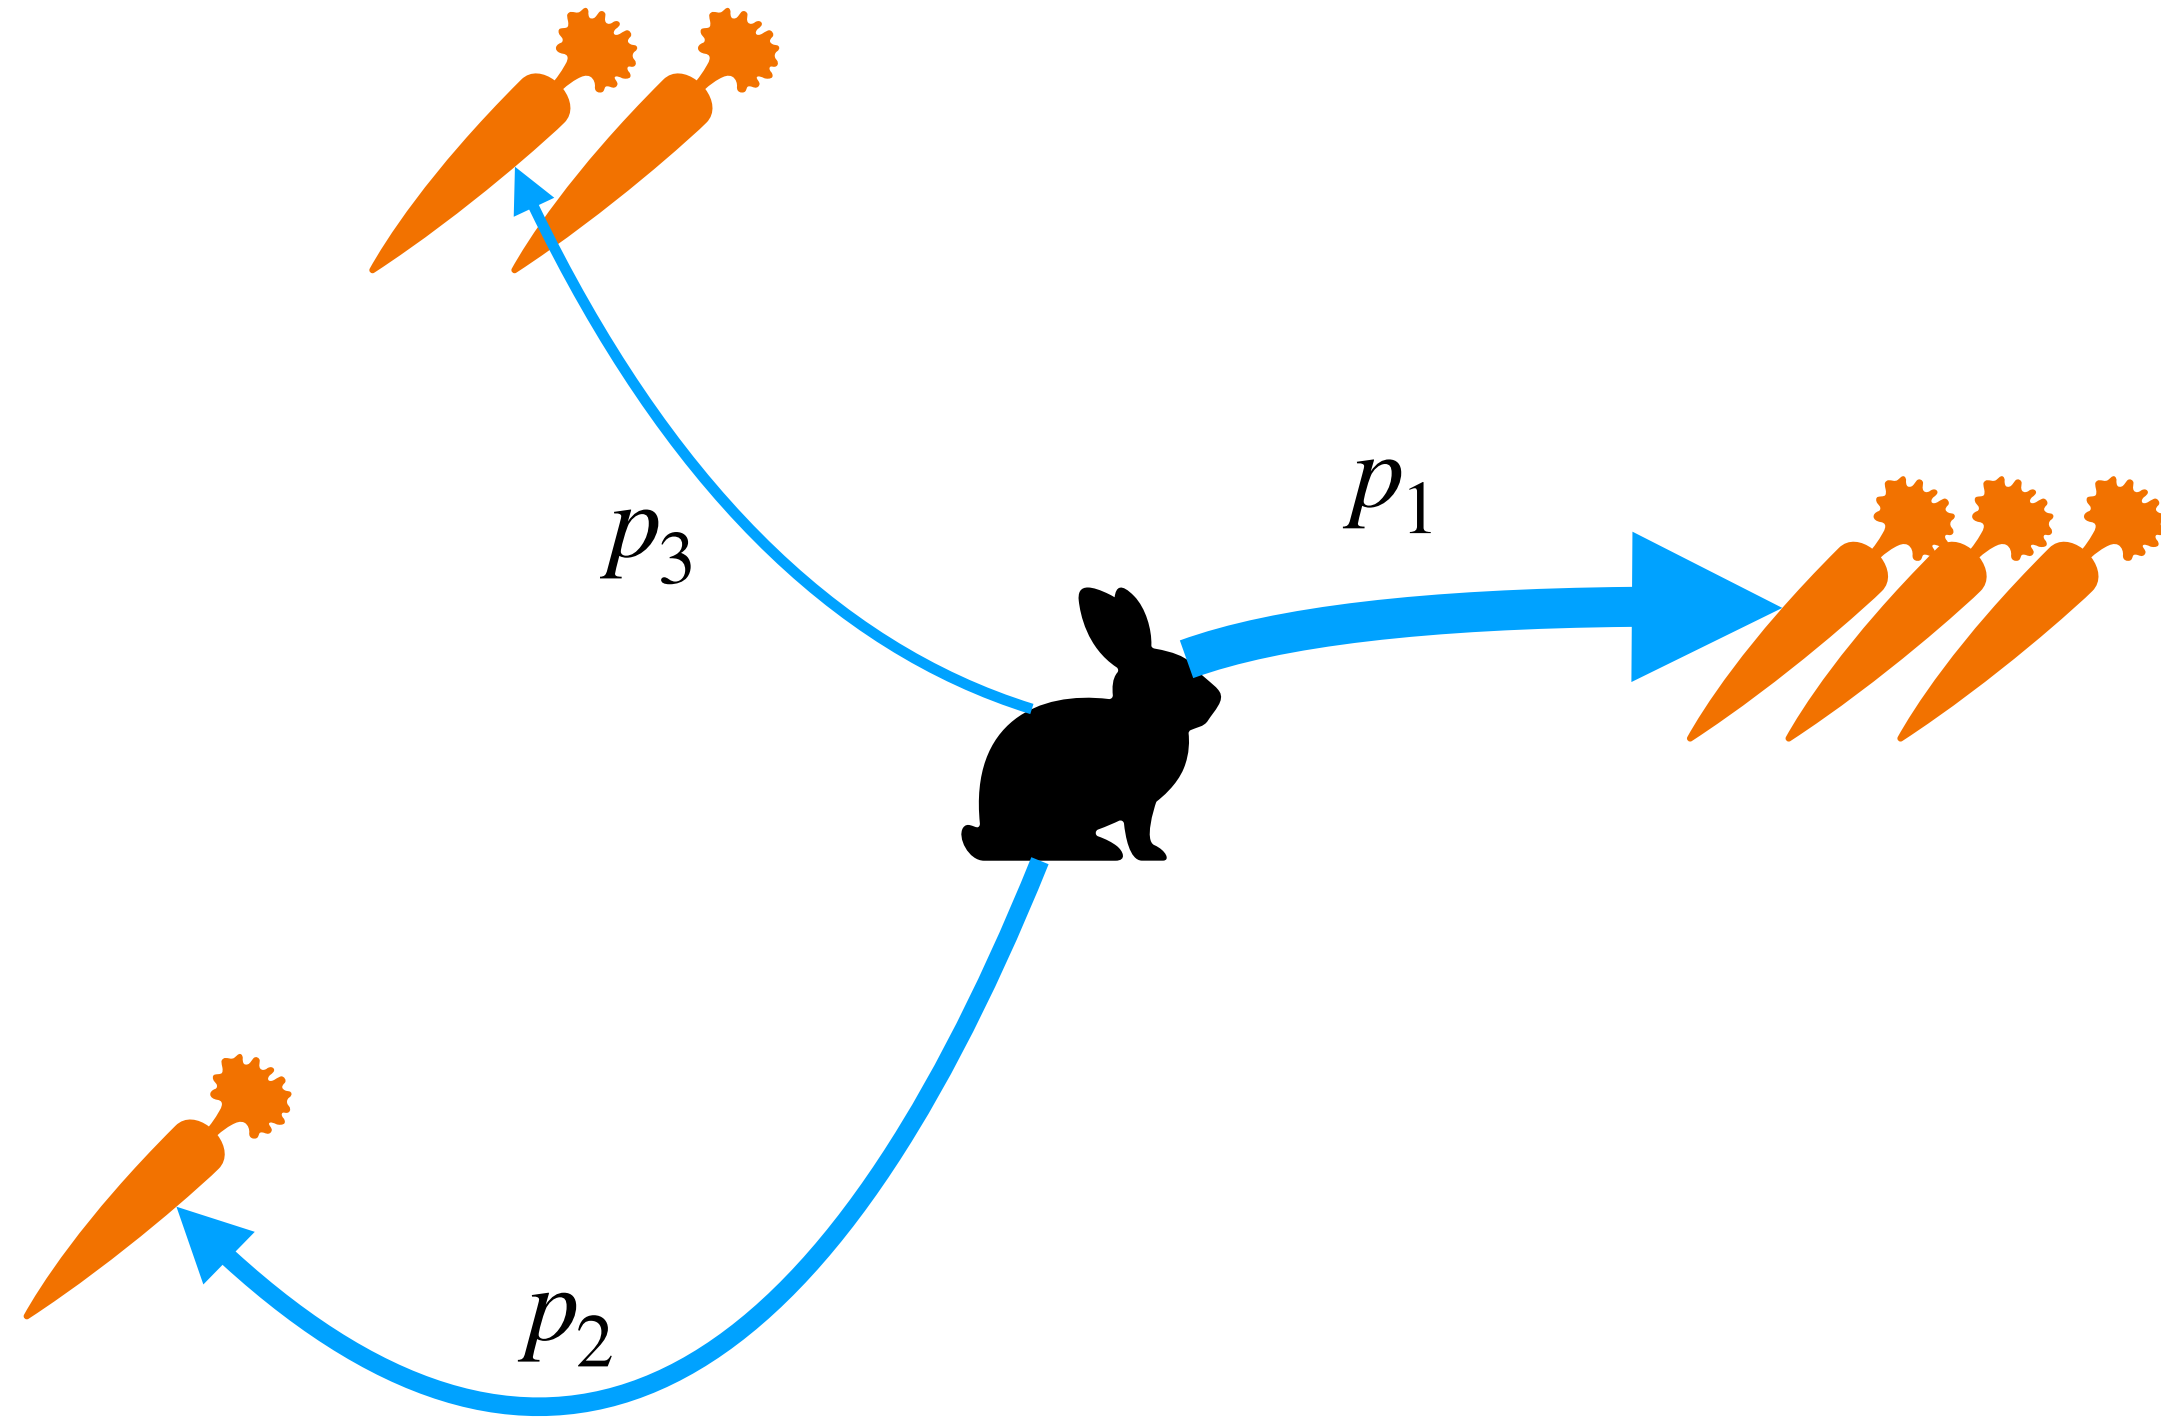
Differentiable agent-based models enable:

1. Fast **simulation** via tensorisation.
2. Fast and accurate Bayesian **calibration** via gradients.
3. Fast and accurate **sensitivity analyses** via gradients.

Papers + slides: arnau.ai/talks

Backup slides

Differentiable Control Flow



$$x' = \text{Argmax} \left(\begin{array}{c} \text{carrot} \text{ carrot} \text{ carrot} \\ \text{carrot} \cdot \text{carrot} \\ \text{carrot} \end{array} \right)$$

$$x' = \text{Softmax} \left(\begin{array}{c} \text{carrot} \text{ carrot} \text{ carrot} \\ \text{carrot} \cdot \text{carrot} \\ \text{carrot} \end{array} \right)$$

Differentiable Stochasticity

Continuous Variables


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

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

Differentiable Stochasticity

Discrete Variables

- Gumbel-Softmax



Jang et al. (2016)