Bayesian Probabilistic Modeling for Ecology

Paul Bürkner

You are a Bayesian if you quantify uncertainty with probability

The Bayes Theorem

$$p(\theta \mid y) = \frac{p(y \mid \theta)\, p(\theta)}{p(y)}$$

Rethinking the Bayes Theorem

$$p(\theta \mid y) \propto p(y \mid \theta)\, p(\theta) = p(y, \theta)$$

Example: Catching Fish

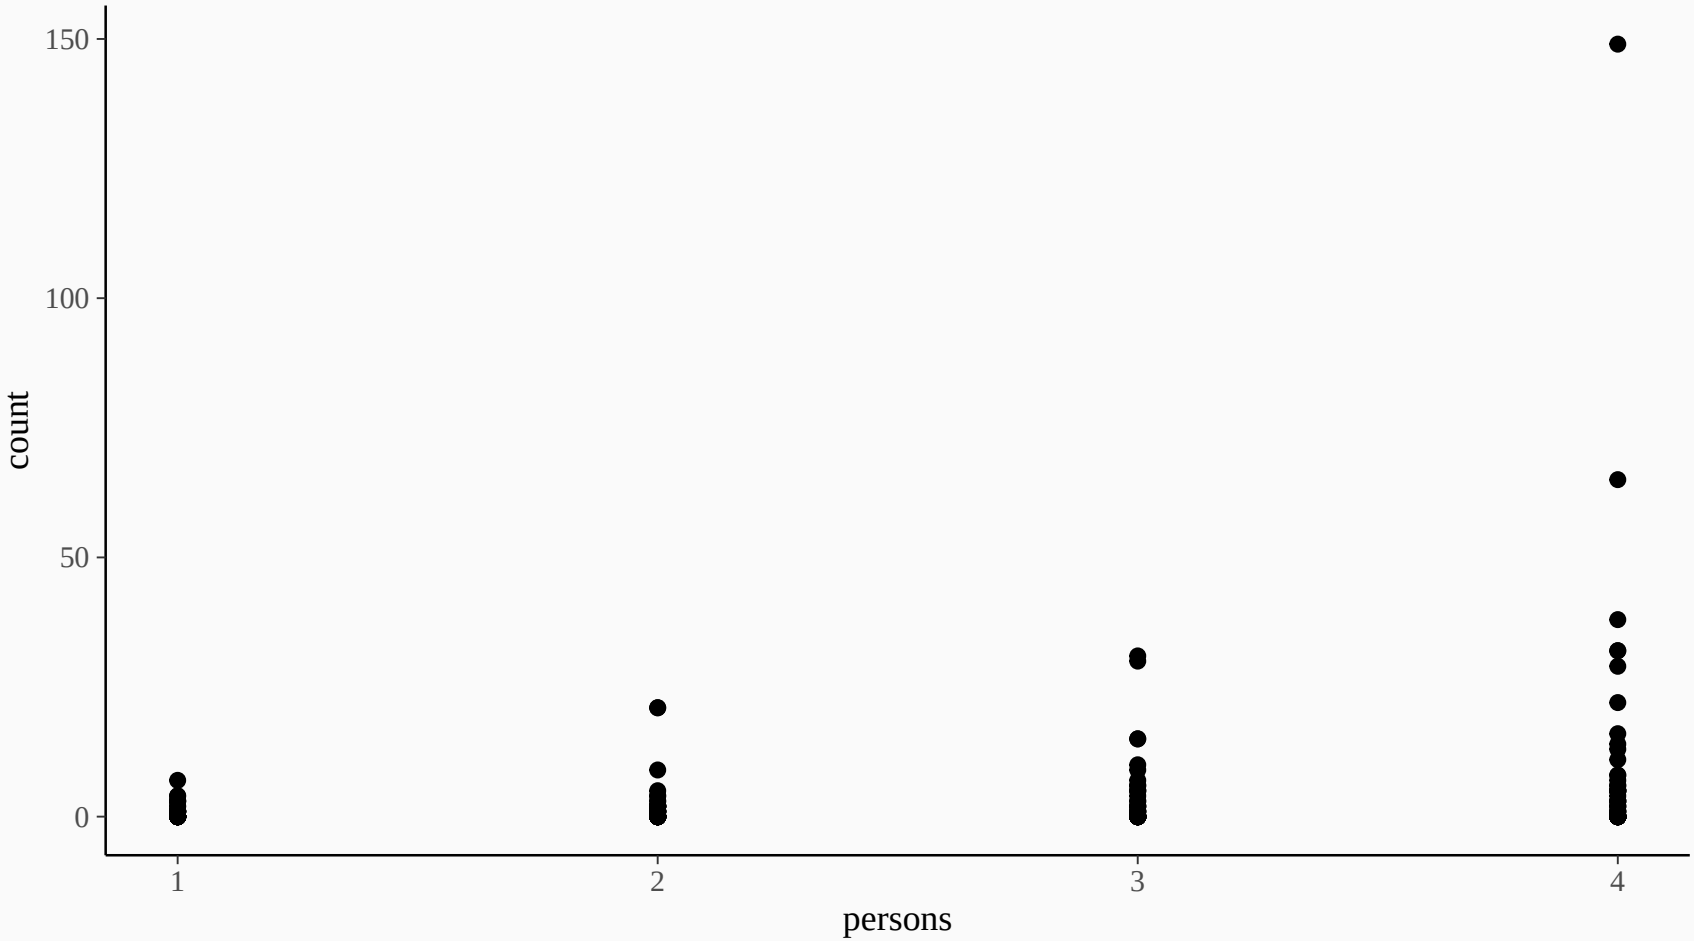

Catching Fish: Modeling

Likelihood:

$$p(y_n \mid \theta) = \text{Poisson}(\exp(\theta_1 + \theta_2 x_n))$$

Example weakly informative priors:

$$p(\theta_1) = \text{normal}(0, 3), \quad p(\theta_2) = \text{normal}(0, 1)$$

Flat ("uninformative") priors:

$$\theta_1 \propto 1, \quad \theta_2 \propto 1$$

Joint model:

$$p(y, \theta) = \left(\prod_n^N p(y_n \mid \theta) \right) p(\theta_1)\, p(\theta_2)$$

The Innocent Marginal Likelihood

$$p(y) = \int p(y \mid \theta)\, p(\theta)\, d\theta$$

Expectations to summarize distributions

(Almost) all we care about are expectations

Expectation of some function f over the posterior $p(\theta \mid y)$:

$$\mathbb{E}_{\theta|y}(f) = \int f(\theta)\, p(\theta \mid y)\, \mathrm{d}\theta$$

Things that are (behave like) expectations:

- Mean
- Variance / standard deviation
- Median
- Other quantiles

Monte-Carlo Estimator

Samples can be used to approximate expectations

Having obtained S random samples $\{\theta^{(s)}\}$ from $p(\theta \mid y)$:

$$\frac{1}{S}\sum_{s=1}^{S} f(\theta^{(s)}) \sim \text{Normal}\left(\mathbb{E}_{\theta|y}(f), \sqrt{\frac{\text{Var}_{\theta|y}(f)}{\text{S}}}\right)$$

Propagation of Uncertainty

Uncertainty can be propagated easily using samples

Suppose, we are interesting in the posterior of $\delta = |\theta_1 - \theta_2|$

Just evaluate $\delta^{(s)} = |\theta_1^{(s)} - \theta_2^{(s)}|$ per sample:

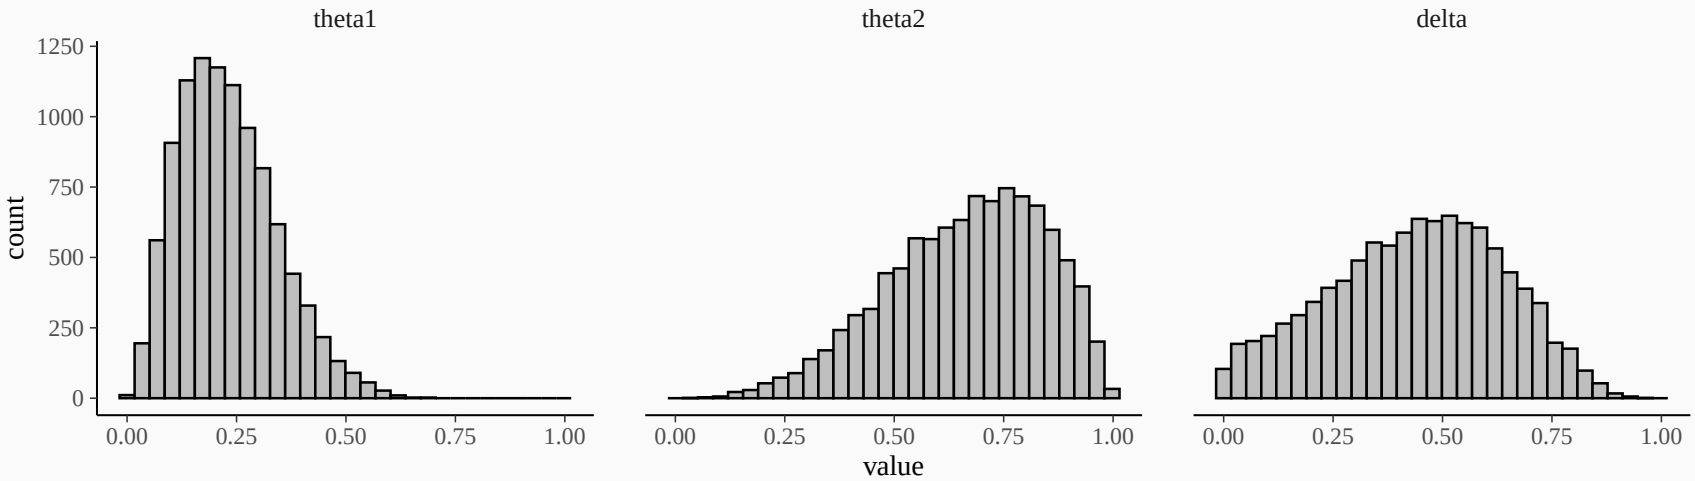

Markov-Chain Monte-Carlo (MCMC) Sampling

We can't simply draw samples from the posterior

A Markov Chain is a sequence of values where the value at position s is based only on the former value at position $s-1$:

$$\theta^{(1)} \rightarrow \theta^{(2)} \rightarrow \theta^{(3)} \rightarrow ... \rightarrow \theta^{(S)}$$

If the *transition distribution* is set up correctly, the values $\theta^{(1)}, ..., \theta^{(S)}$ will represent (dependent) samples from the posterior

MCMC Sampling: A Single Chain (10 Iterations)

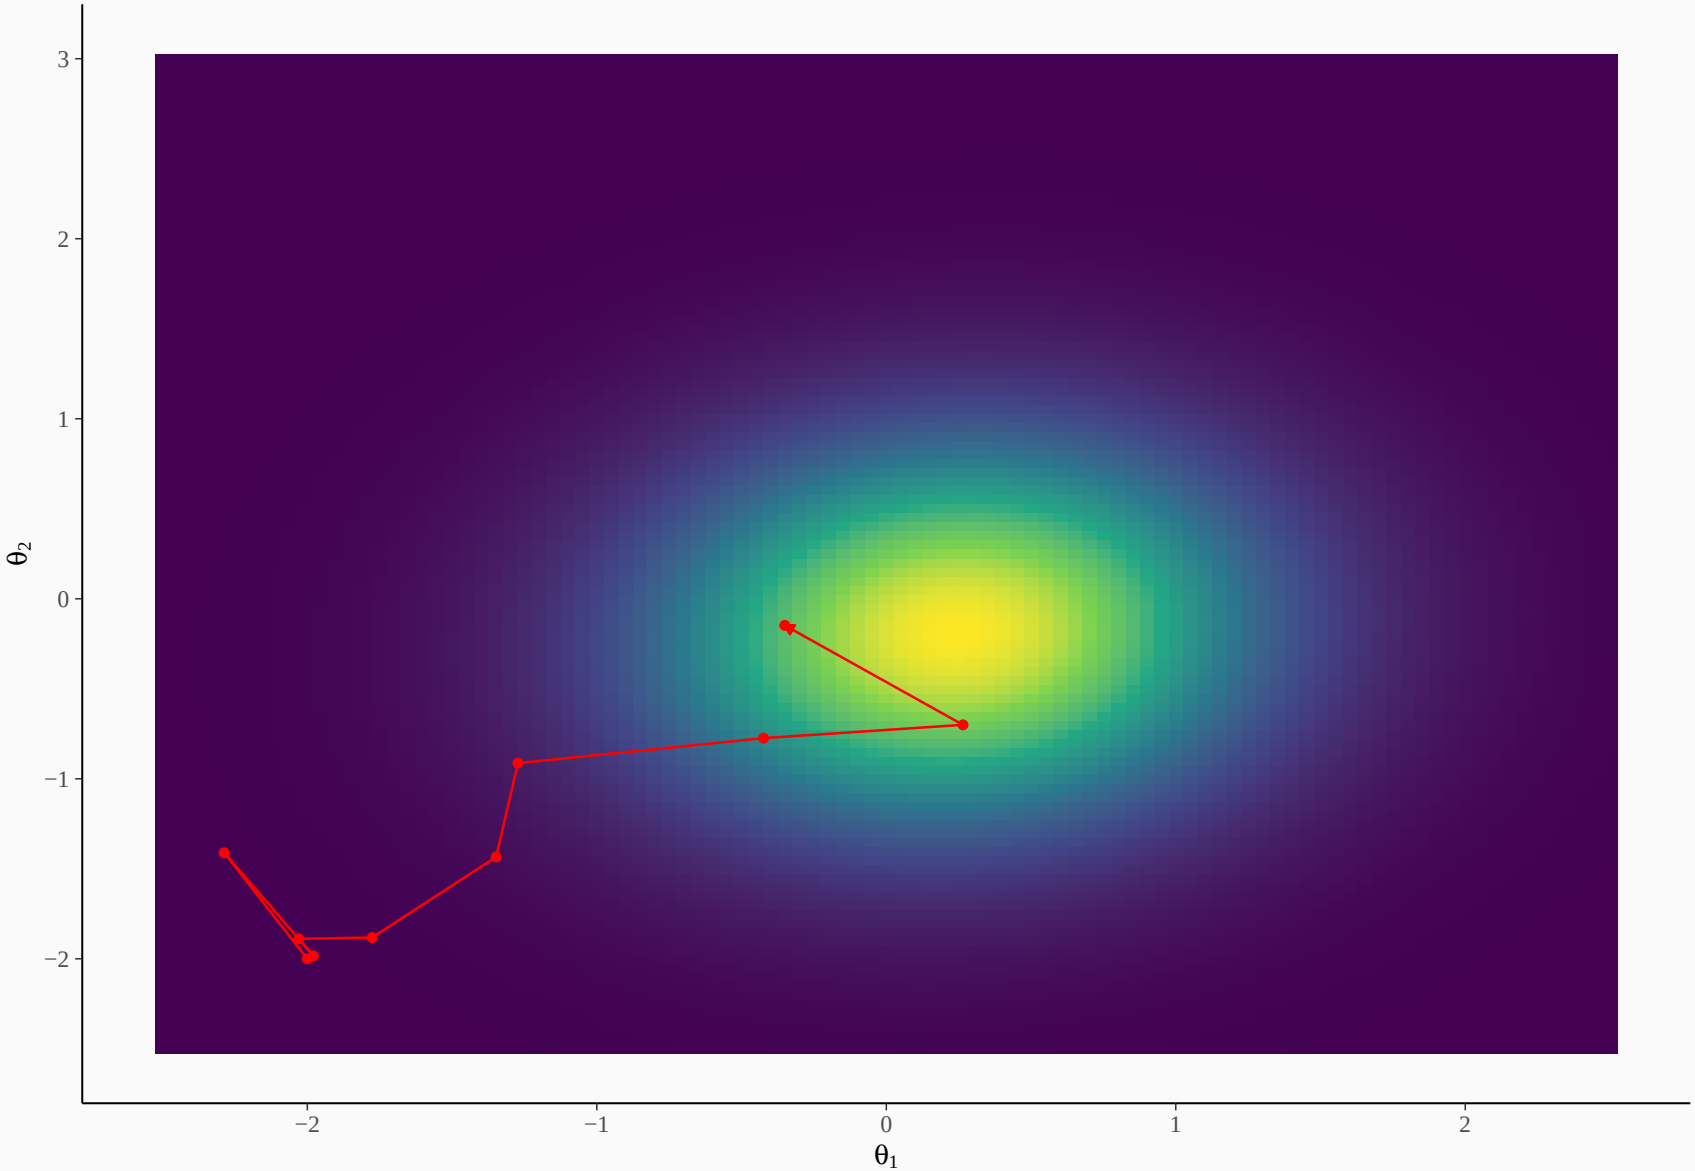

MCMC Sampling: A Single Chain (50 Iterations)

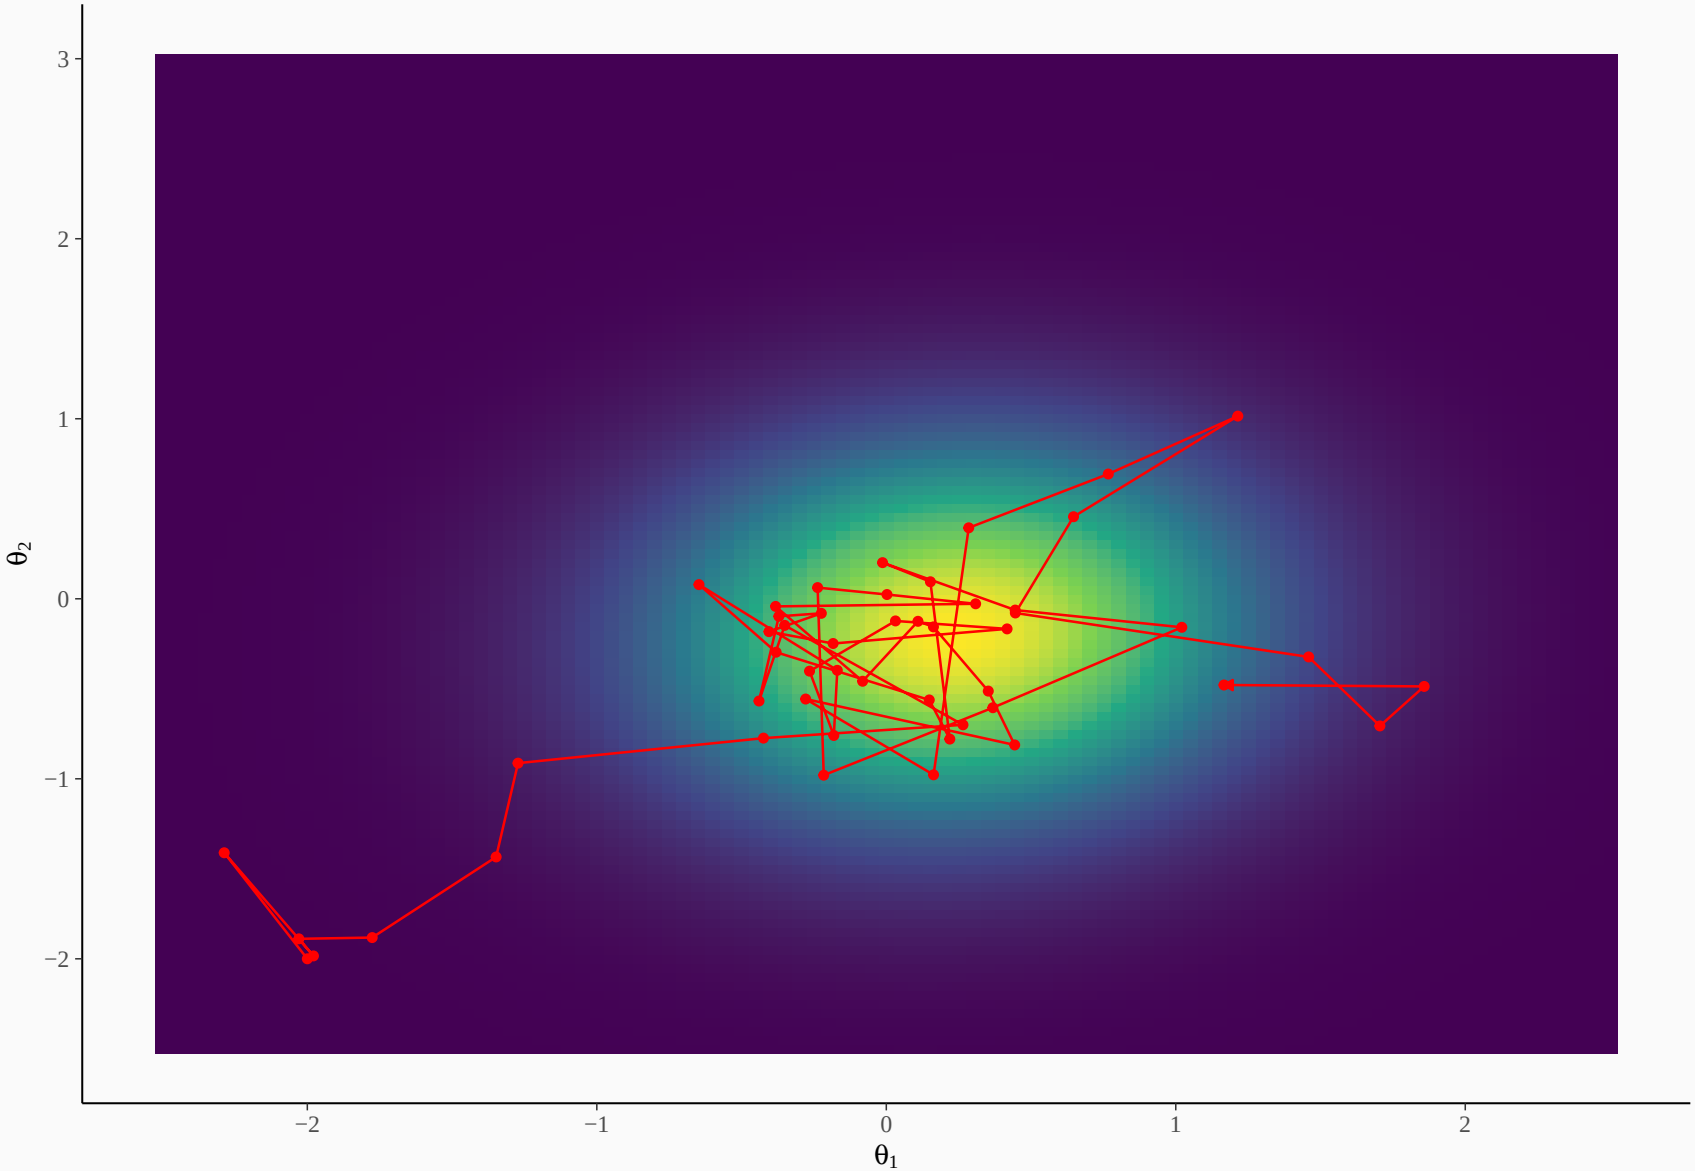

MCMC Sampling: A Single Chain (1000 Iterations)

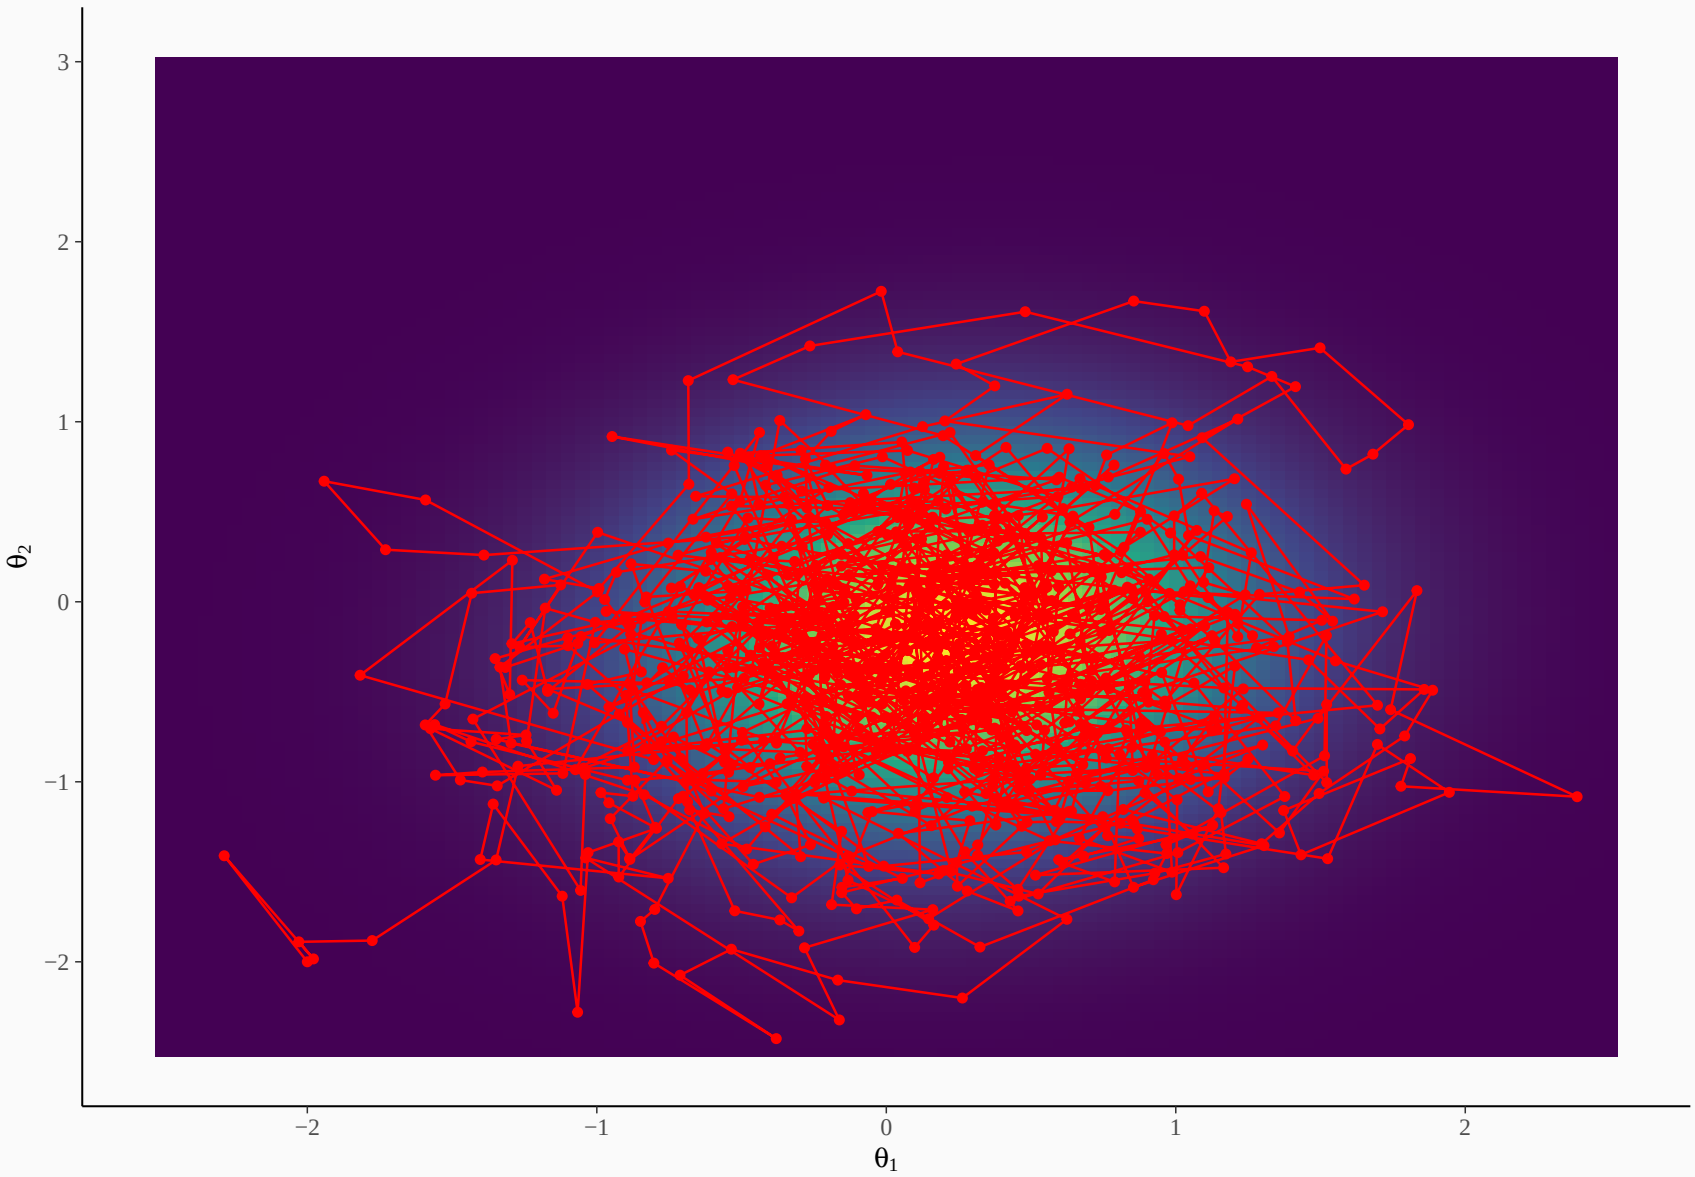

Markov-Chain Monto-Carlo Estimator

MCMC samples can be used to approximate expectations

Assume well-behaved MCMC samples $\{\theta^{(s)}\}$ over $p(\theta \mid y)$:

$$\frac{1}{S}\sum_{s=1}^{S} f(\theta^{(s)}) \sim \text{Normal}\left(\mathbb{E}_{\theta|y}(f), \sqrt{\frac{\text{Var}_{\theta|y}(f)}{\text{ESS}}}\right)$$

Trace Plots: Visualizing a Single Chain

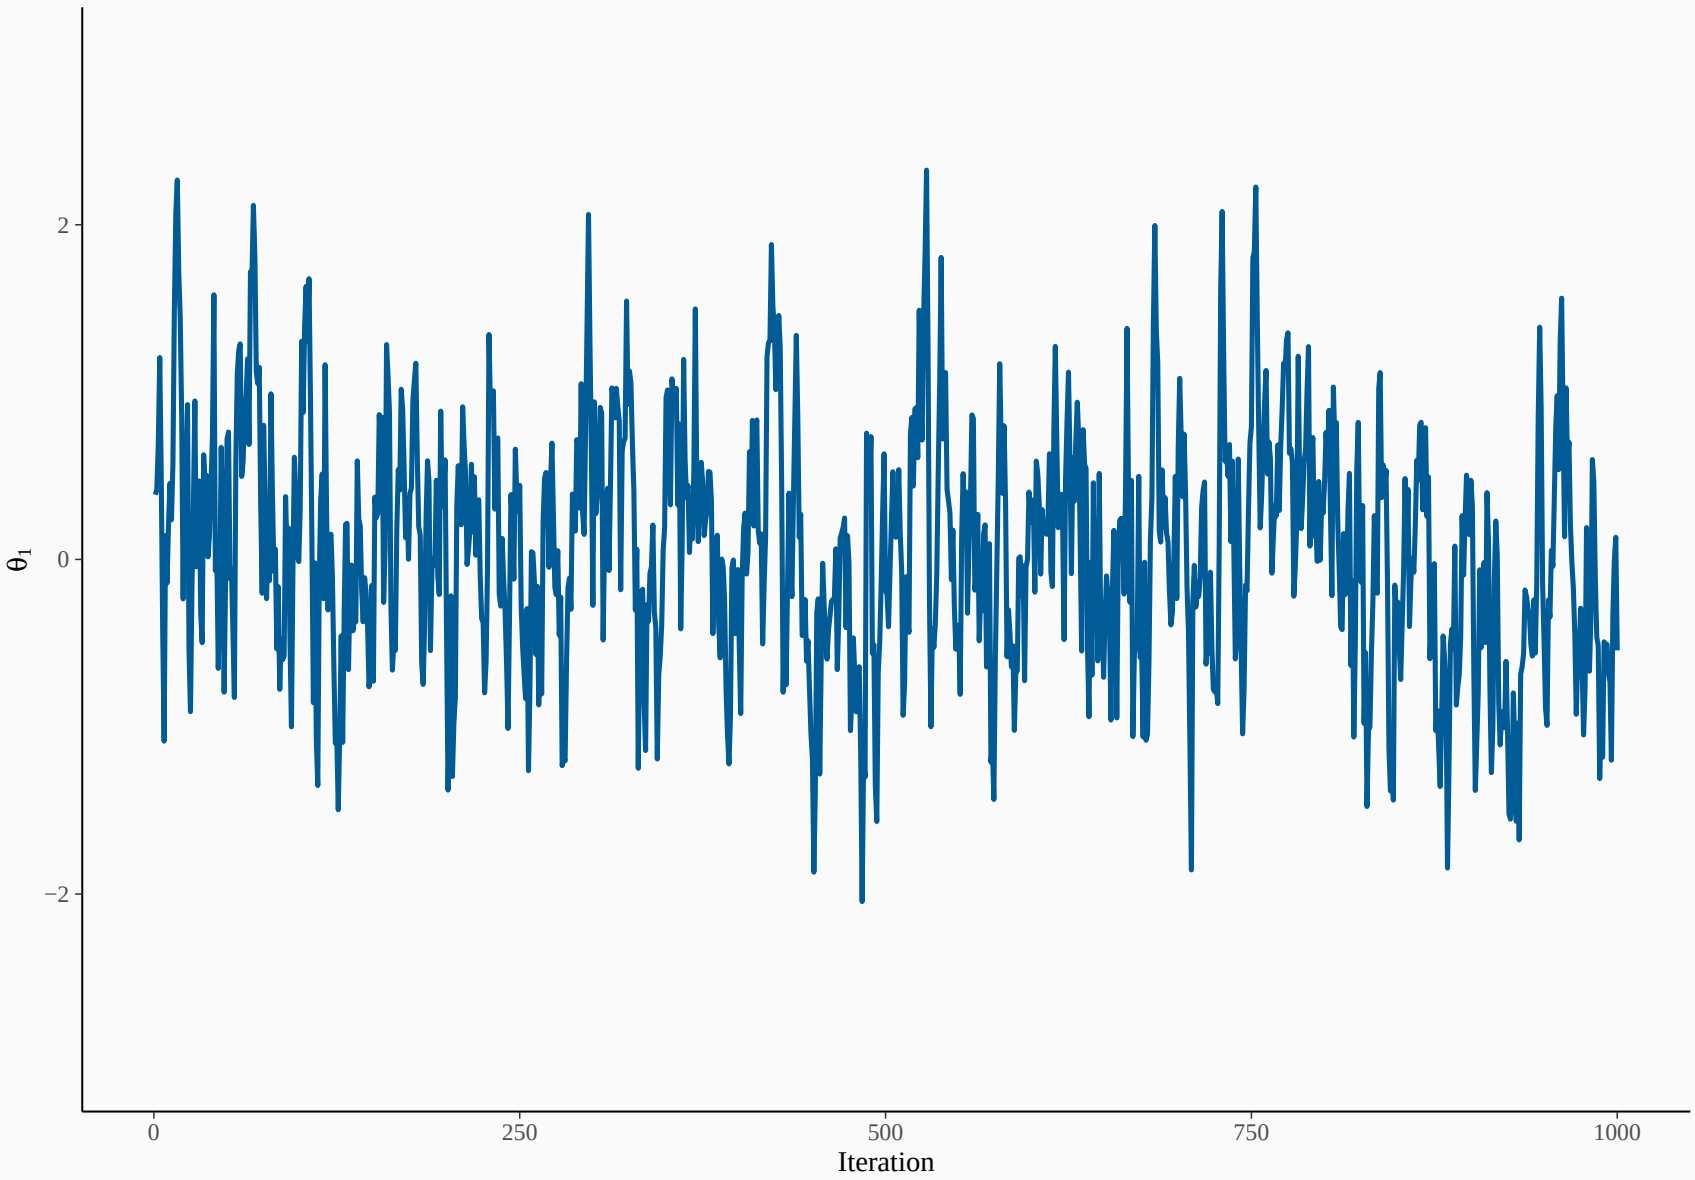

Trace Plots: Visualizing Multiple Chains

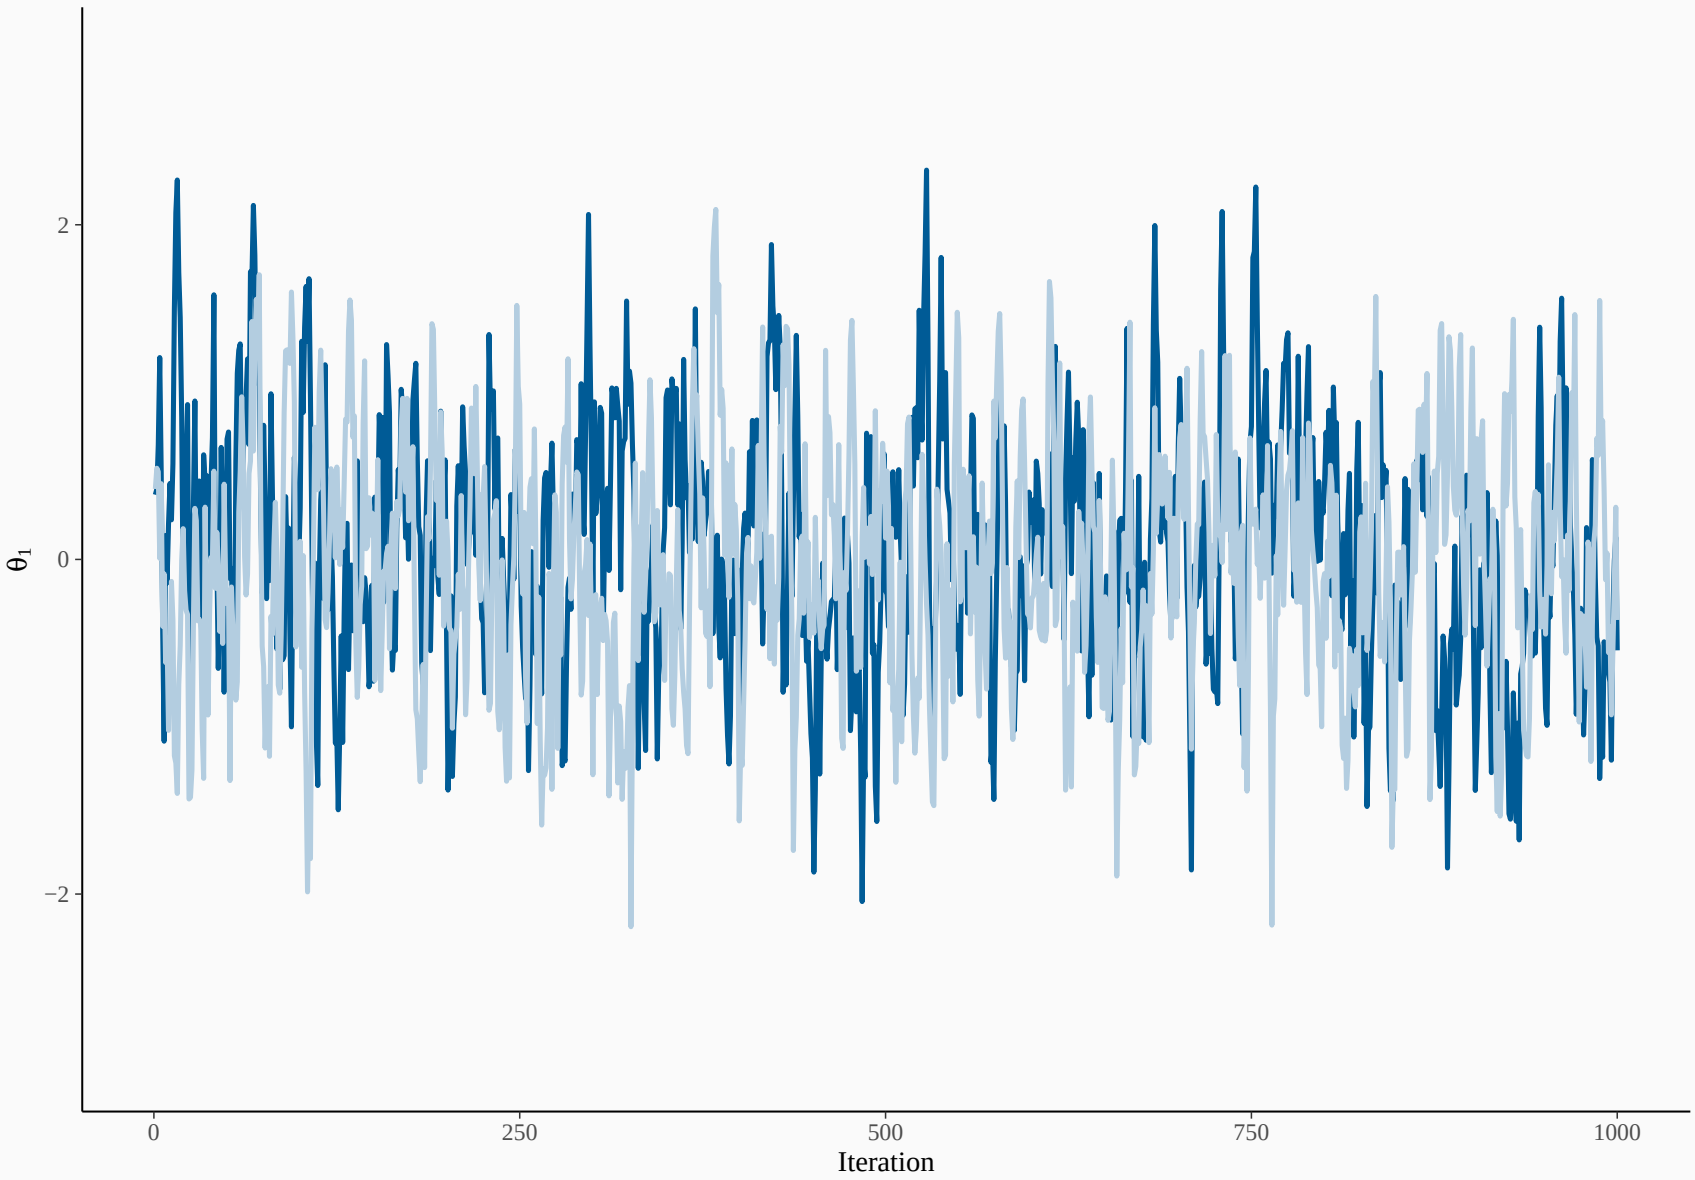

Probabilistic Programming Languages

Catching Fish: Results

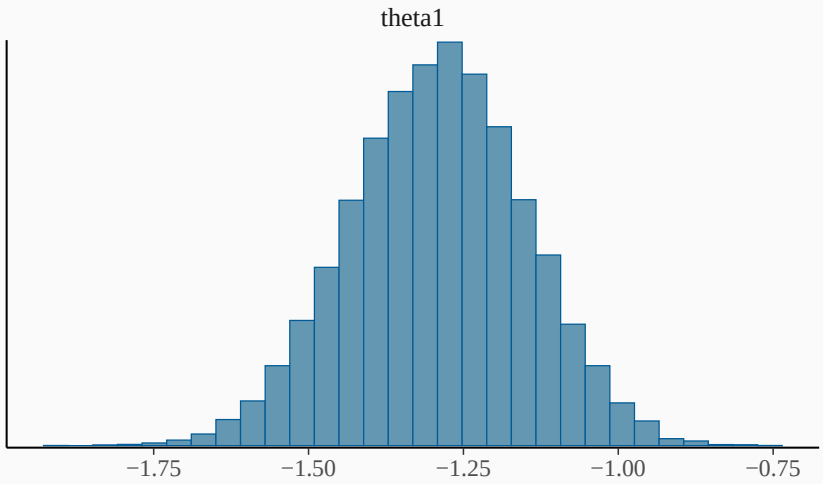

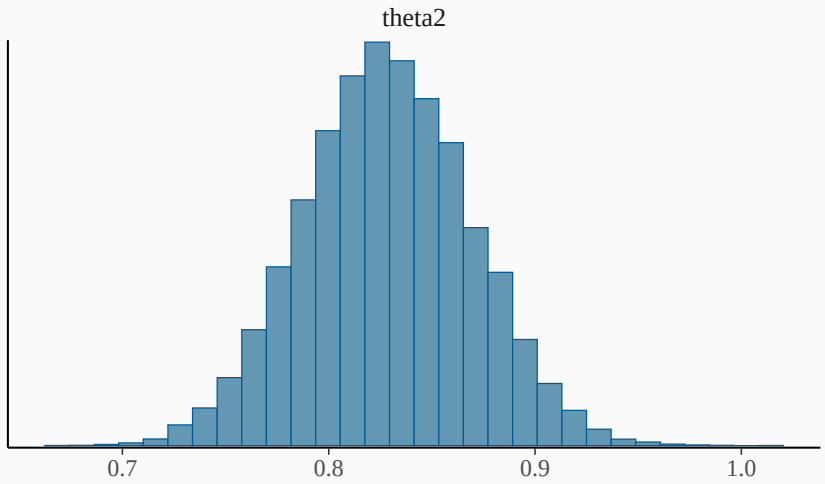

variable	mean	median	sd	q5	q95
theta1	-1.29	-1.29	0.14	-1.53	-1.06
theta2	0.83	0.83	0.04	0.76	0.90

Catching Fish: Expected posterior predictions

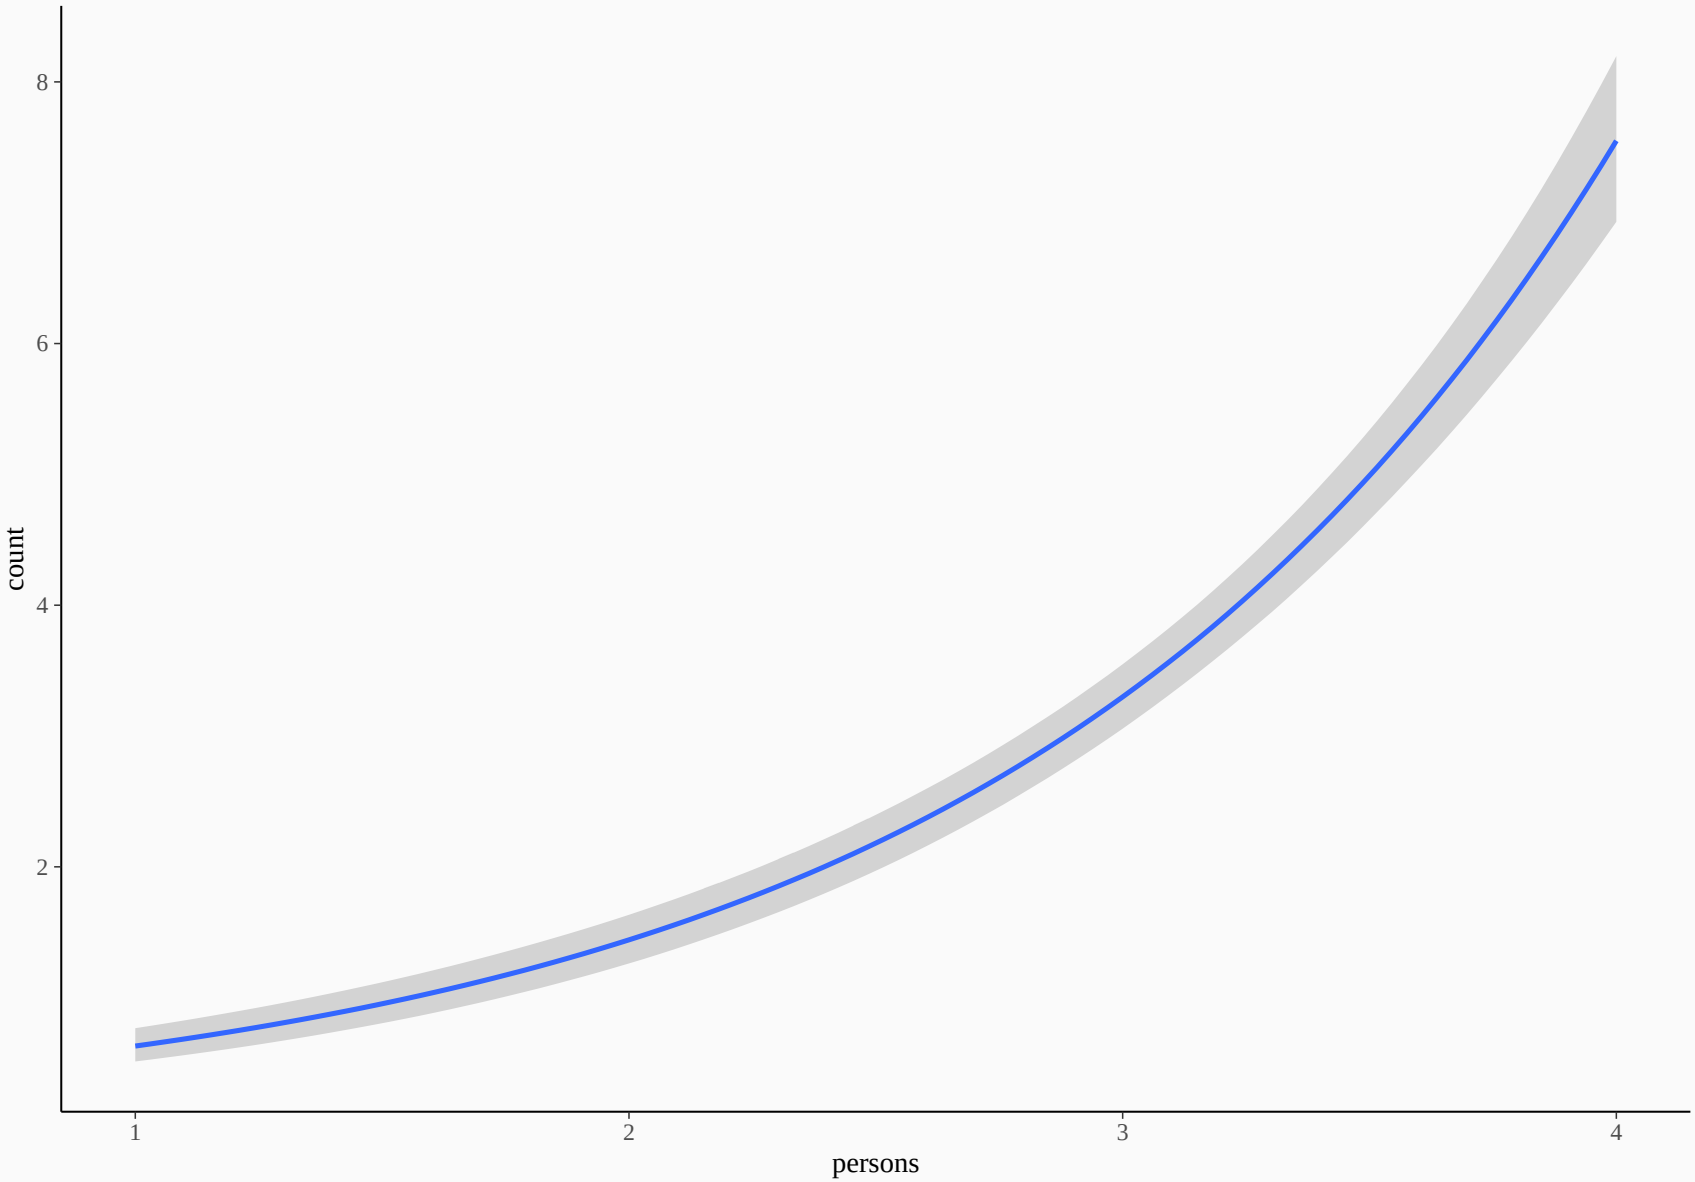

Catching Fish: Posterior prediction

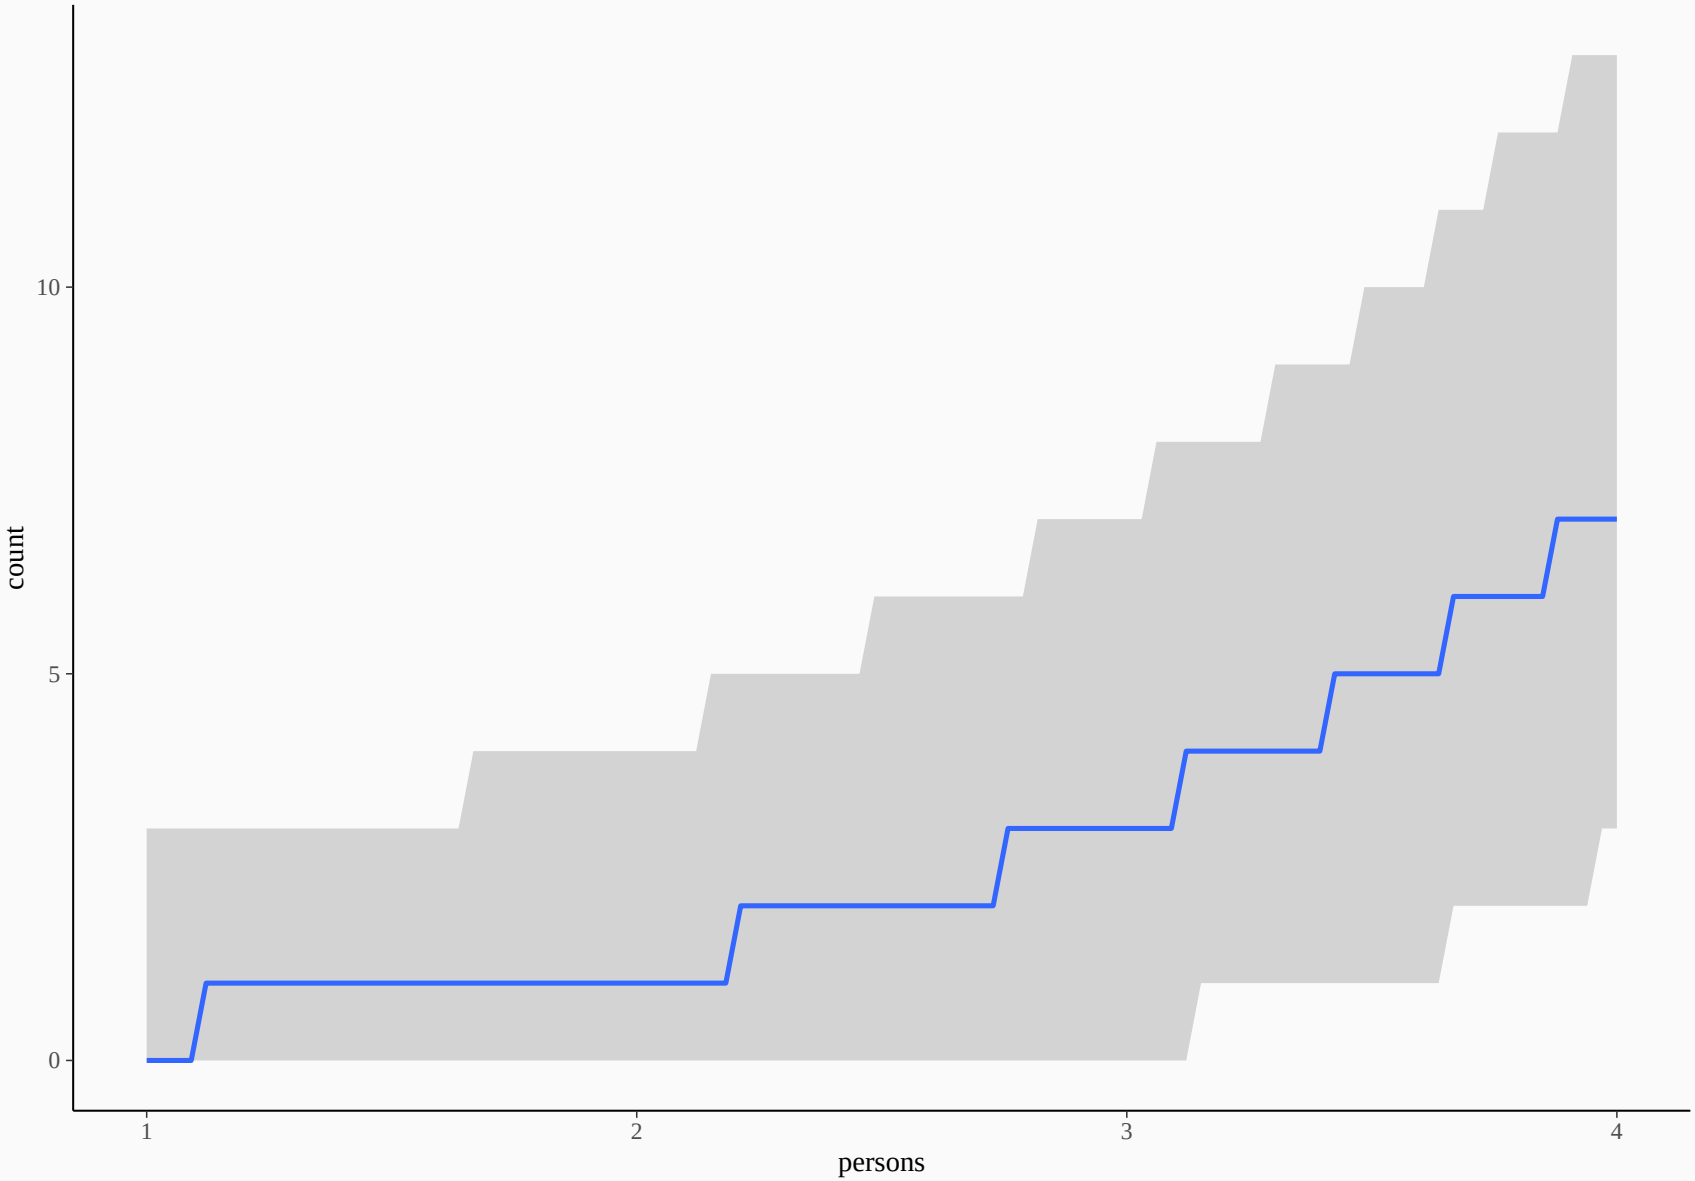

What I like and don't like about Bayesian inference

What I like:

- Intuitive approach to expressing uncertainty
- Ability to incorporate prior information
- A lot of modeling flexibility
- Joint posterior distribution of parameters
- Easy propagation of uncertainty

What I don't like:

- Slow Speed of model estimation

A brief look into my own research

What actually is a Bayesian Model?

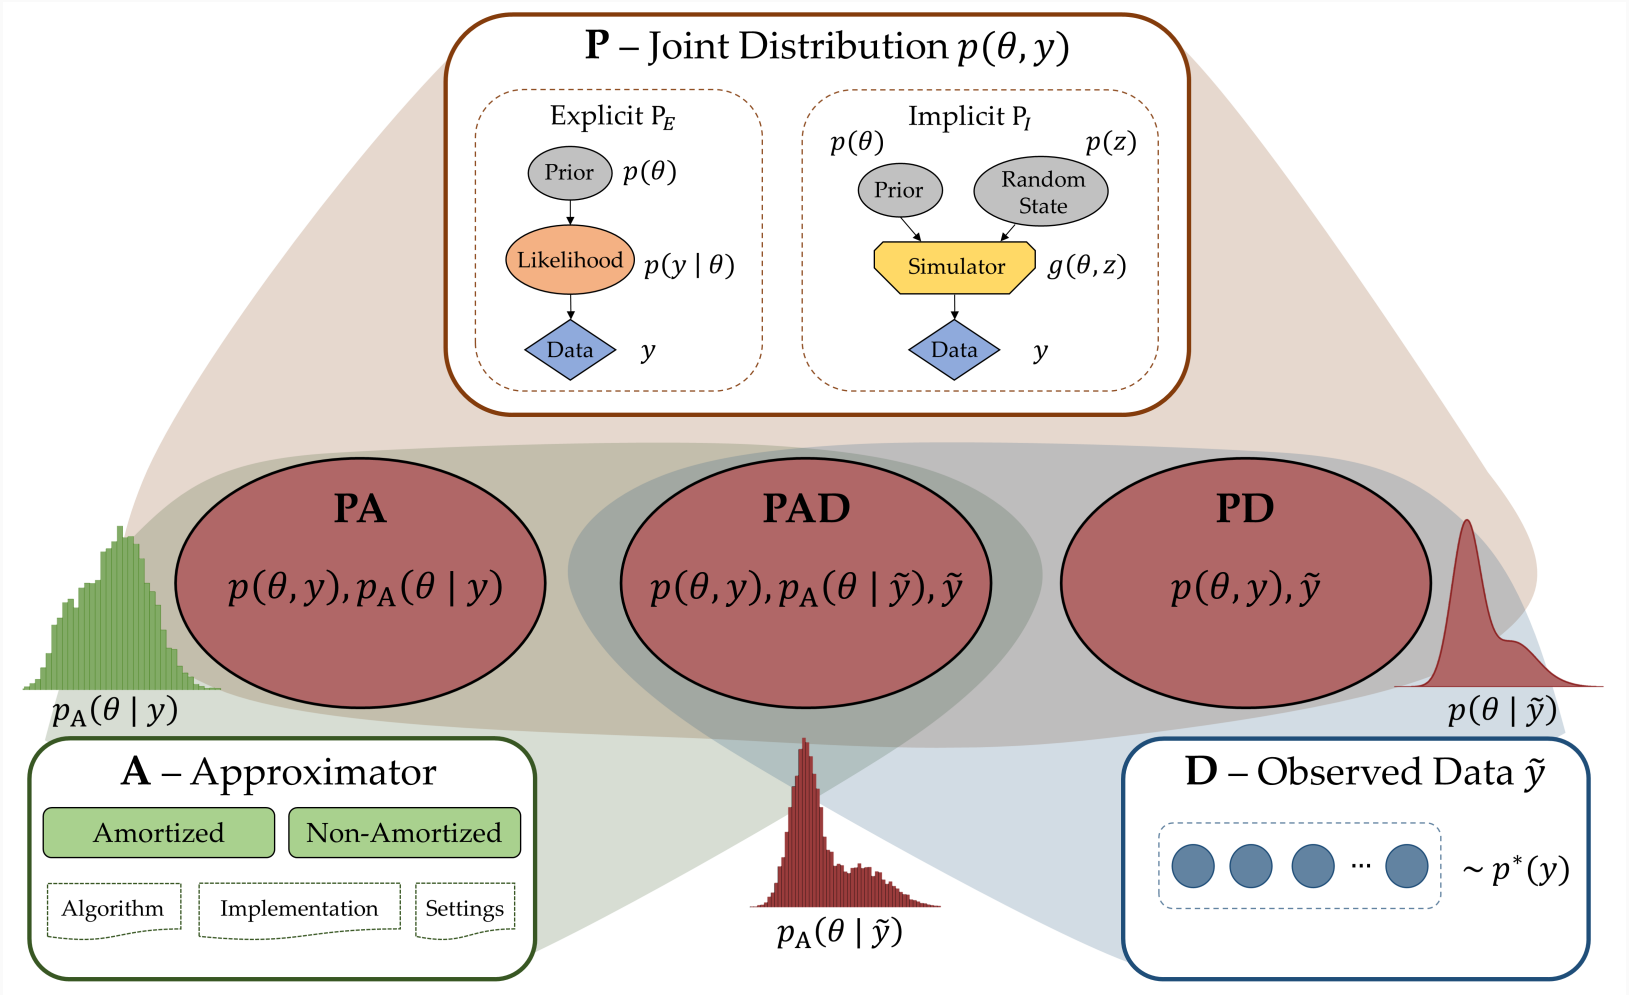

Uncertainty of Uncertainty (Meta-Uncertainty)

How can we combine Bayesian and frequentist uncertainties?

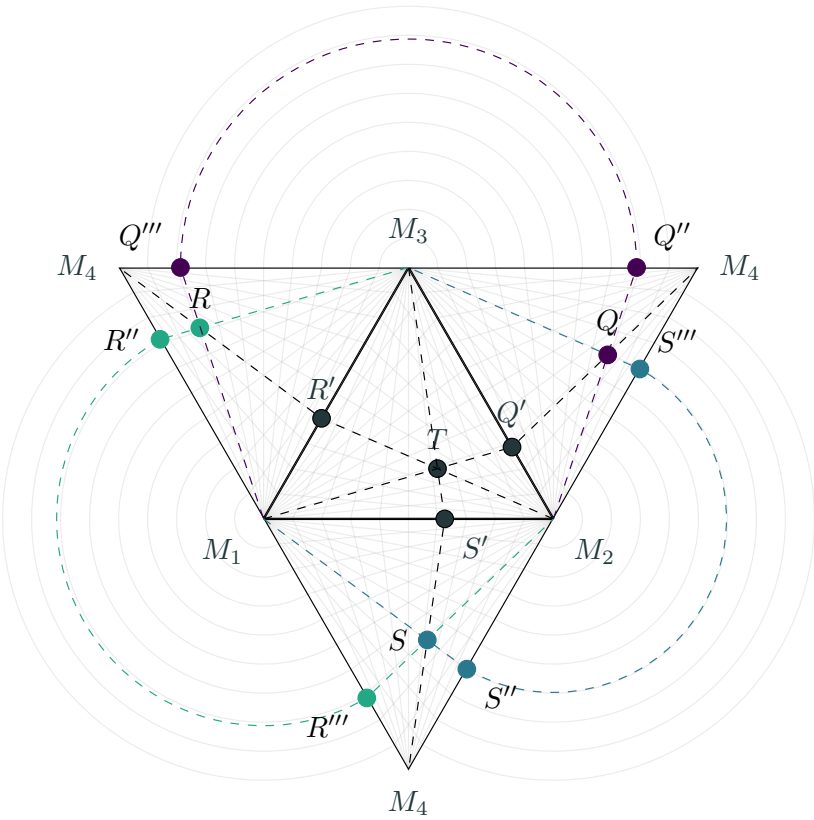

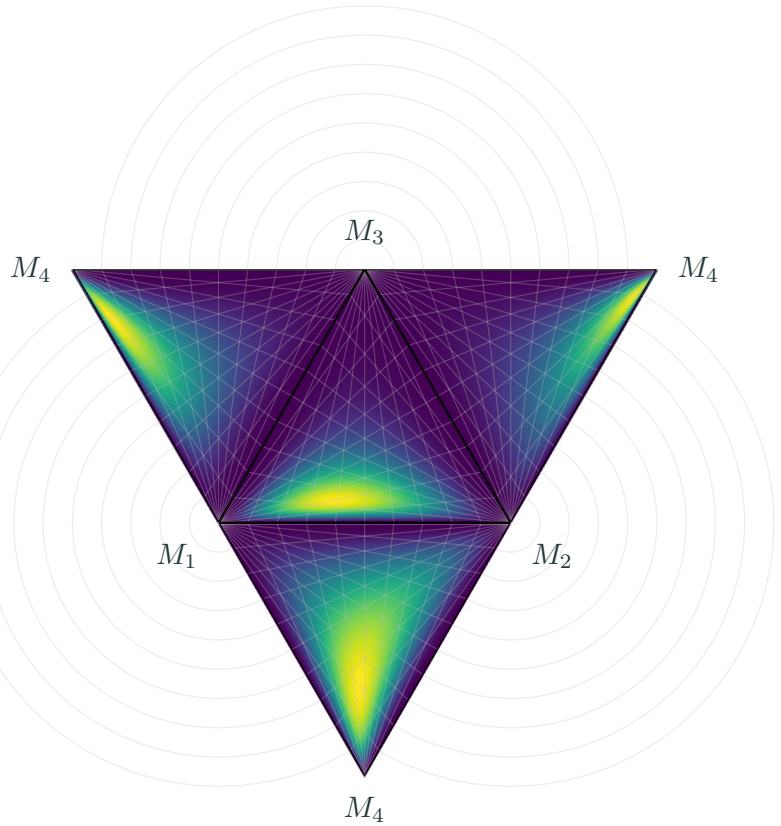

Non-Amortized (Standard) Inference

How can we improved the standard inference setting?

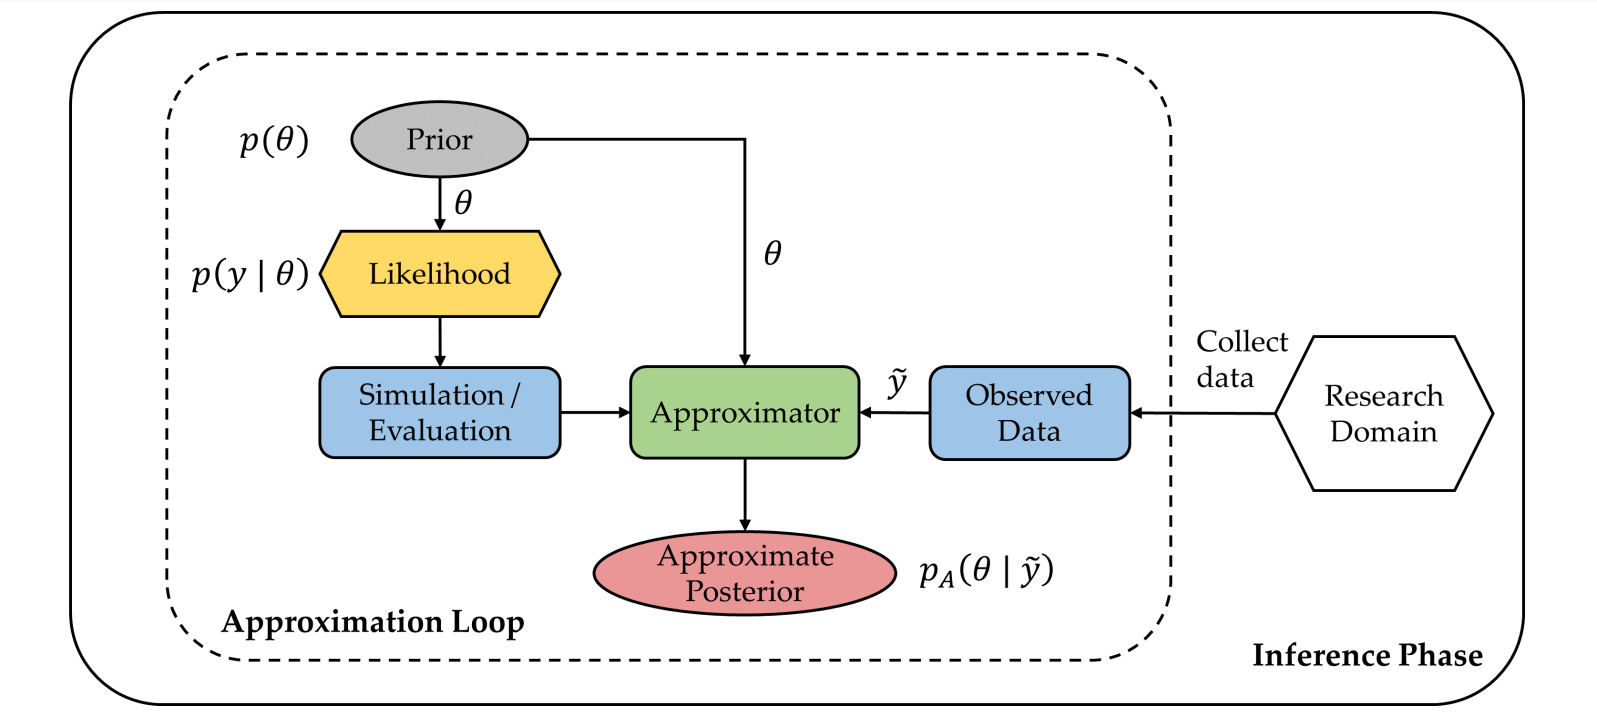

Amortized Inference

How far can we scale amortized inference?

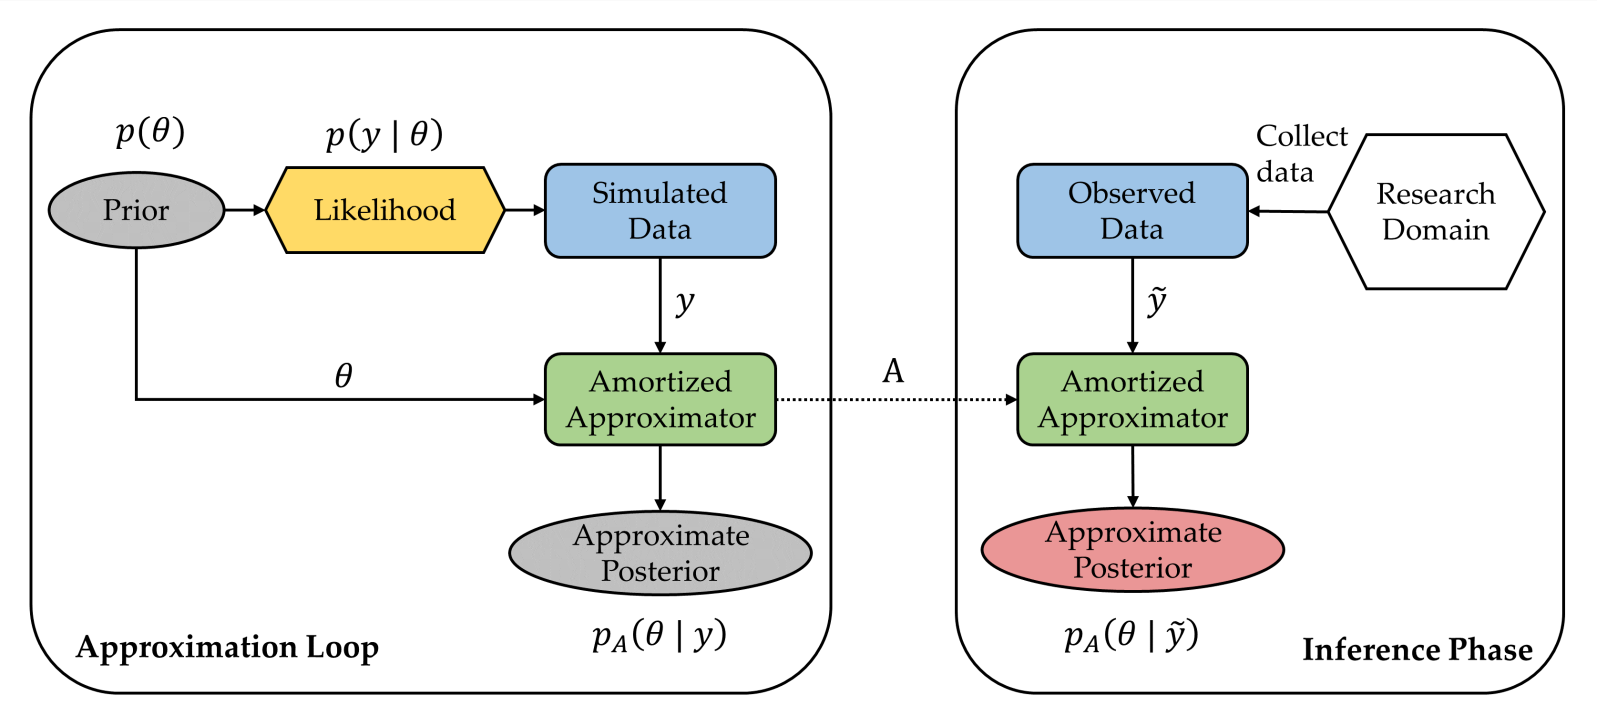

More about me and my research

- Website: https://paul-buerkner.github.io/
- Email: `paul.buerkner@gmail.com`
- Twitter: `@paulbuerkner`
- Mastodon: `fosstodon.org/@paul_buerkner`