

Bayesian Methodology for Particle Physics with pyhf

CHEP23, 08.05.2023

Matthew Feickert, Lukas Heinrich, Malin Horstmann (malin.horstmann@tum.de)

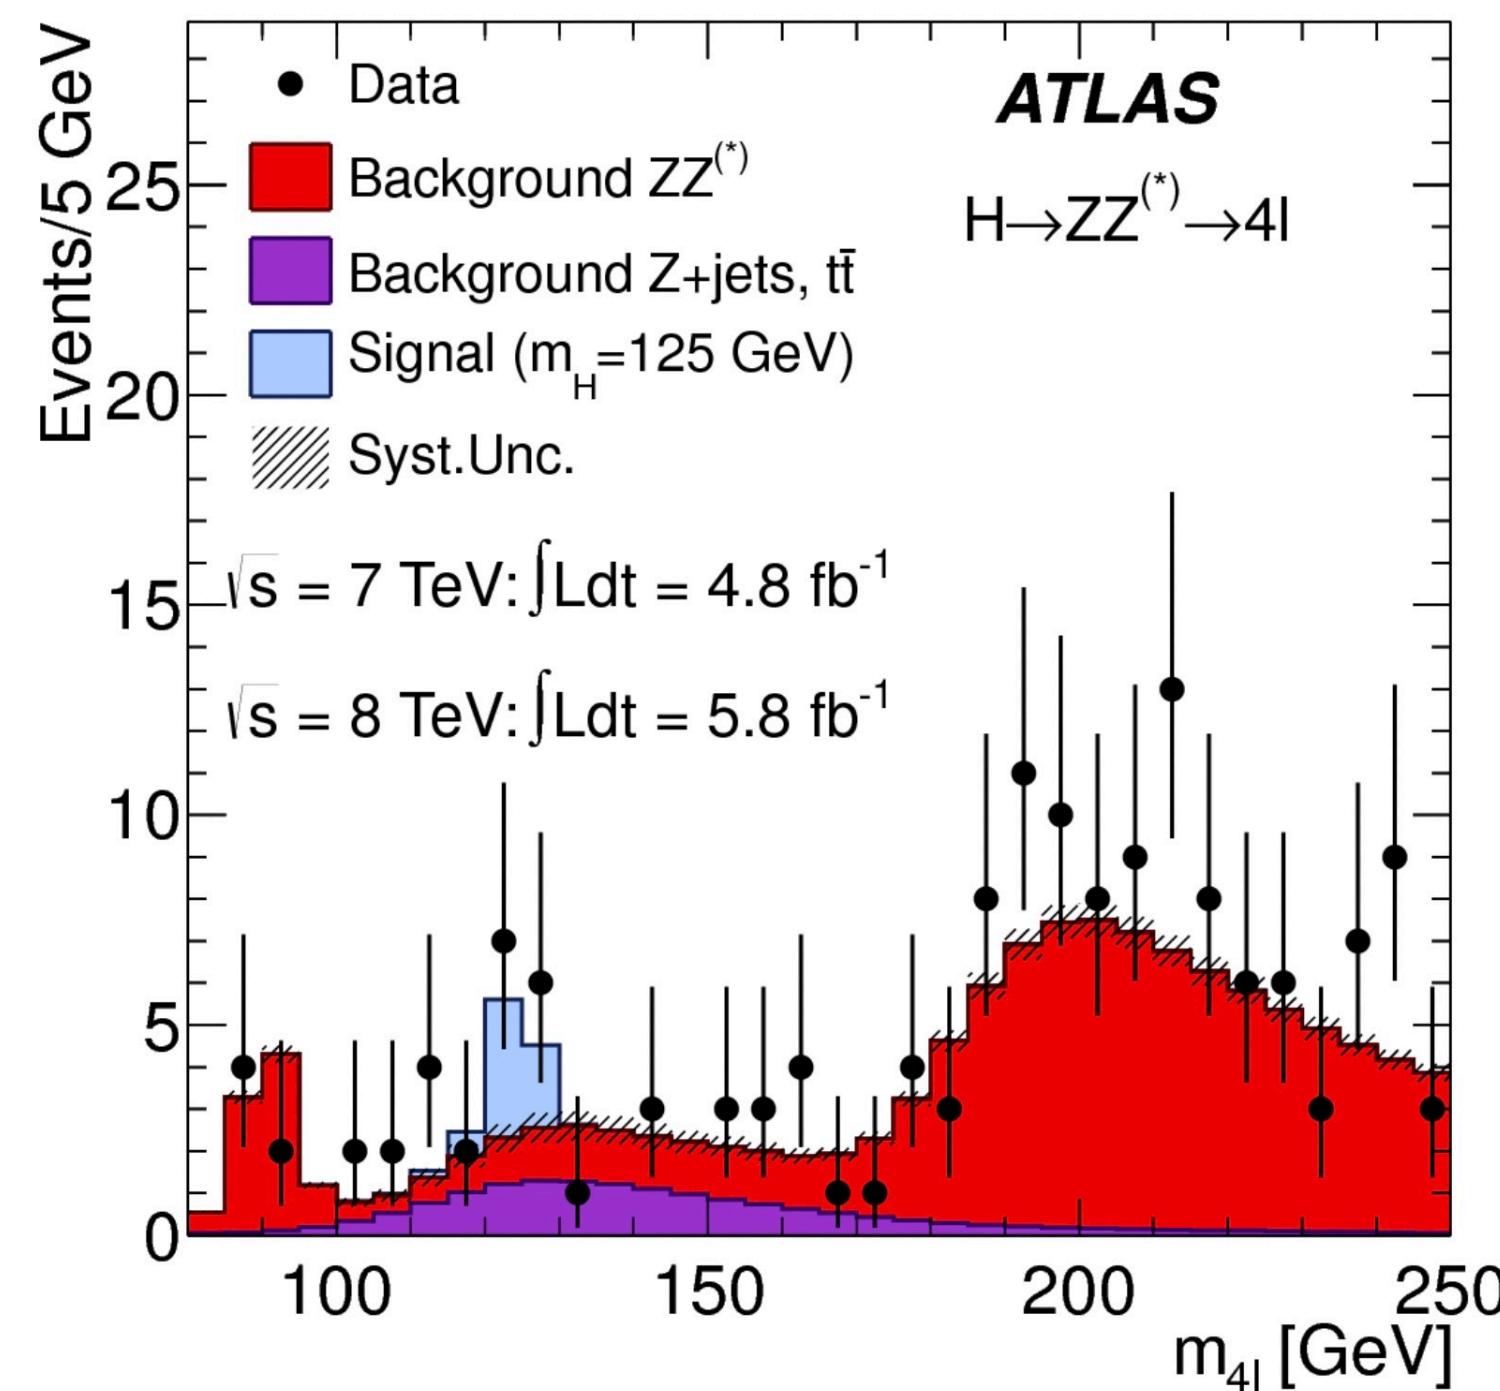


MAX-PLANCK-INSTITUT
FÜR PHYSIK



HistFactory

- We need a data-generating model for inferring parameters of nature
- HistFactory: Toolbox for creating these models for binned analyses



HistFactory

- We need a data-generating model for inferring parameters of nature
- HistFactory: Toolbox for creating these models for binned analyses
- Probability model:

$$p(\vec{n}, \vec{a} | \vec{\eta}, \vec{\chi}) = \prod_{c \in \text{Channels}} \prod_{b \in \text{Bins}} \text{Pois} (n_{cb} | \nu_{cb}(\vec{\eta}, \vec{\chi})) \prod_{\chi \in \vec{\chi}} c_\chi (a_\chi | \chi)$$

Diagram illustrating the components of the probability model:

- $\vec{\eta}$: Unconstrained Parameters (incl. POI)
- $\vec{\chi}$: Nuisance parameters
- \vec{n} : Bin counts
- \vec{a} : Auxiliary data
- Poisson-distributed bin counts
- Constraint terms

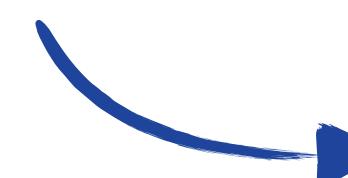
Annotations with arrows:

- A curved arrow points from $\vec{\eta}$ to the term $\nu_{cb}(\vec{\eta}, \vec{\chi})$.
- A curved arrow points from $\vec{\chi}$ to the term $c_\chi (a_\chi | \chi)$.
- A curved arrow points from \vec{n} to the term n_{cb} .
- A curved arrow points from \vec{a} to the term a_χ .
- A blue arrow points from the text "Poisson-distributed bin counts" to the term n_{cb} .
- A red arrow points from the text "Constraint terms" to the term $c_\chi (a_\chi | \chi)$.

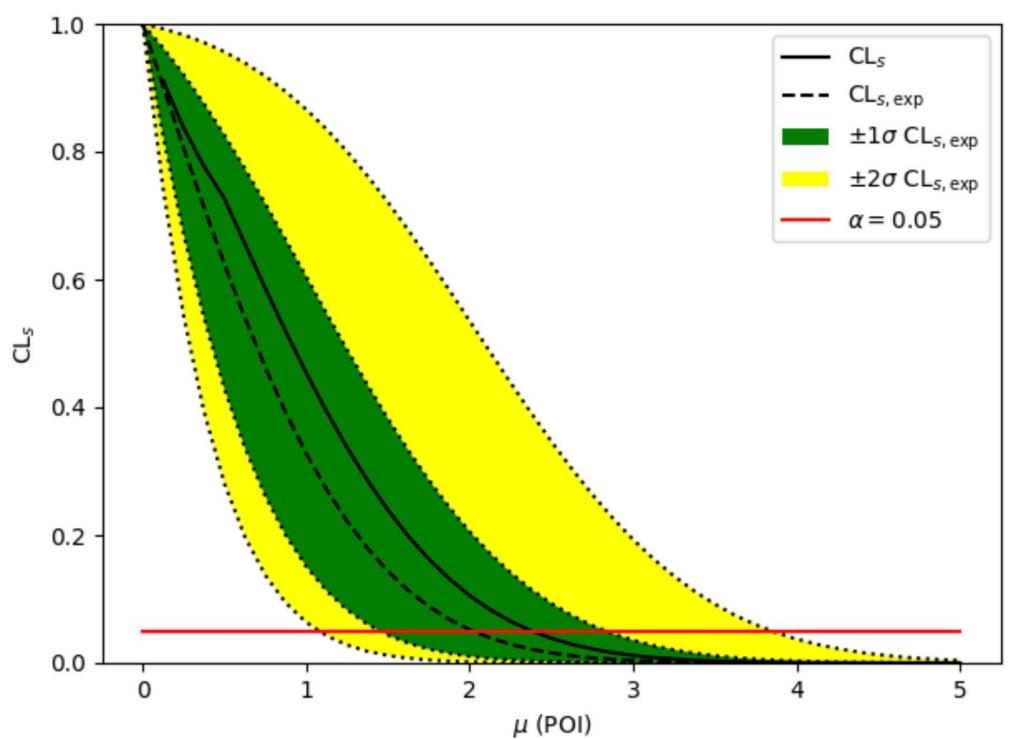
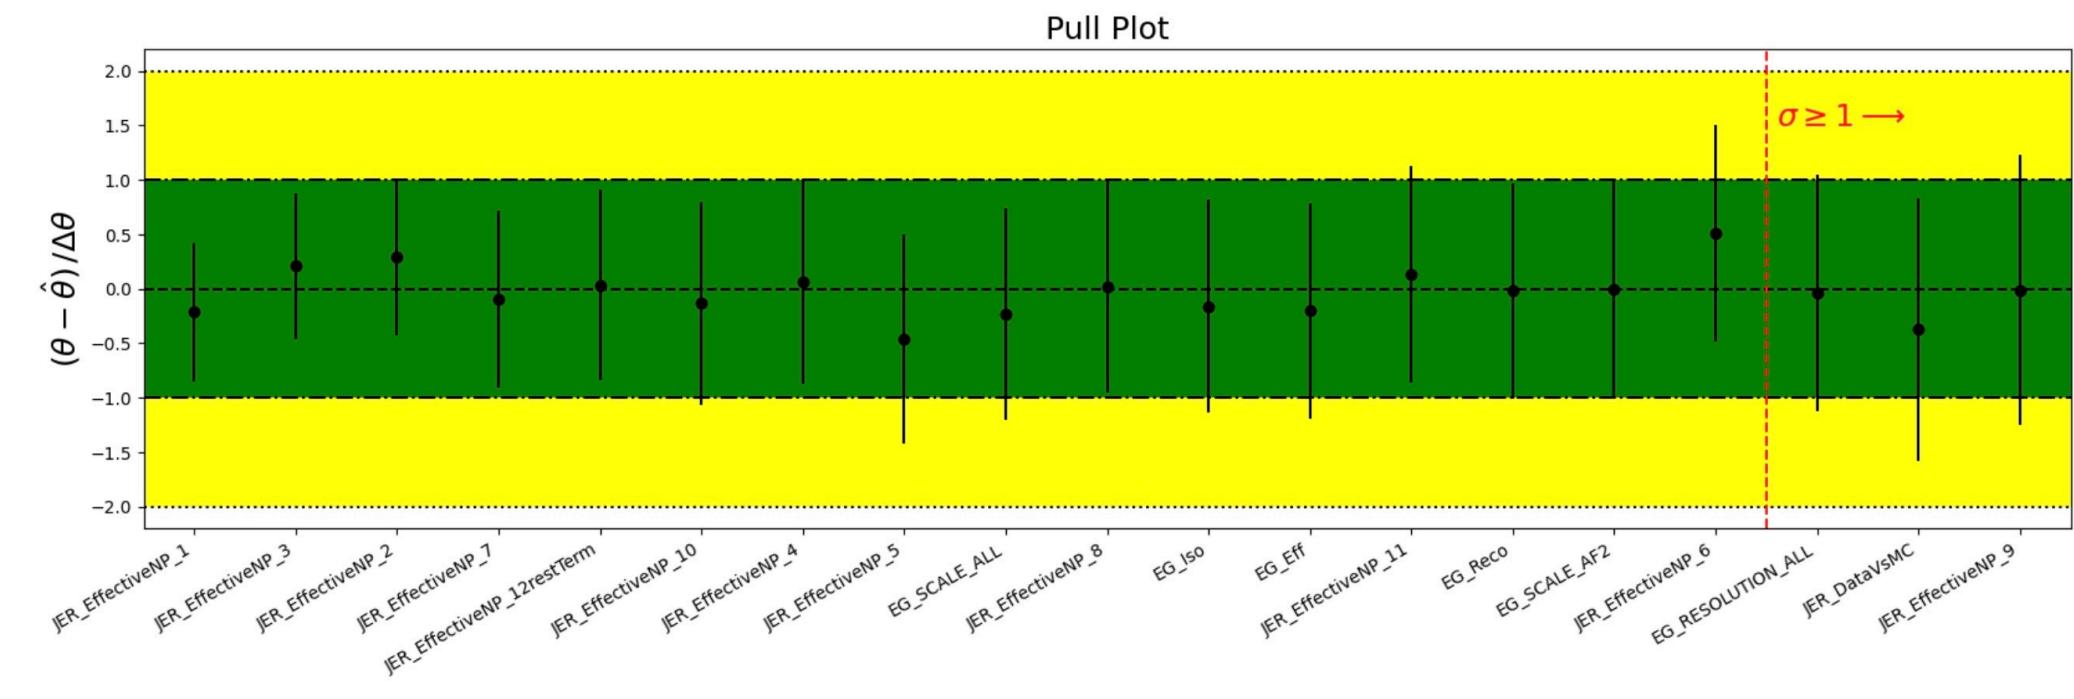
pyhf



- A pure Python implementation of the HistFactory statistical model, including tools for Frequentist inference
- Supports auto-differentiation via different computational backends (e.g. jax)



Before this work:



After this work: Priors and Posteriors

Bayesian Inference

- Updating a prior belief - after taking observations \vec{n}, \vec{a} into consideration - via Bayes theorem:

$$p(\vec{\eta}, \vec{\chi} | \vec{n}, \vec{a}) \approx p(\vec{n}, \vec{a} | \vec{\eta}, \vec{\chi}) p(\vec{\eta}) p(\vec{\chi})$$

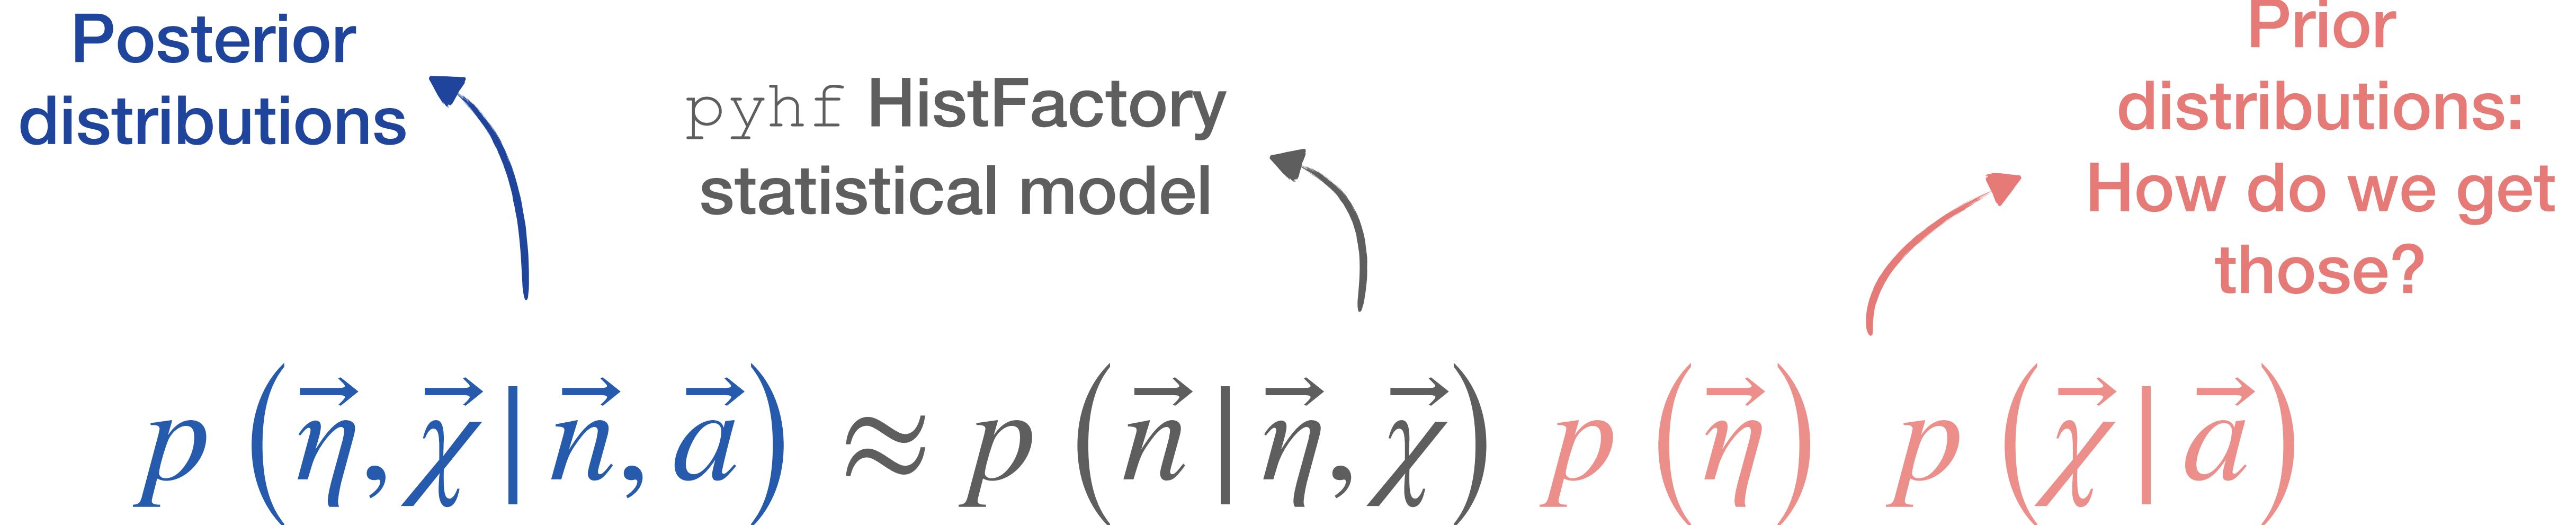
Posterior distributions pyhf HistFactory statistical model Prior distributions

```
graph TD; Eq[p(η̂, χ̂ | n̂, â) ≈ p(n̂, â | η̂, χ̂) p(η̂) p(χ̂)] -- "Blue arrow" --> Posterior[Posterior distributions]; Eq -- "Blue arrow" --> Likelihood[pyhf HistFactory statistical model]; Eq -- "Blue arrow" --> Priors[Prior distributions]; Posterior -- "Blue arrow" --> Eq; Likelihood -- "Grey arrow" --> Eq; Priors -- "Red arrow" --> Eq; Priors -- "Red arrow" --> Eq;
```

Bayesian Inference

- Updating a prior belief - after taking observations \vec{n}, \vec{a} into consideration - via Bayes theorem:

$$p(\vec{\eta}, \vec{\chi} | \vec{n}, \vec{a}) \approx p(\vec{n}, \vec{a} | \vec{\eta}, \vec{\chi}) p(\vec{\eta}) p(\vec{\chi})$$



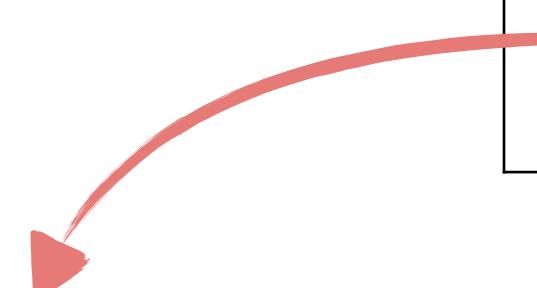
$$p(\vec{\eta}, \vec{\chi} | \vec{n}, \vec{a}) \approx p(\vec{n} | \vec{\eta}, \vec{\chi}) \ p(\vec{\eta}) \ p(\vec{\chi} | \vec{a})$$

Building Prior Distributions

- Turn information from **auxiliary measurements** (means, rates, uncertainties) into prior distributions for $\vec{\chi}$ using **conjugate priors**:

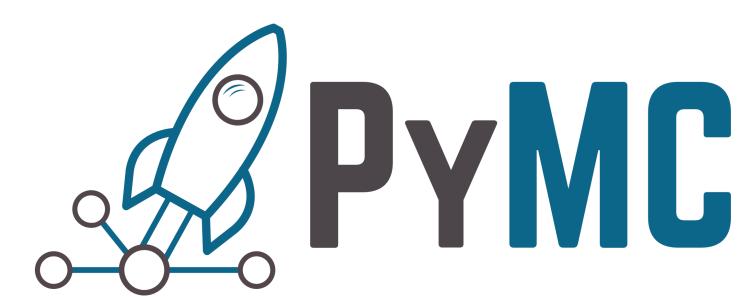
$$p(\vec{\chi} | \vec{a}) \approx p(\vec{a} | \vec{\chi}) \text{ pur}(\vec{\chi})$$

$p(a \chi)$	$\propto \text{Poiss}(a \chi)$	$\propto \mathcal{N}(a \mu = 0, \sigma = 1)$
$p(\chi a)$	$\propto \Gamma(\chi \alpha = a, \beta = a)$	$\propto \mathcal{N}(\chi \mu = 0, \sigma = 1)$



Priors for the nuisance parameters $\vec{\chi}$

pyhf + PyMC



- PyMC: Probabilistic programming Python library for Bayesian analysis based on MCMC methodology
- Already includes a wide range of MCMC techniques, cross-checks, ...
- PyMC allows for the implementation of external models
- Example inference code:

```
import pyhf
import pymc

with pyhf.infer.bayes(model, prior, data):
    posterior = pymc.sample(10_000)
    posterior_predictive = pymc.sample_posterior_predictive(posterior)
    prior = pymc.sample_prior_predictive(10_000)
```

Example Bayesian Workflow

Model

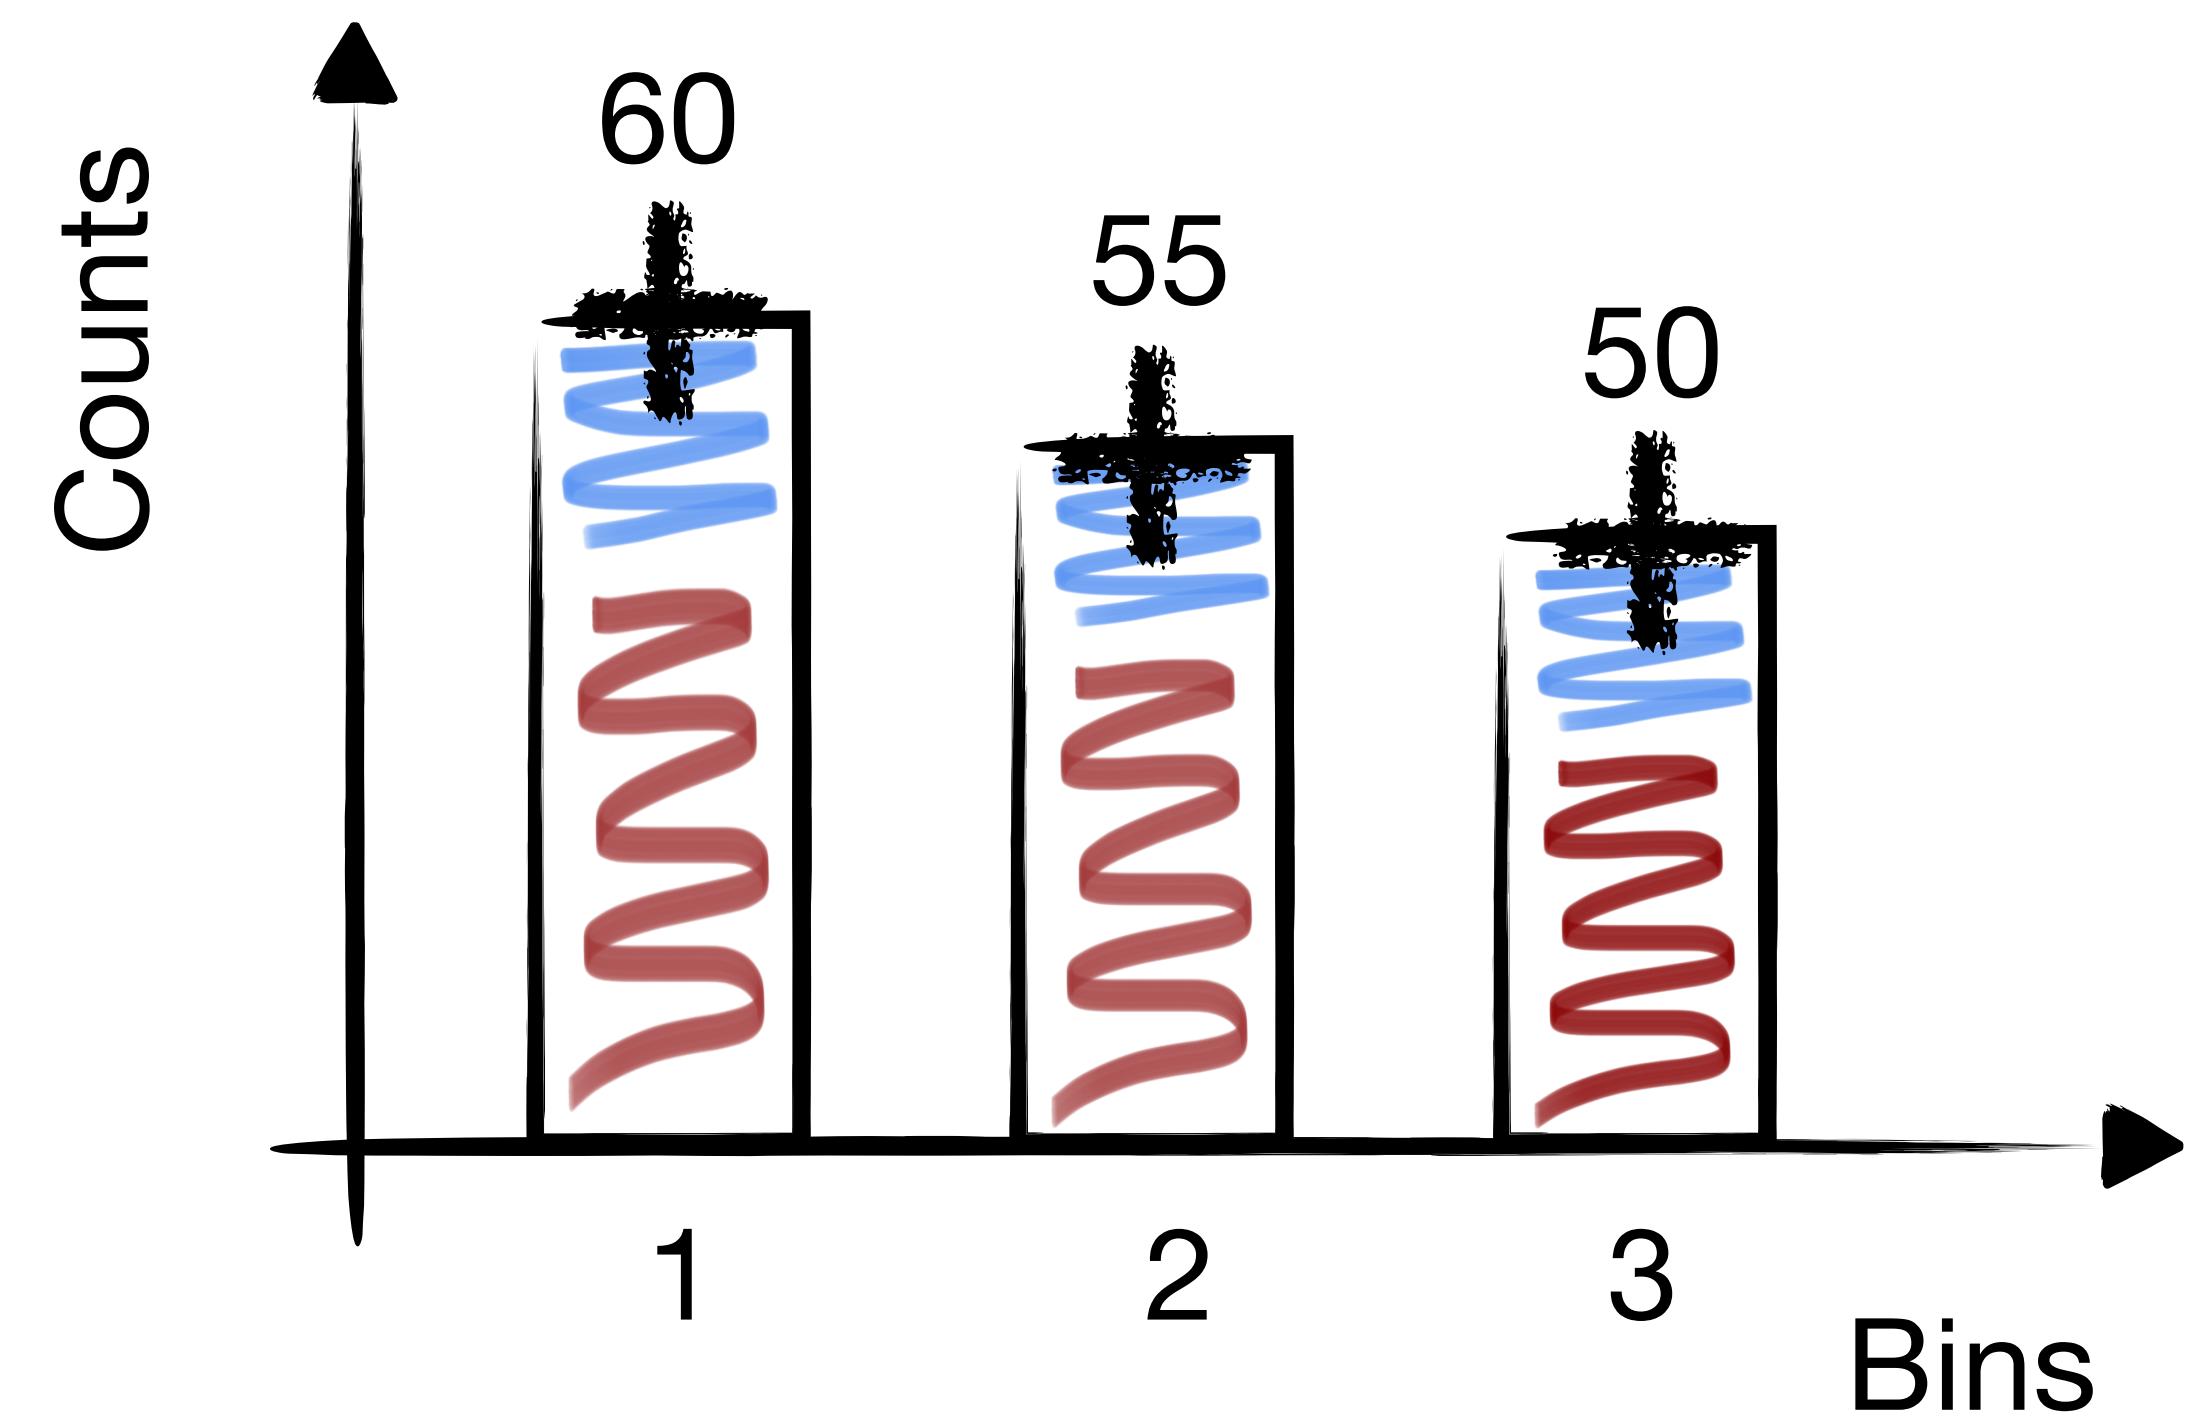
- Bin counts: $N = \eta s + \chi b$



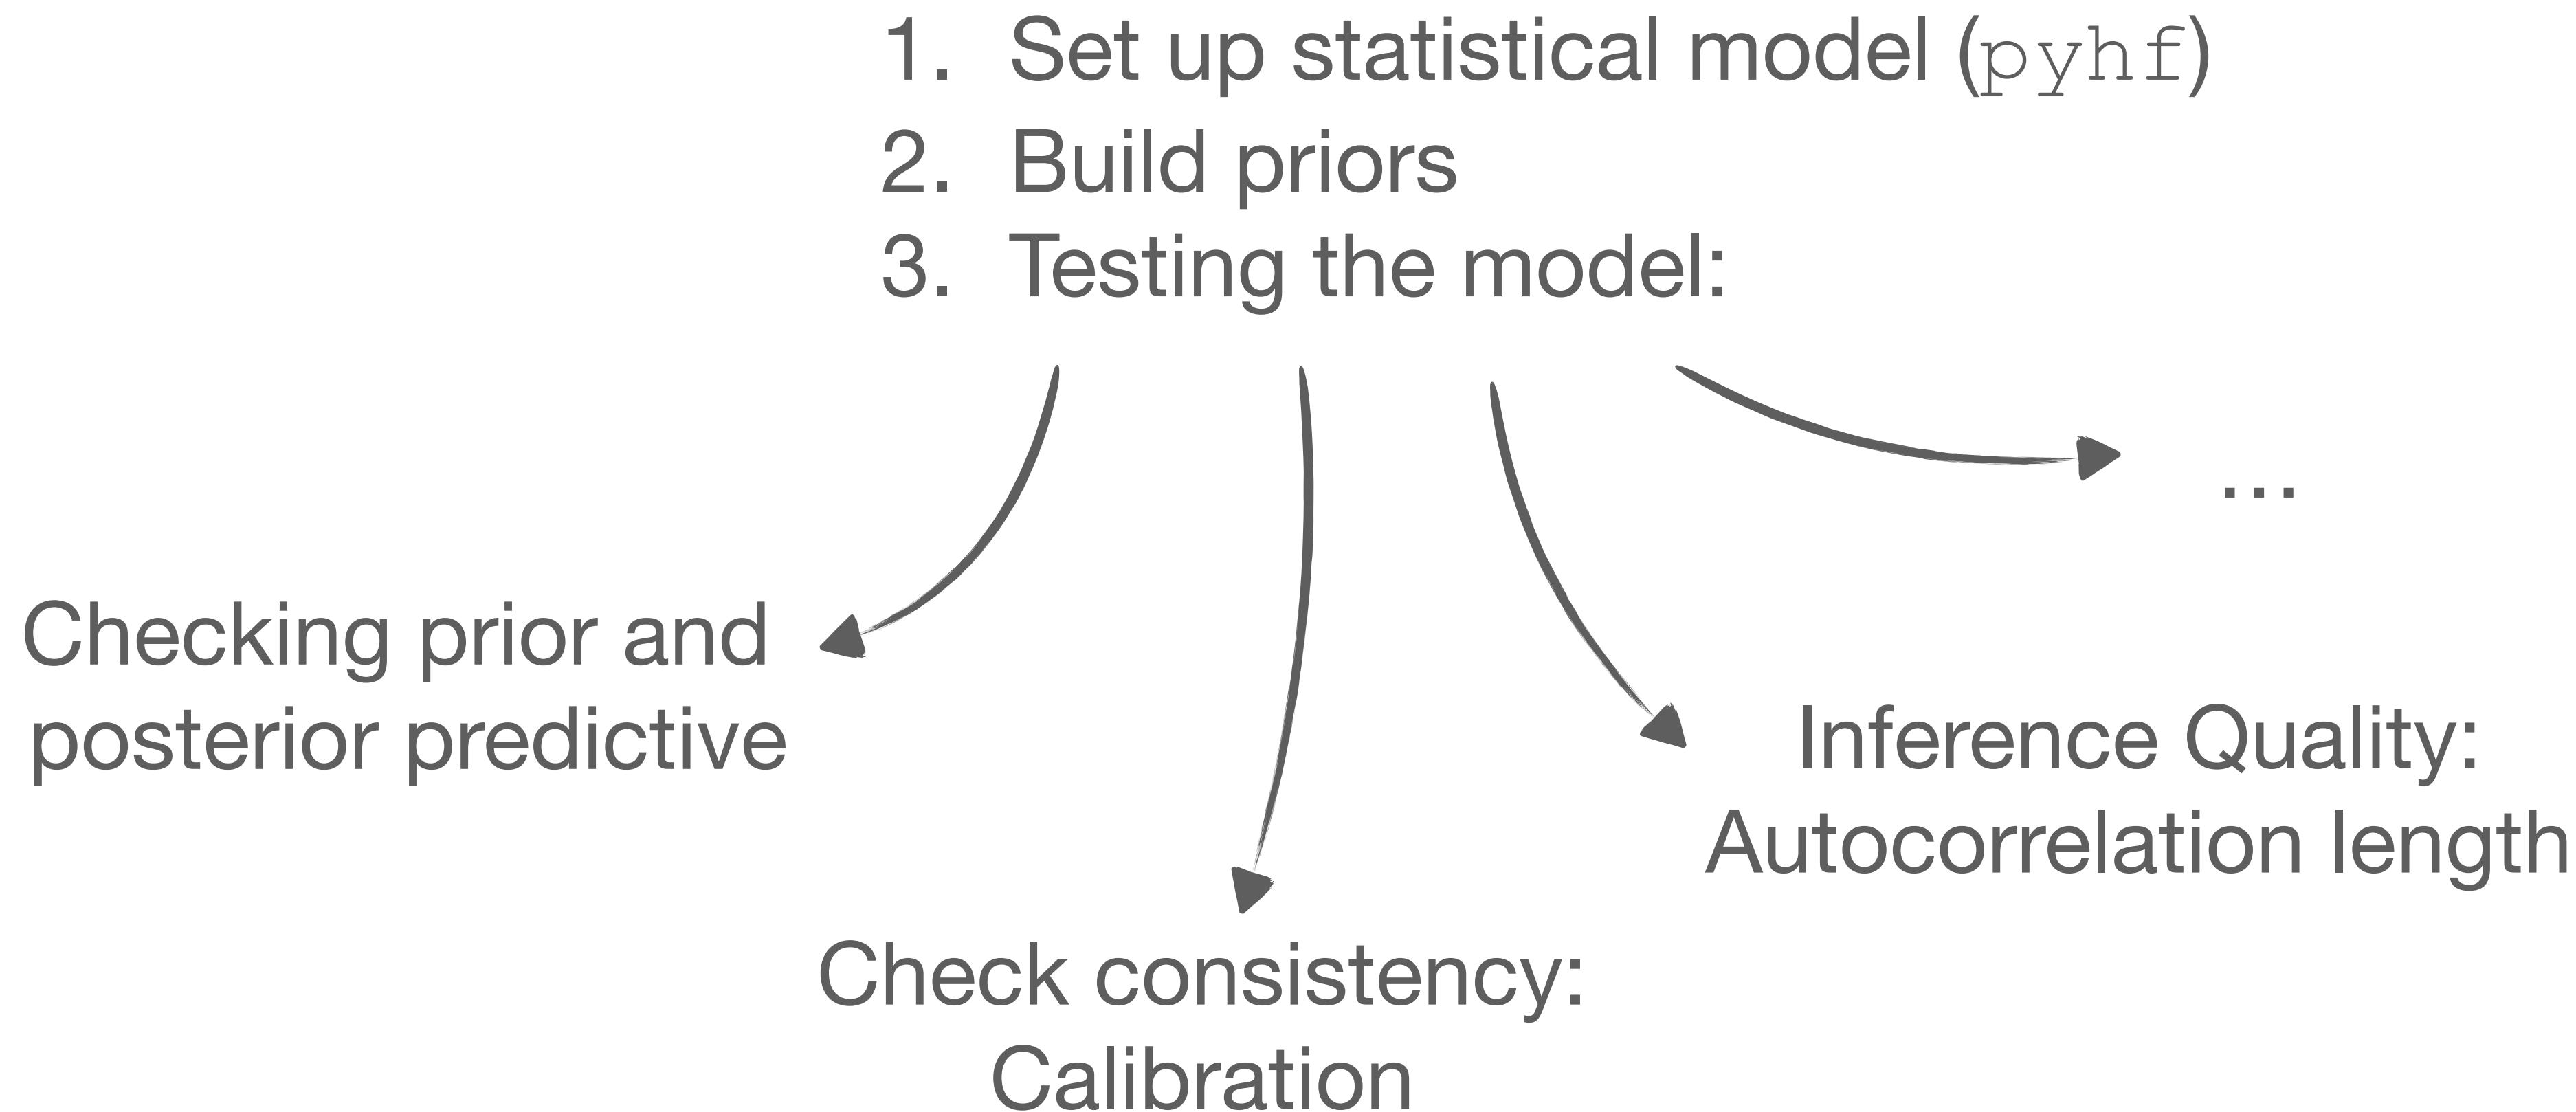
Background (Nuisance): [50,40,30]

Signal Strength (Unconstrained): [10, 15, 20]

- Data: [60, 55, 50]



Example Bayesian Workflow

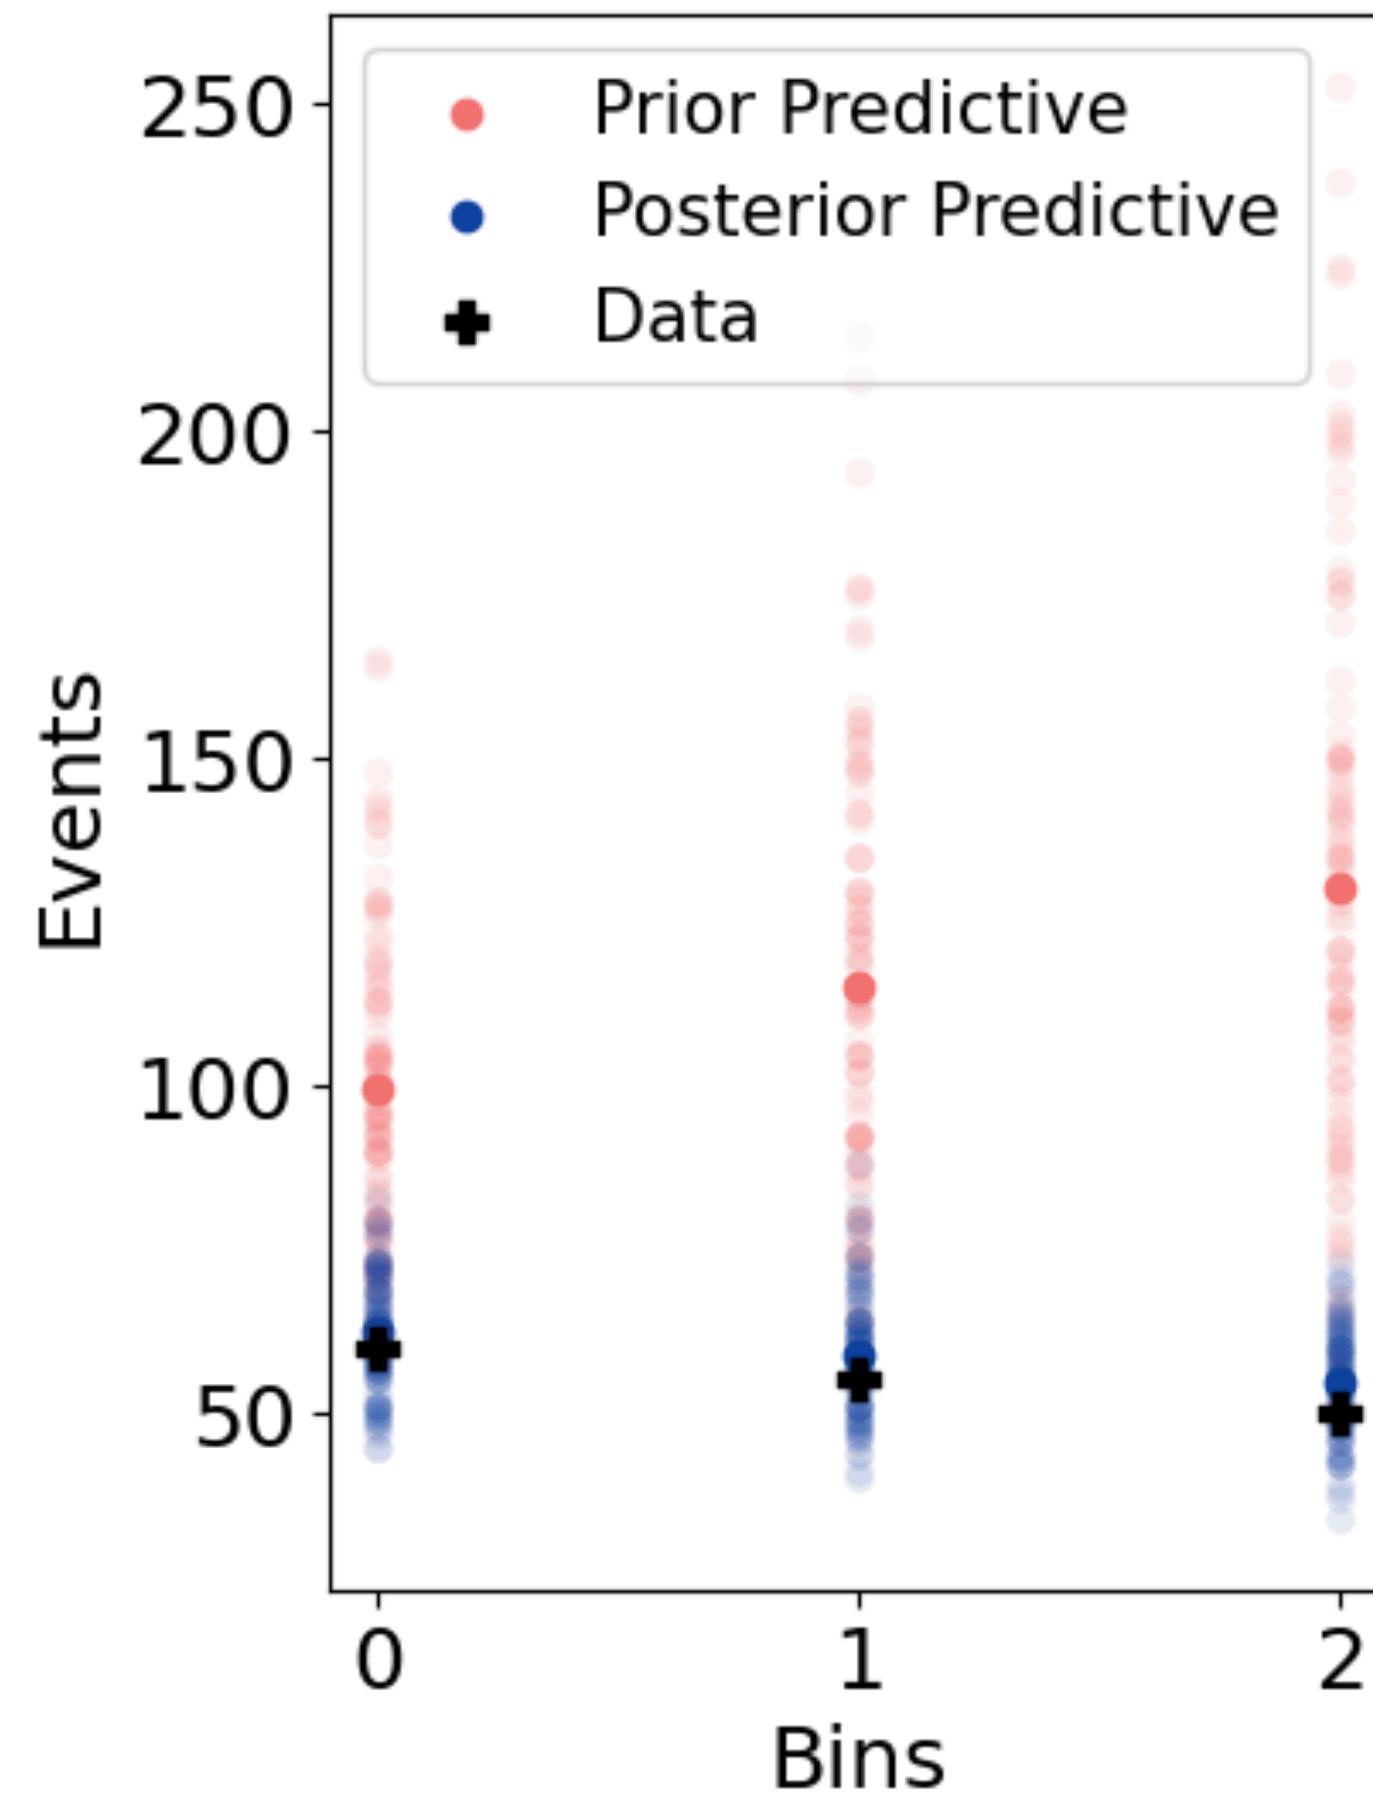
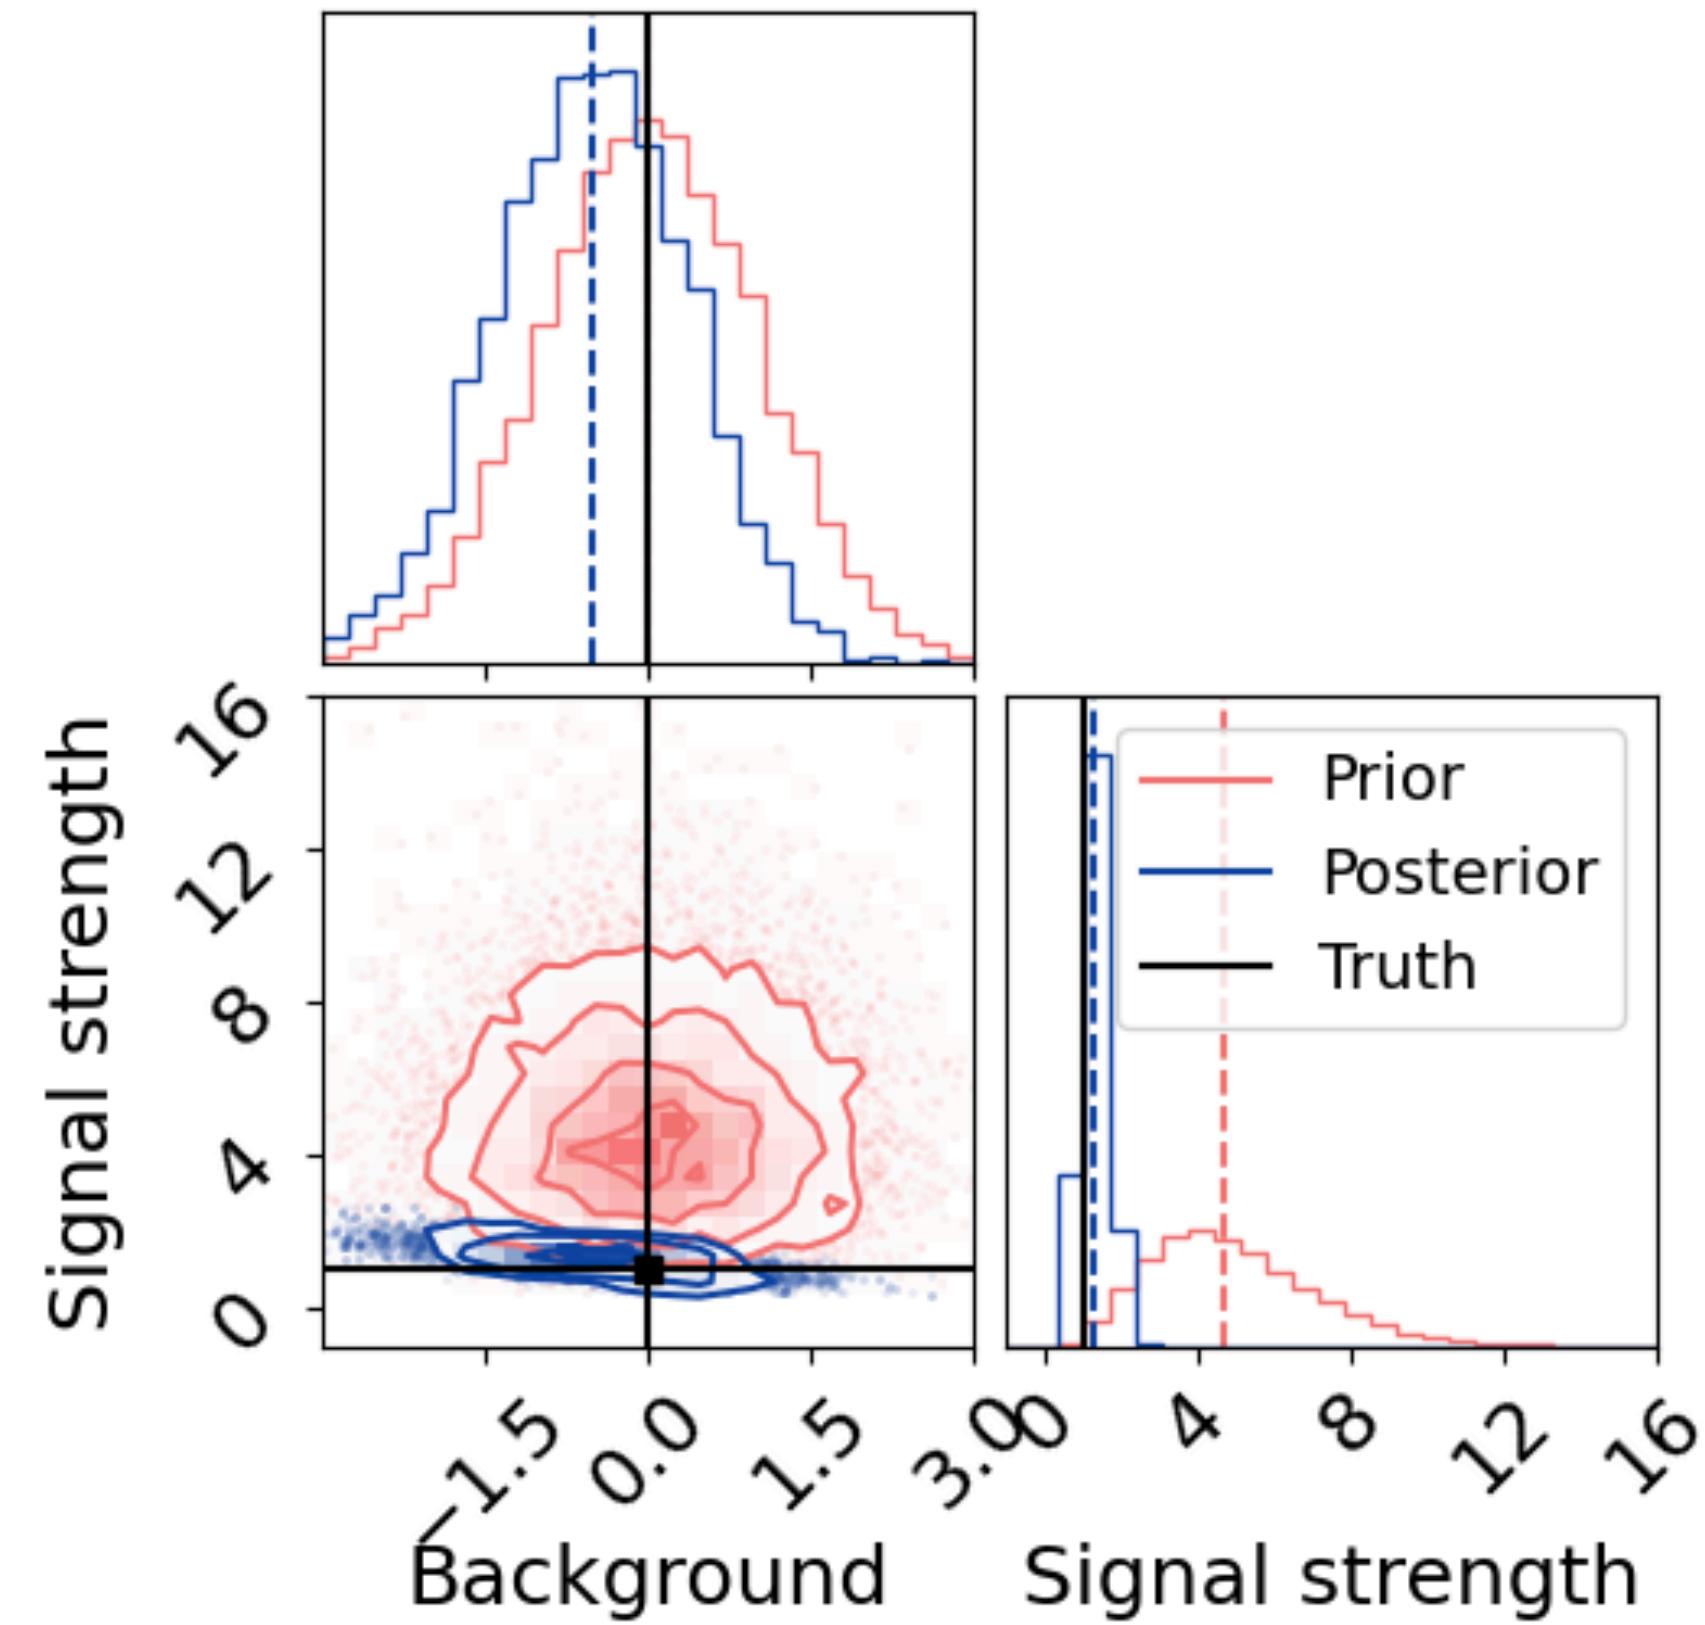


$$p(\vec{\eta}, \vec{\chi} | \vec{n}, \vec{a}) \approx p(\vec{n} | \vec{\eta}, \vec{\chi}) p(\vec{\eta}) p(\vec{\chi} | \vec{a})$$

$$N = \eta s + \chi b$$

Example Bayesian Workflow

- Prior and posterior predictive:



Example Bayesian Workflow

Calibration: Testing Computational Faithfulness

- Average posterior draws with observations sampled from prior predictive should converge to prior distribution:

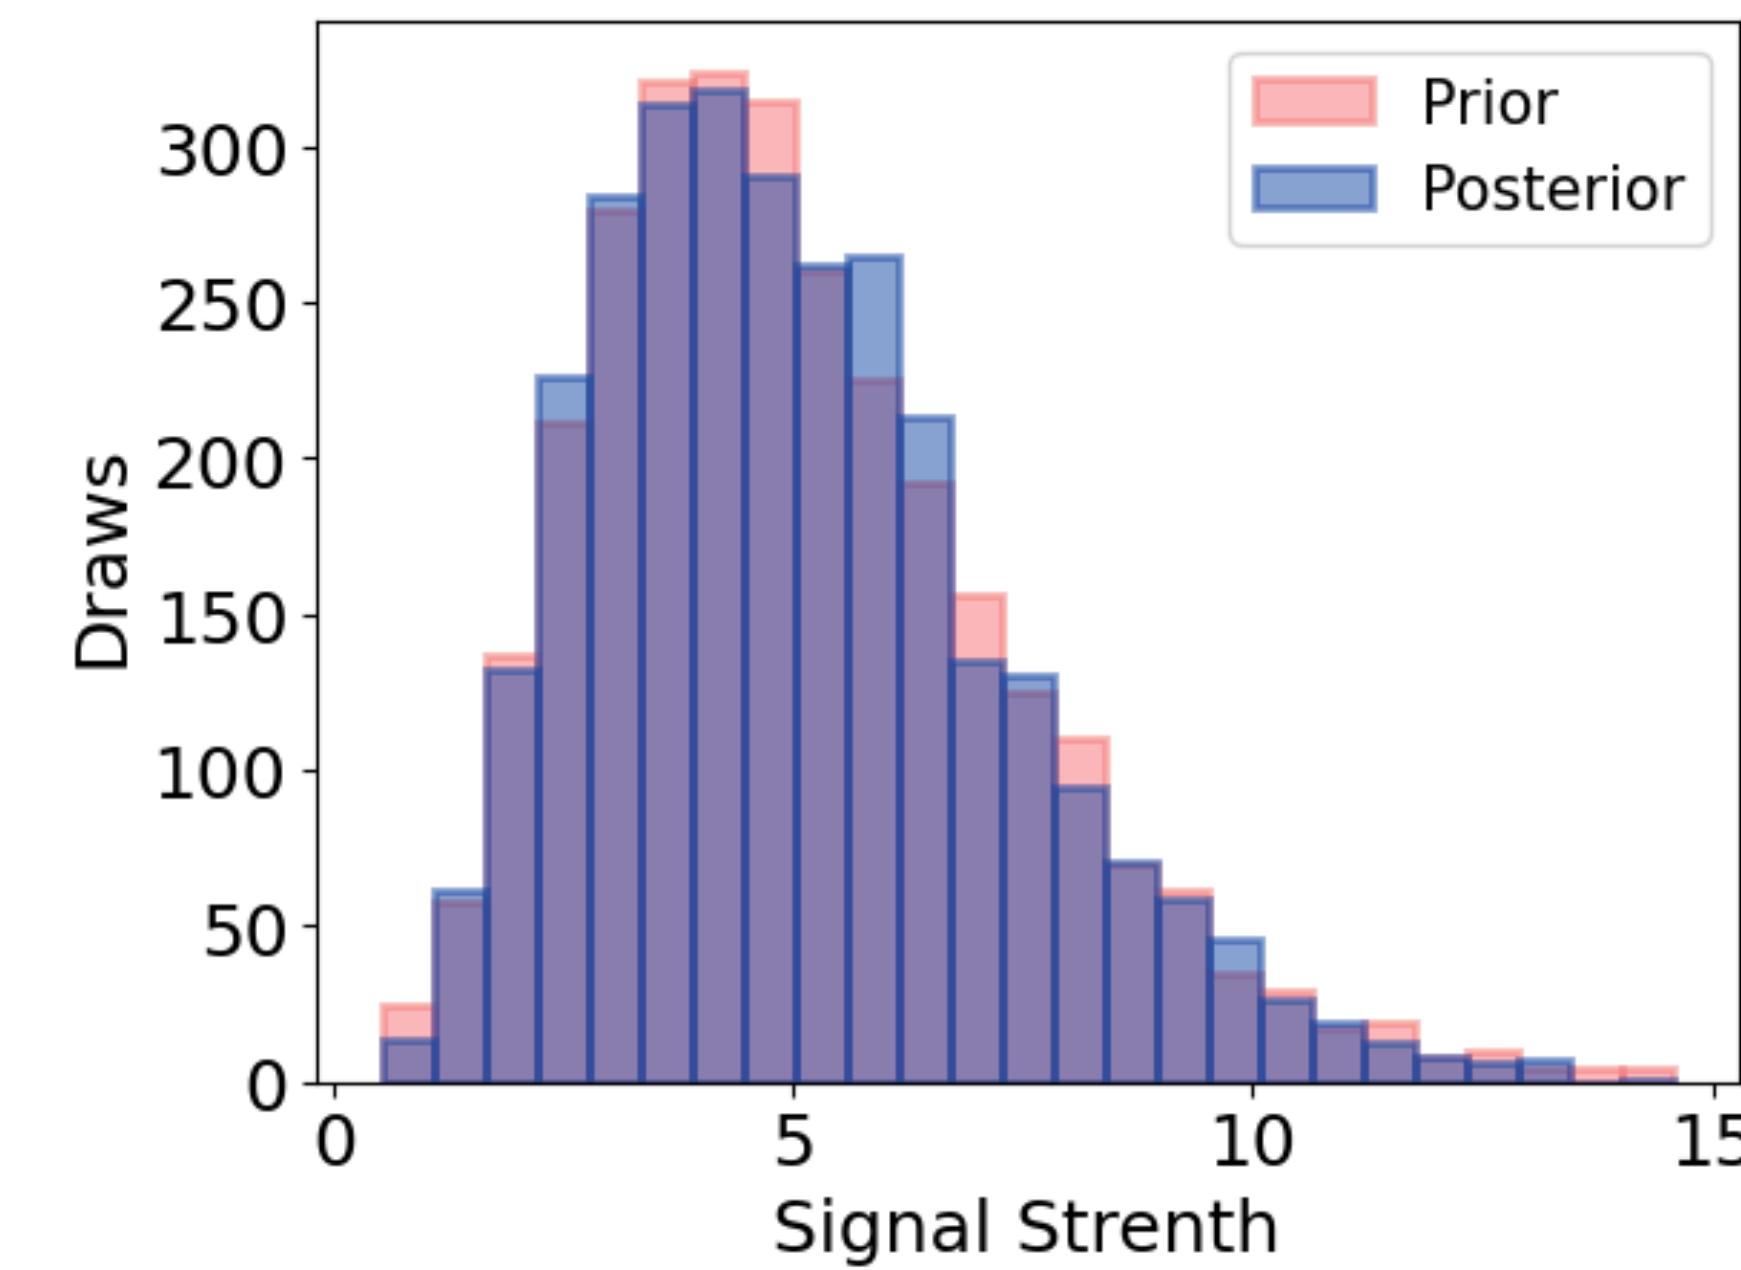
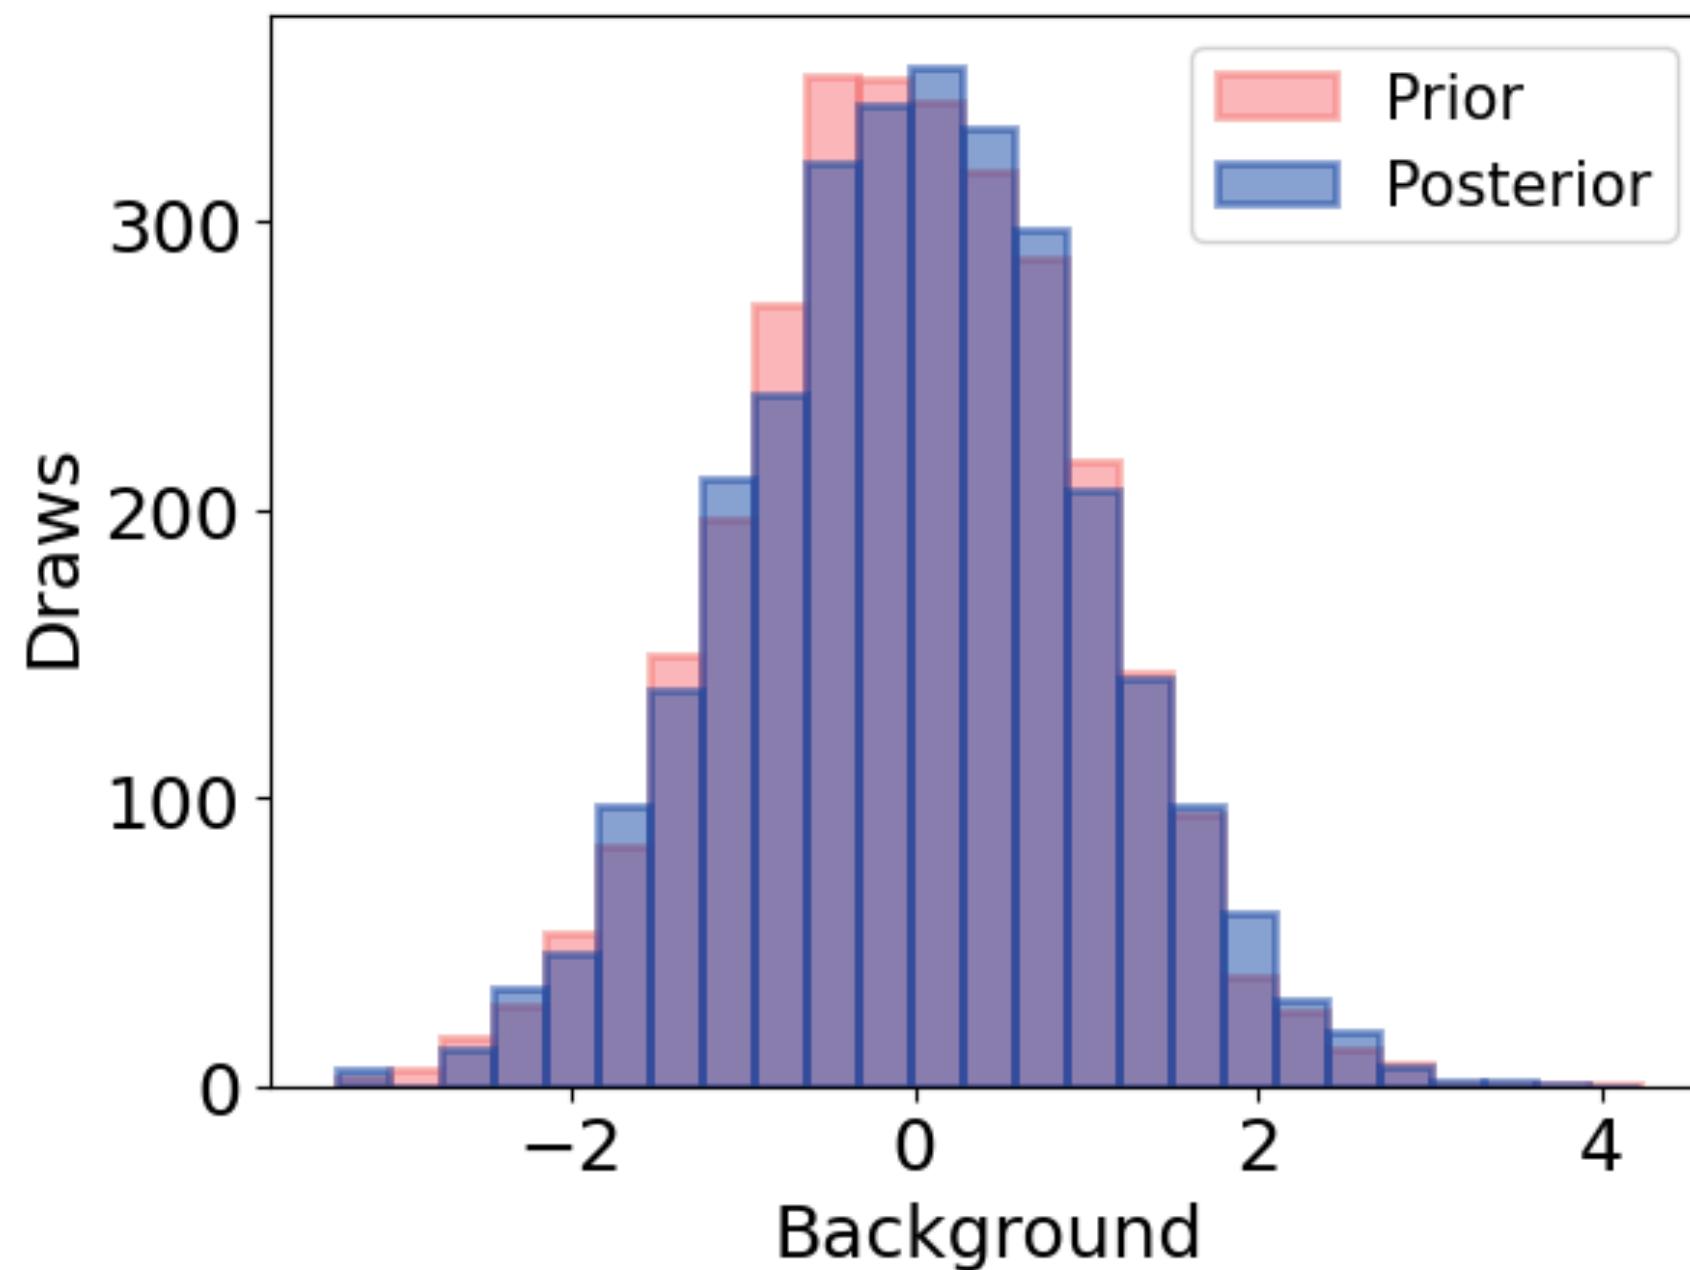
$$p(\theta) \stackrel{!}{\approx} \int dy d\theta' p(\theta|y) p(y|\theta')$$

Example Bayesian Workflow

Calibration: Testing Computational Faithfulness

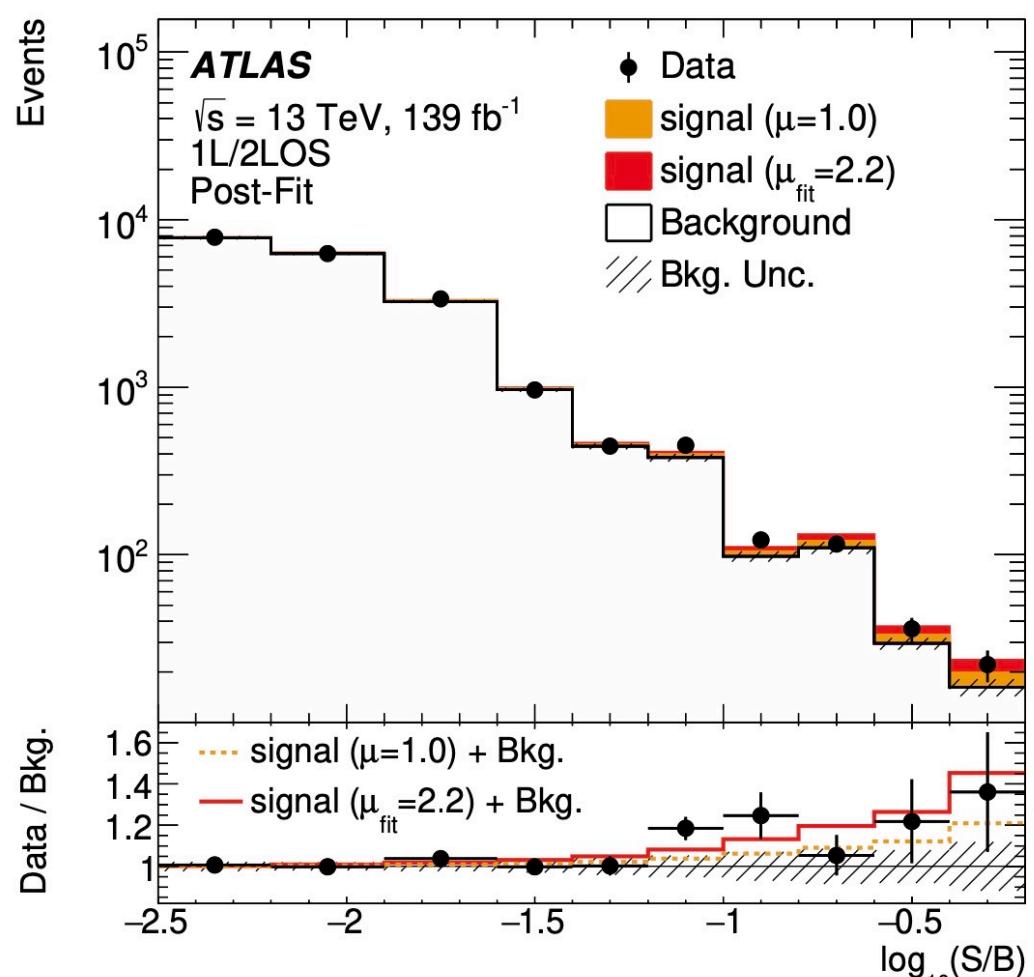
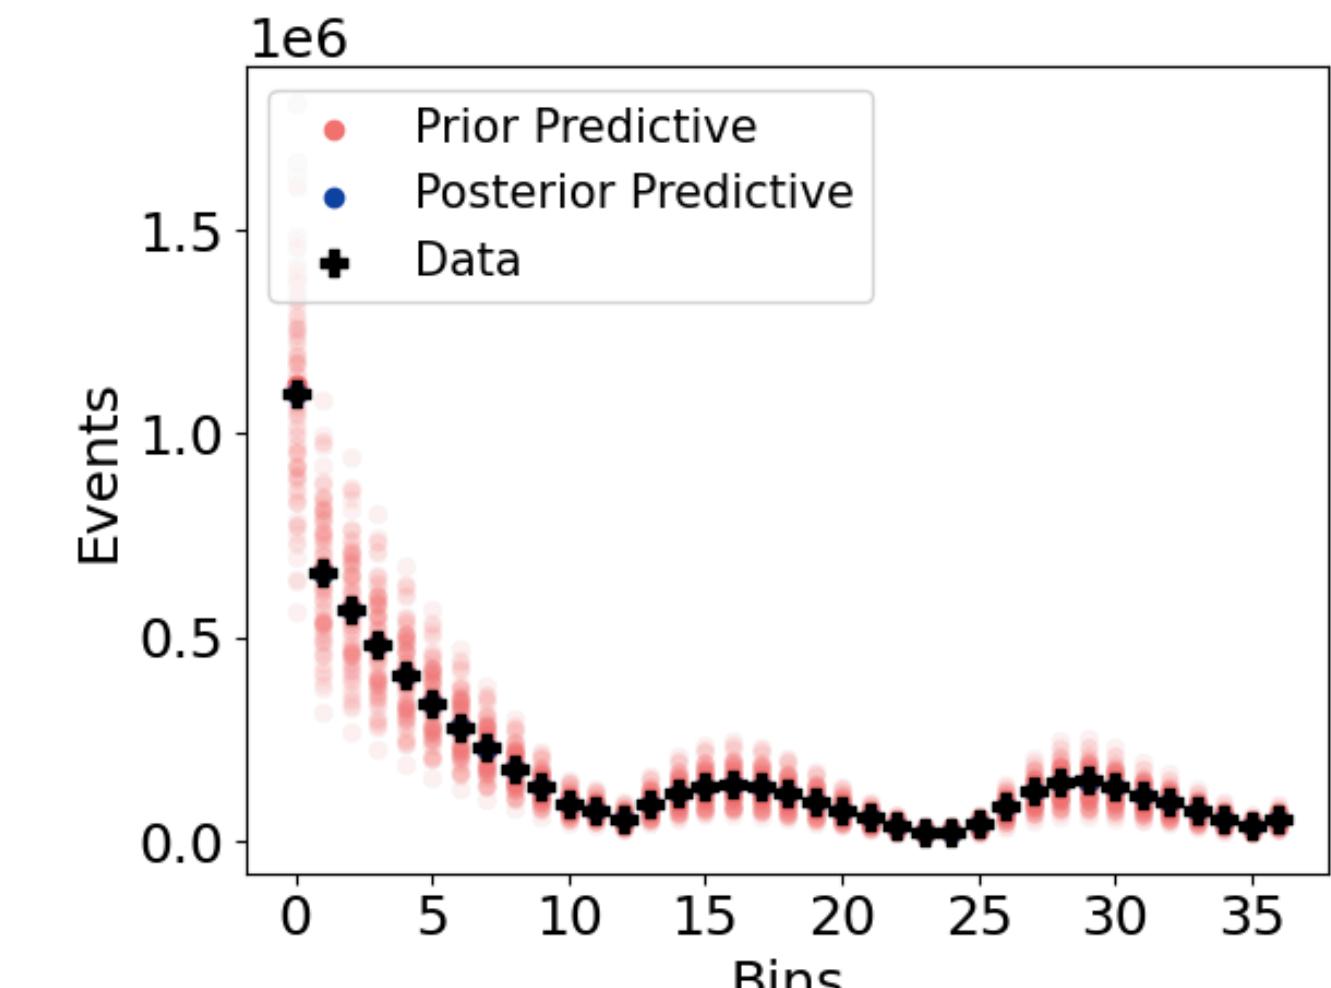
- Average posterior draws with observations sampled from prior predictive should converge to prior distribution:

$$p(\theta) \stackrel{!}{\approx} \int dy d\theta' p(\theta|y) p(y|\theta')$$



Conclusion and Outlook

- All of this work is open-source and available under:  DOI [10.5281/zenodo.7886632](https://doi.org/10.5281/zenodo.7886632)
github.com/malin-horstmann/pyhf_pymc
- Parallel Bayesian and Frequentist inference for HistFactory models is now possible
- Next steps:
 - Integrate into pyhf
 - Continued tests for robustness (i.e. parallel chains)
- An ATLAS public likelihood example:
<https://doi.org/10.17182/hepdata.105039>
- Want to try it?



Backup

Computation Times

	6000 (= 1000)	12_000 (= 1000)
Metropolis		11 s
NUTS	14 s	

Building Prior Distributions for $\vec{\eta}, \vec{\chi}$

- Turn information from auxiliary measurements (means, rates, uncertainties) into prior distributions for $\vec{\chi}$:

$$p(\vec{\chi}) \approx p(\vec{a}|\vec{\chi}) p_{\text{Ur}}(\vec{\chi})$$


- By using rules for conjugate priors, this inference turns trivial:

$p(a \chi)$	$\propto \text{Poiss}(a \chi)$	$\propto \mathcal{N}(a \mu=0, \sigma=1)$
$p(\chi)$	$\propto \Gamma(\chi \alpha=a, \beta=a)$	$\propto \mathcal{N}(\chi \mu=0, \sigma=1)$

Implementing external models in PyMC

- Class: PyTensor Op

```
class VJPOp(Op):
    ...
    ...
    itypes = [pt.dvector,pt.dvector]
    otypes = [pt.dvector]

    def perform(self, node, inputs, outputs):
        (parameters, tangent_vector) = inputs
        results = jitted_vjp_expData(parameters, tangent_vector)
        outputs[0][0] = np.asarray(results)

vjp_op = VJPOp()
```

```
class ExpDataOp(Op):
    ...
    ...
    itypes = [pt.dvector]
    otypes = [pt.dvector]

    def perform(self, node, inputs, outputs):
        (parameters, ) = inputs
        results = jitted_processed_expData(parameters)
        outputs[0][0] = np.asarray(results)

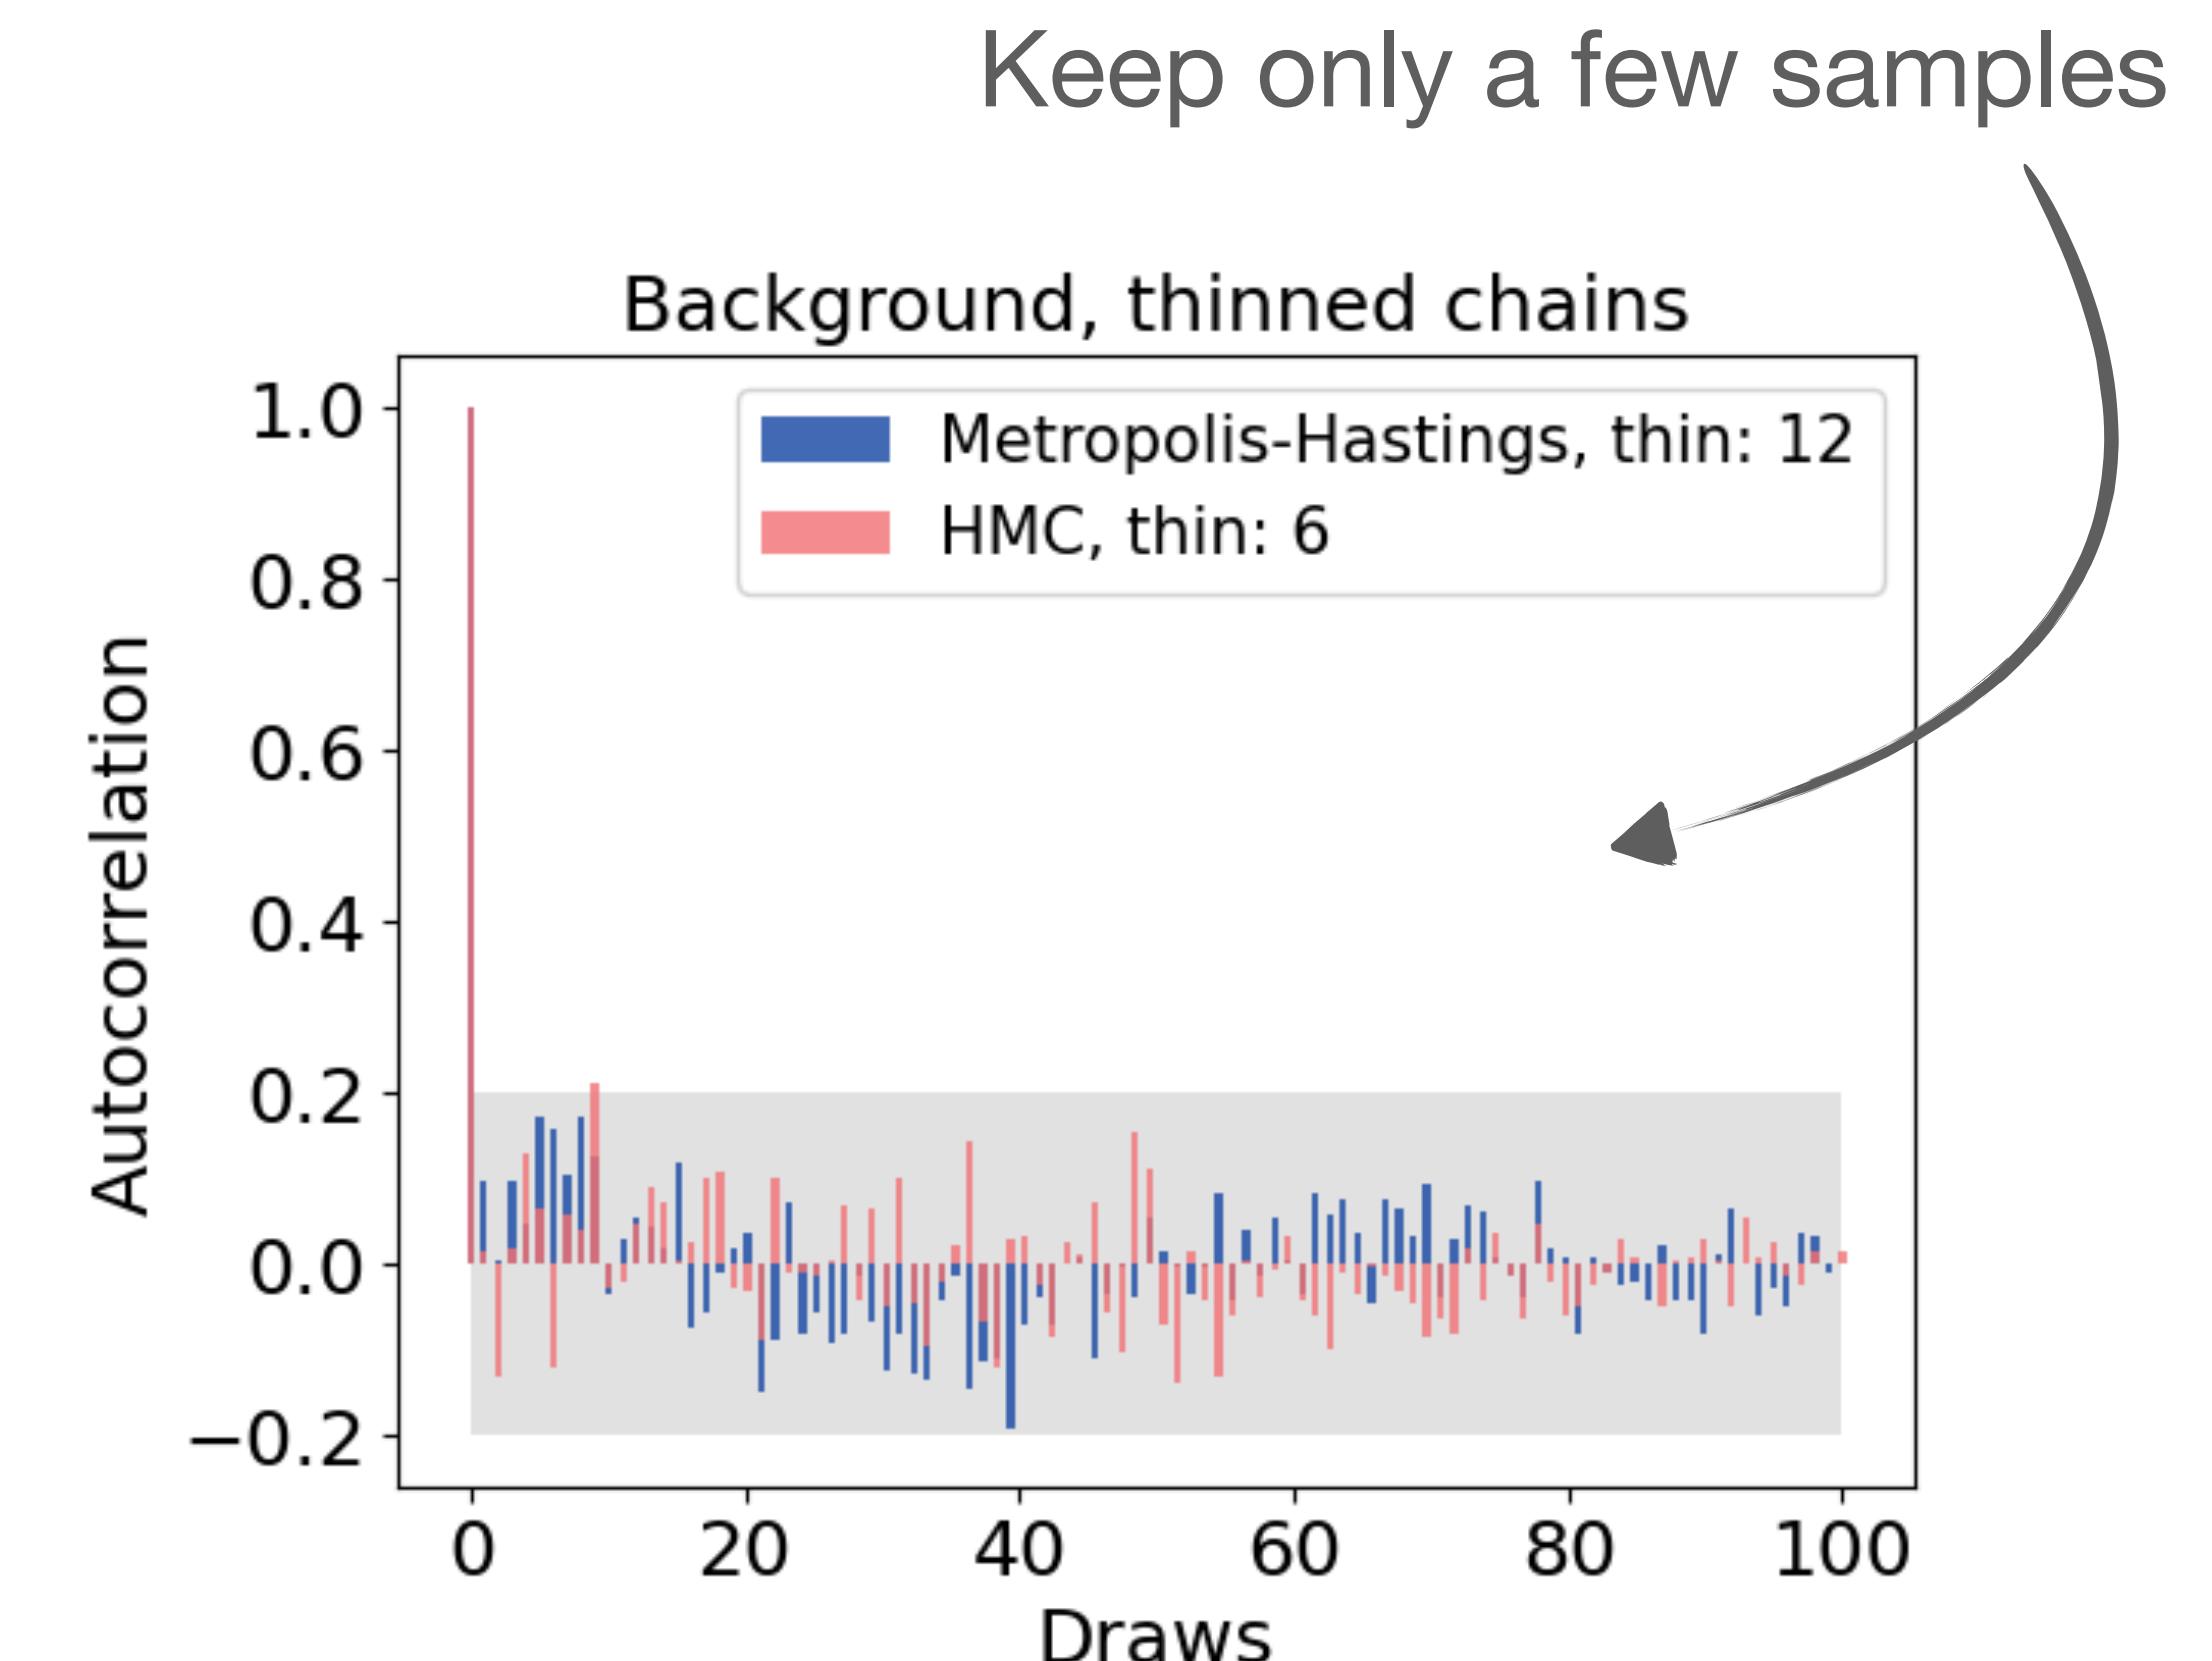
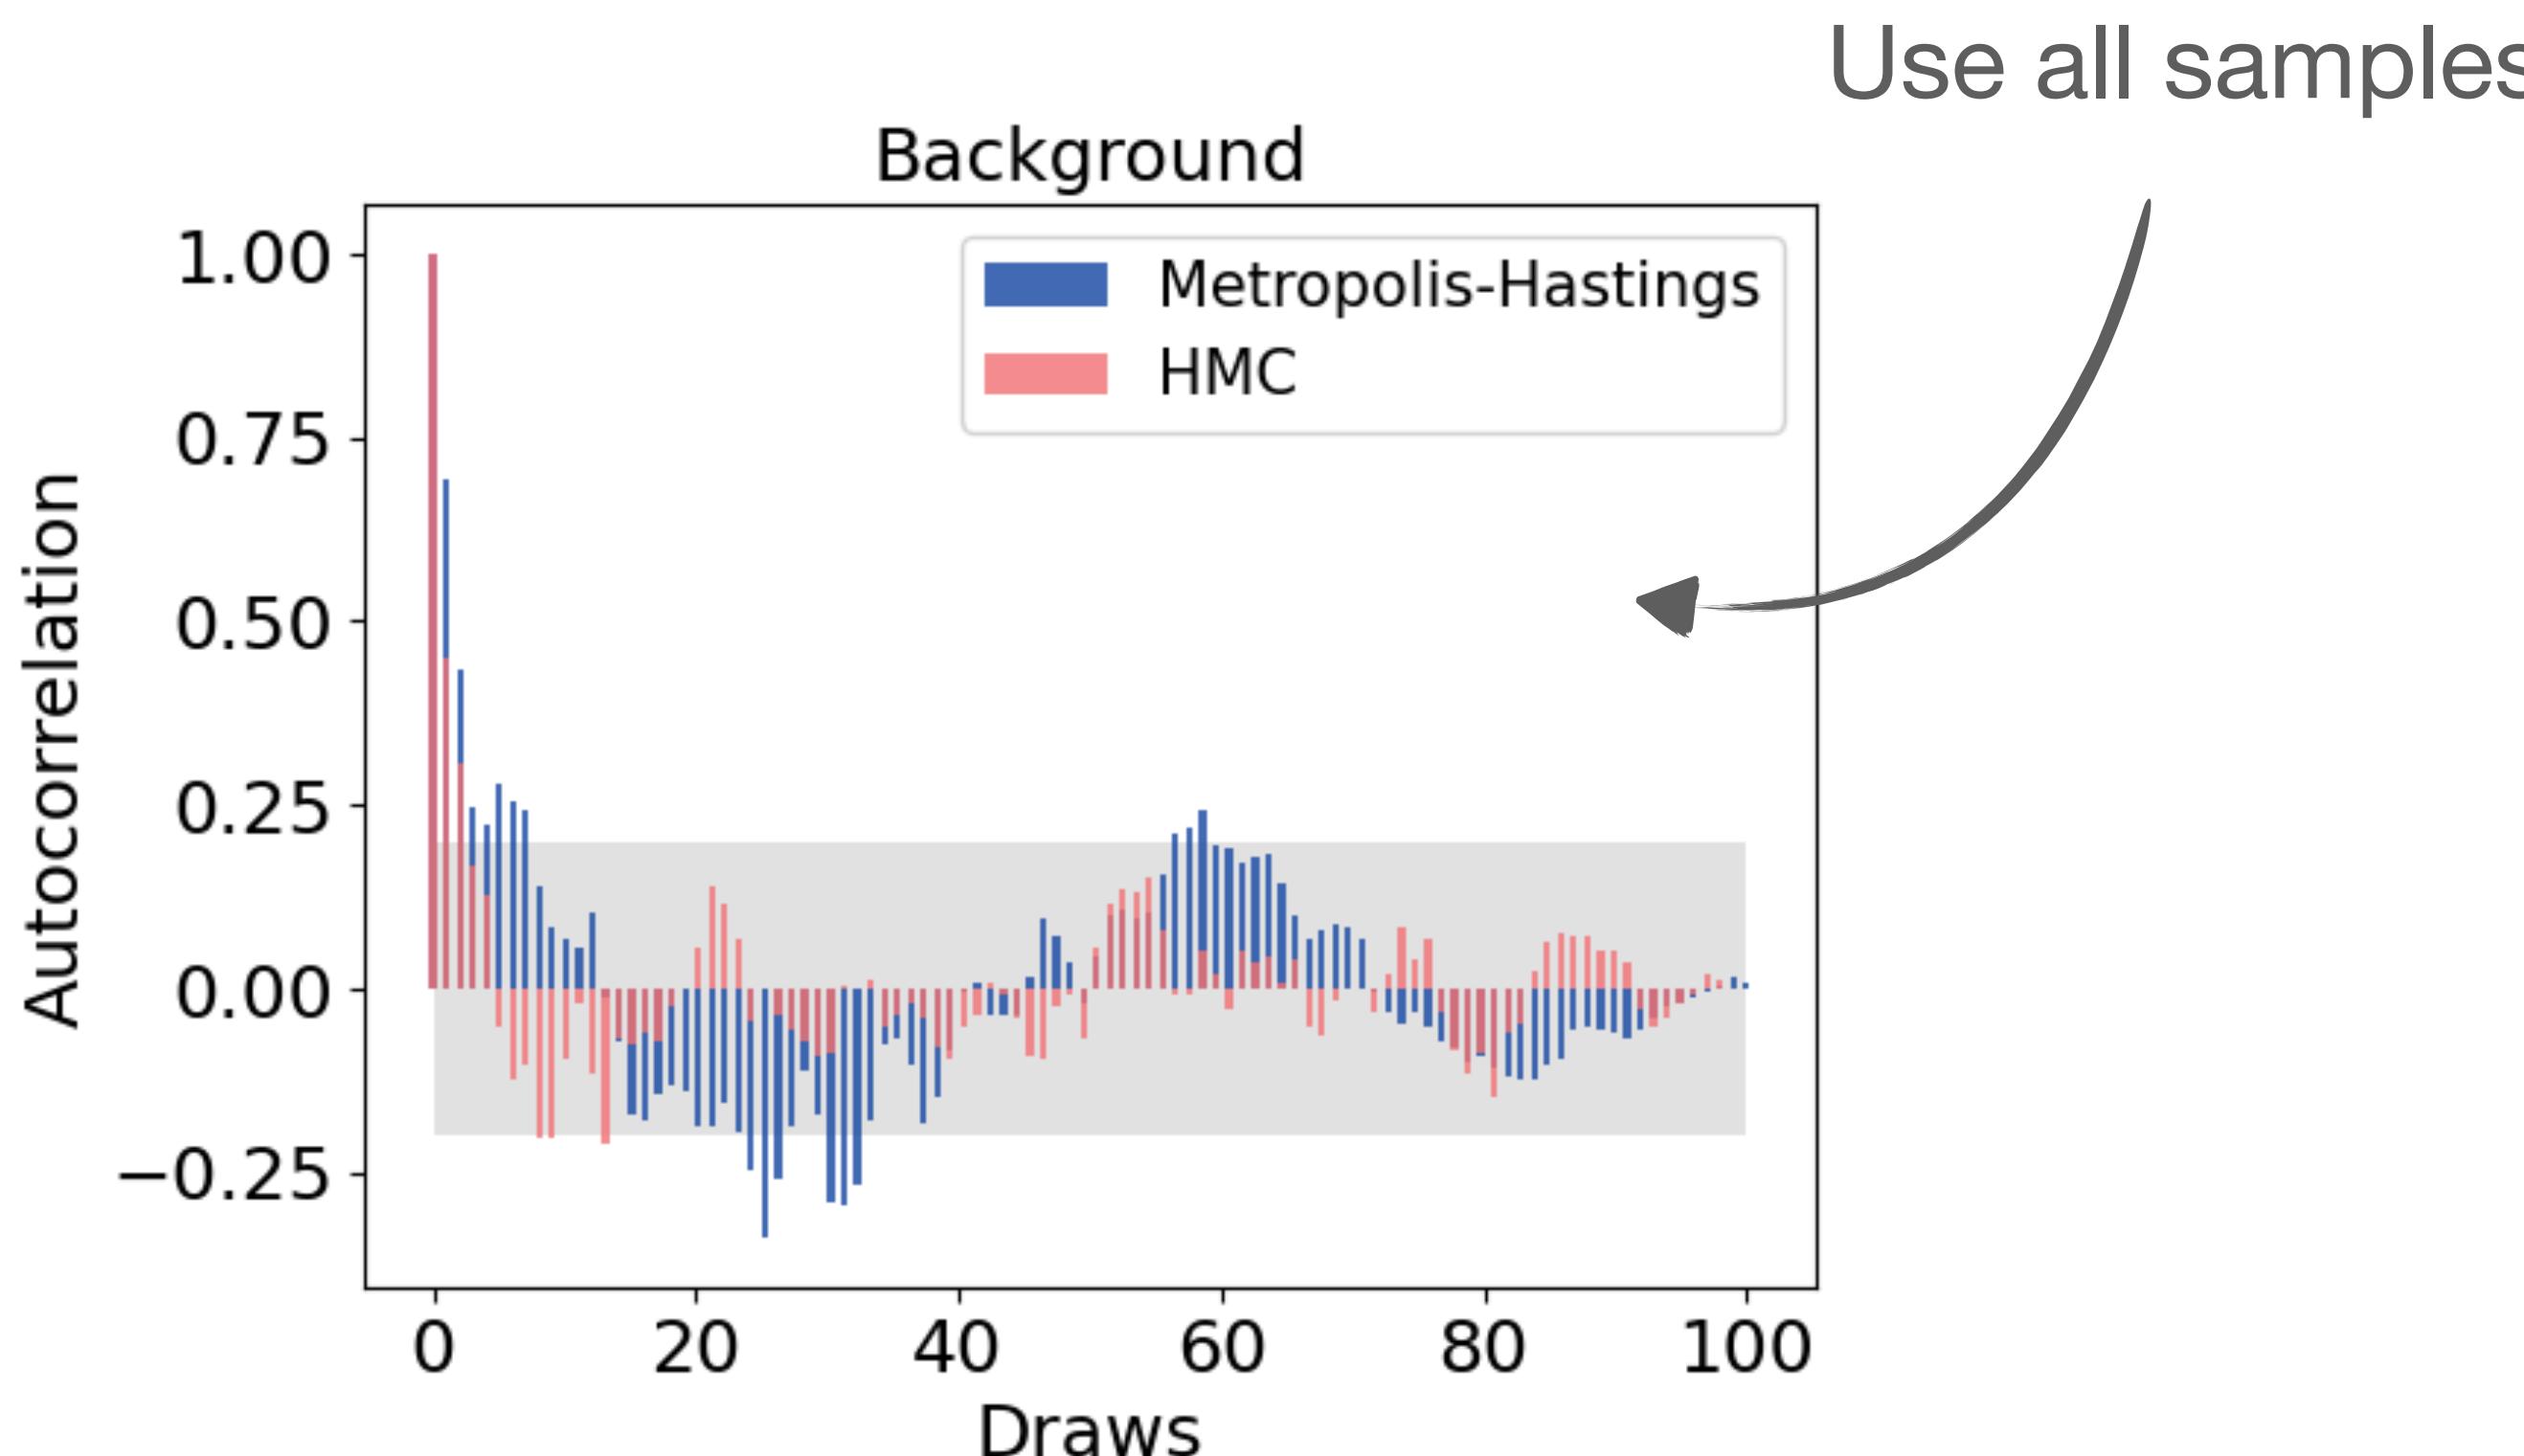
    def grad(self, inputs, output_gradients):
        (parameters,) = inputs
        (tangent_vector,) = output_gradients
        return [vjp_op(parameters, tangent_vector)]

expData_op = ExpDataOp() |
```

Example Bayesian Workflow

Autocorrelation: A Measure for the Inference Quality

- Gradient-based sampling (HMC) in general higher quality (but computationally heavier)
- Improve quality by thinning chains:



Autocorrelation

- $\text{ACF}(\tau) = \text{IFFT}(\text{FFT}(X(\tau))\text{FFT}^*(X(\tau)))$ (Wiener-Khinchin Theorem)

Ur-Priors:

- Poisson auxiliary measurement:

$$p(\chi | a) \propto \text{Poiss}(a | \chi) p_{\text{ur}}(\chi) \text{ and}$$

$$p_{\text{ur}}(\chi) = \Gamma(\chi | \alpha, \beta)$$

- Possible choices of ur-prior:
 - Mean and uncertainty from observation from auxiliary measurement
 - Uninformative

