Introduction to Bayesian Statistics

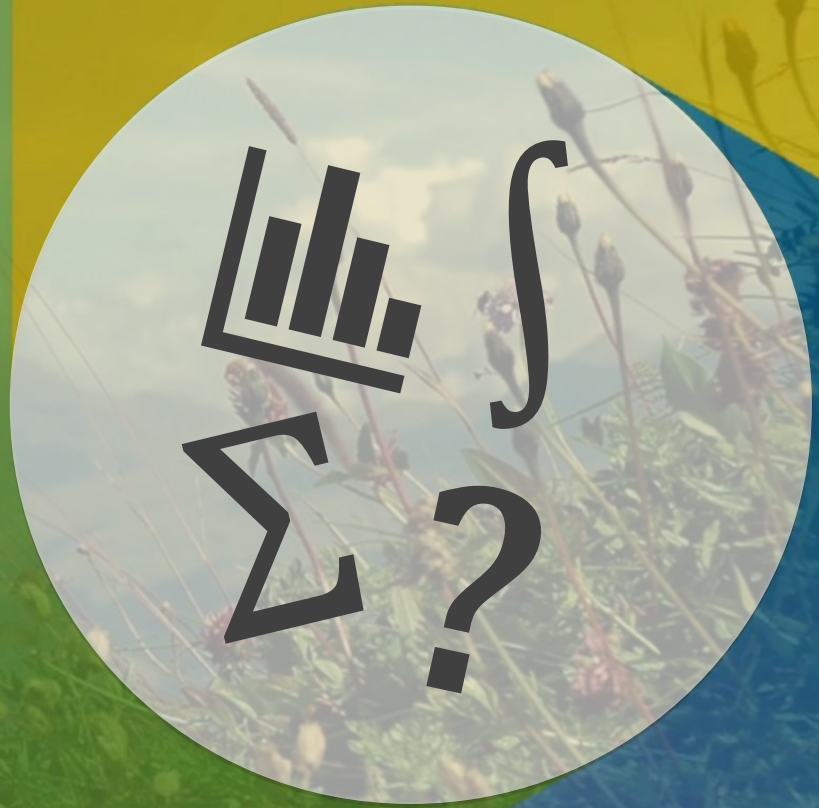

Part 8
Conclusions

Benjamin Rosenbaum

iDiv 2025

Statistical modeling

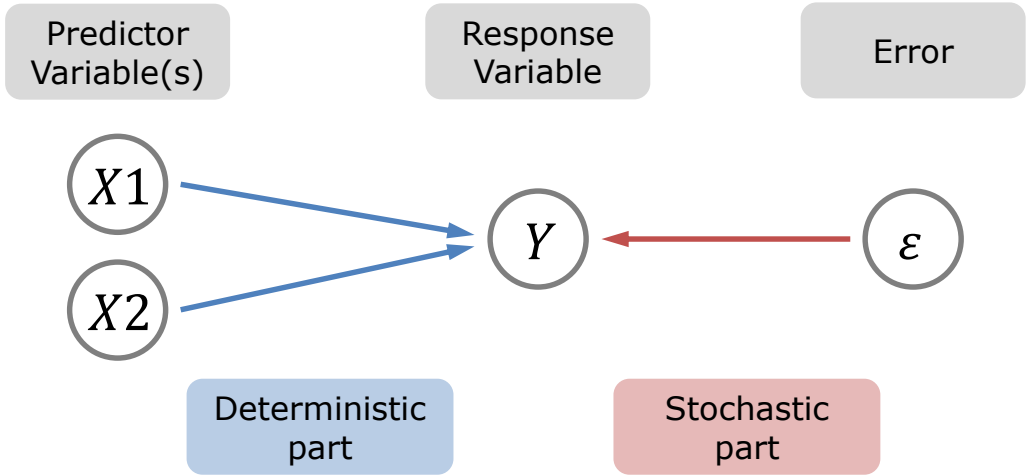

- Deterministic part:
  Prediction model, e.g. mean regression line
- Stochastic part:
  The prediction model cannot explain response perfectly, include random error
- Deterministic and stochastic parts both have **parameters** (e.g. effect sizes)

Statistical modeling

Model

Prediction

Probability
$P(\boldsymbol{y}|\theta)$

Likelihood
$L(\boldsymbol{\theta}|y) = P(y|\boldsymbol{\theta})$

Inference ?

Data / Nature

We want to make **quantitative statements** on our research questions:

What does the data tell me about my model?

Statistical modeling

Model

Prediction

Probability
$P(\boldsymbol{y}|\theta)$

Likelihood
$L(\boldsymbol{\theta}|y) = P(y|\boldsymbol{\theta})$

Inference ?

Data / Nature

Frequentist statistics cannot (mathematically) do direct inference $P(\boldsymbol{\theta}|y)$,

and requires a (methodological) detour via **NHST** $P(y|\boldsymbol{\theta} = \mathbf{0})$

Statistical modeling

Model

Prediction

Probability
$P(\boldsymbol{y}|\theta)$

Probability
$P(\boldsymbol{\theta}|y)$

Inference !

Data / Nature

Bayesian statistics can (conceptually) do direct inference $P(\boldsymbol{\theta}|y)$, but requires a (computational) detour via **MCMC**

Bayesian workflow

1) Research question (hypotheses)
2) Data collection
3) Statistical model
4) Prior distribution choice
5) Model fitting (MCMC)
6) Evaluate model output

Revise model

7) Quantitative statements on hypotheses
8) Reporting

→ Workflow not that different from frequentist statistics.

(1) Research question
(2) Data collection

Your responsibility. You are the experts!

Don‘t let statistical methods limit your creativity in asking important questions.

But at least think of possible analyses before planning your study / experiments.

The Bayesian 3d printer is more flexible than the frequentist toolbox. But still limited to the framework of statistical modeling, and model identifiability

iDiv Ecotron
Source: Schmidt et al. (2021) Ecol. Evol. https://doi.org/10.1002/ece3.8198

(3) Statistical model

- Driven by research question / hypotheses!

Deterministic part (mean fitted response)

- Start small. Especially when using Stan.
  - few predictors
  - without interactions
  - simple or no random effects structure
- Not as „forward model selection" …
- … but to get an idea how the model works

Stochastic part (residual distribution & link function)

- Ideally already provided by data type

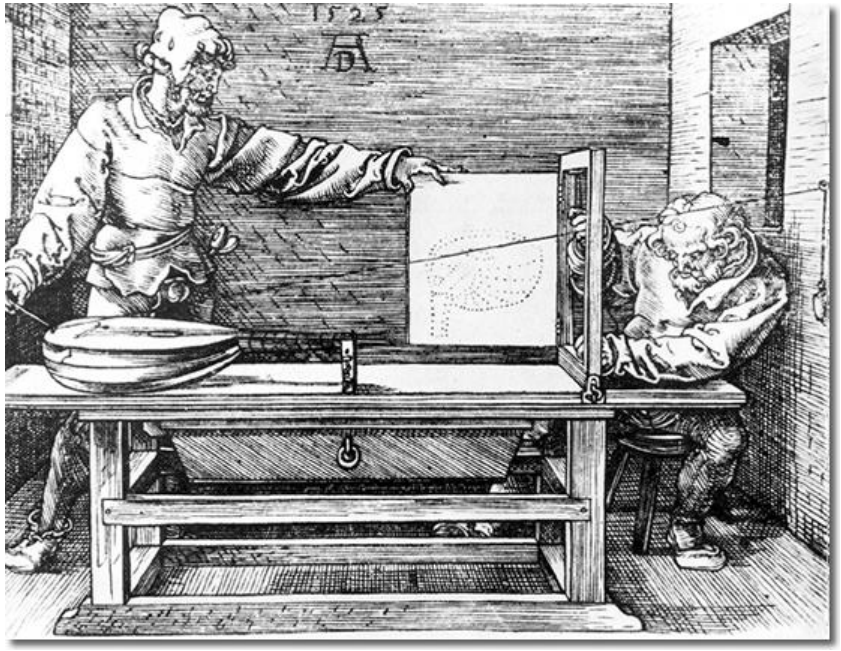

(4) Prior distribution

- Flat, weak, or informative priors?
- You can always do better than flat priors!
- Prior predictive checks
- Often easier with scaled predictors
- brms defaults for intercept and sd often good
- Priors for effects etc. are your choice
- `default_prior(y~x+…, data=…)`

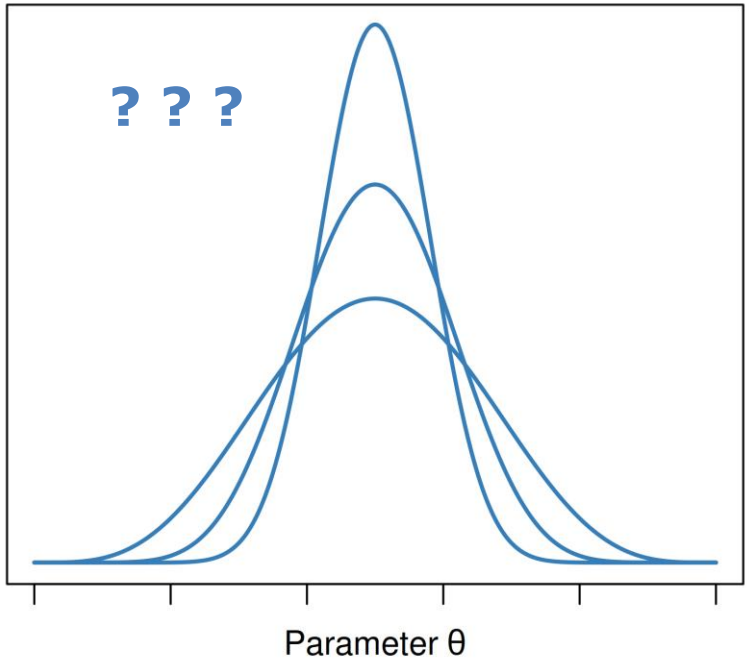

(5) Model fitting

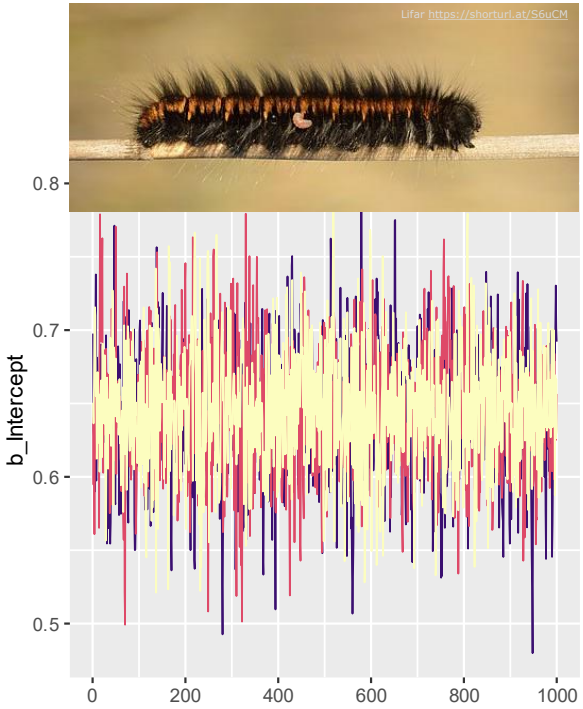

- Run several chains
- `cores=4` will run in parallel, saves time for big models

Check convergence

- visual inspection of traceplots
- check R_hat<1.01
- compare n_eff to n_total

  → few hundreds already good for parameter means,
  but a few thousands needed for parameter quantiles

- take warning messages seriously: https://mc-stan.org/misc/warnings.html

Convergence tells you if MCMC sucessfully approximated the posterior distribution.

Convergence does not tell you if it's a good or bad model !

(5) Model fitting

What to do if the model didn't converge?

- Provide initial values for parameters: `init=…`
- Run longer chains: `warmup=…` and `iter=…`
- Increase sampling accuracy: `control=list(adapt_delta=0.9)`
- Regularization through stricter priors

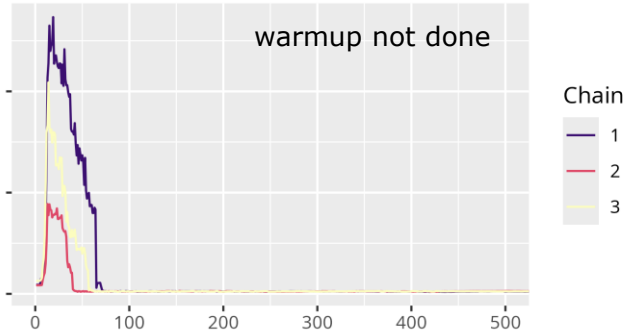

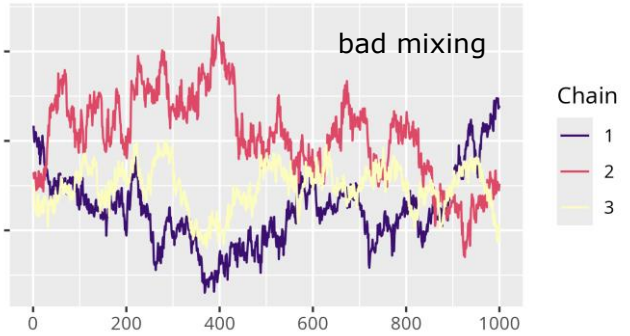

(6) Evaluate model

If, and only if MCMC converged, plot fitted / predicted against data

Options available from brms and performance package

Bayesian statistics does **not** free you from checking model assumptions

Conditional effects

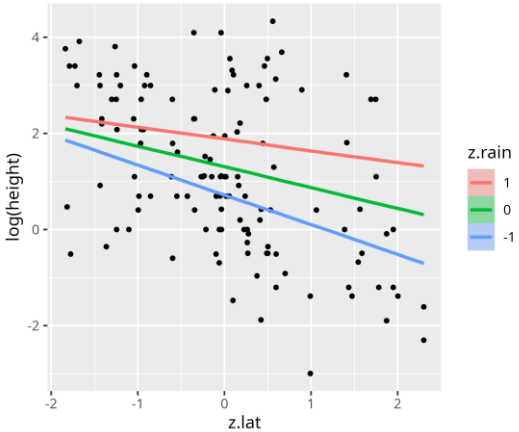

Observed vs. predicted

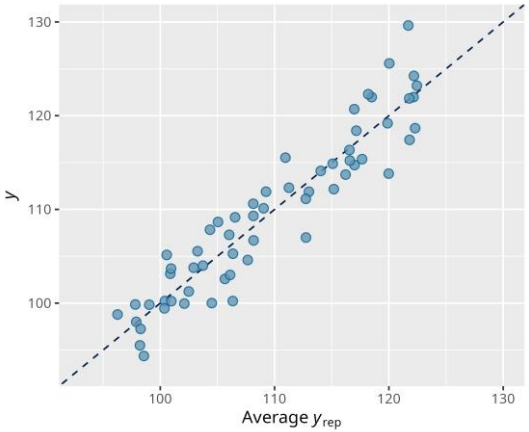

Posterior predictive check

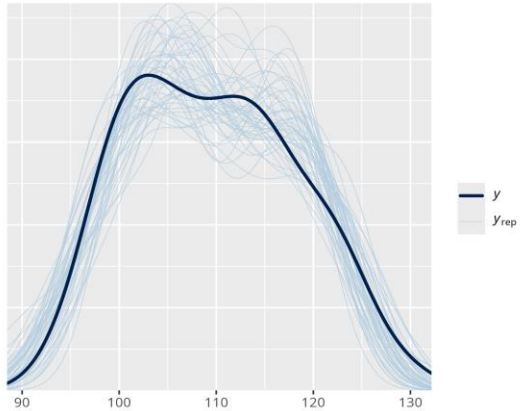

(6) Evaluate model

In Bayesian stats, **everything is a distribution.**

Always use posterior distribution of parameters
(not just point estimates) to compute predictions

Fitted distribution / credible intervals: deterministic model part

Predicted distribution / prediction intervals: stochastic model part

(7) Quantitative statements on hypotheses

- **Statistical inference:**

  What do the data tell me about my model?

- Read $P(\theta|y)$ literally:

  Quantitative statements on model (parameters θ)

  and their derived quantities, given the data y.

- $P(b > 0)$, $P(\mu_1 > \mu_2)$, $P\left(\frac{\mu_1+\mu_2}{2} > \mu_3\right)$, $P(\sigma_1 > \sigma_2)$, …

- „Post-hoc" analysis (emmeans) just computes and

  compares quantities from model parameters

  (level-specific means or slopes, across-level means

  or slopes, pairwise comparisons, group comparisons, …)

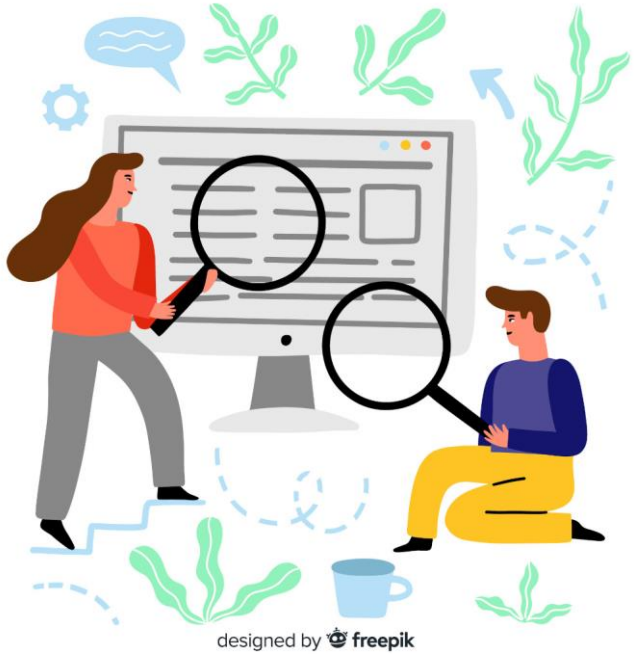

(7) Quantitative statements on hypotheses

Model comparison with LOO

- compares expected predictive performance „elpd“ of several models
- penalizes model complexity
- LOOIC is similar to AIC (lower values = better)

Good scientific practice:

Model candidates driven by hypotheses.

Don‘t just fit lots of models with different predictor combinations!

„Fishing for evidence“

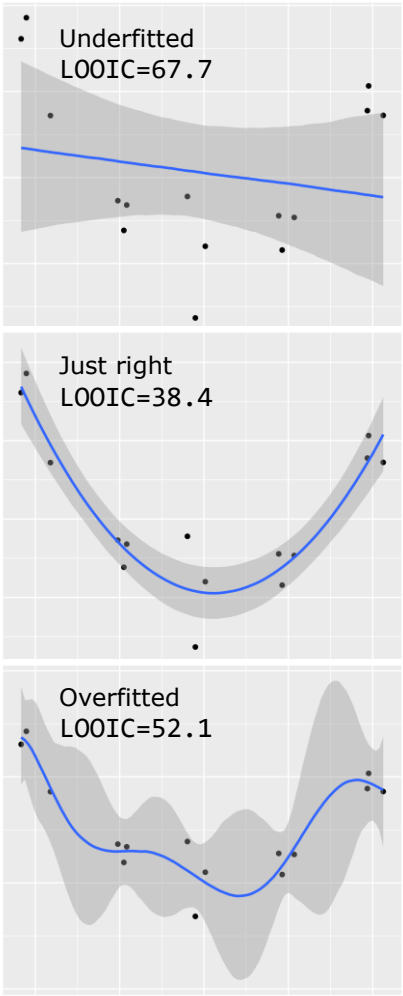

(8) Reporting

Methods:

We used Bayesian statistics for estimating the parameters of models (1)-(3) using the "brms" R-package. We chose normal distributions with zero mean and a standard deviation of 2 as weakly informative priors for effect sizes and brms default priors for intercept and standard deviations.

[Put priors & their justification in an SI table if you have many parameters]

We ran 4 MCMC chains of 5000 iterations with an adaptation phase of 2000 samples (12,000 posterior samples in total). We verified convergence by Gelman-Rubin statistics (Rhat<1.01) and adequate effective sampling size (n_eff). For model comparison we used leave-one-out cross validation (R-package "loo").

(8) Reporting

Results:

According to LOO, model (3) had the best predictive performance (ELPD difference to model (2): 5352.2 with standard error 709.2). We found a postive interaction effect of x1 with x2 (posterior probability P(bx1:x2>0)=0.970) and a negative effect of b3 (P(bx3<0)=1.000) which confirms our hypothesis X.

Parameter	mean	sd	2.5%	97.5%
Intercept	8.051	0.724	6.645	8.533
x1	0.815	0.269	0.255	1.310
x2	-0.263	0.370	-1.039	0.409
x1:x2	0.322	0.156	0.024	0.627
x3	-0.270	0.059	-0.387	-0.230

„report" package support brms:
https://easystats.github.io/report/reference/report.brmsfit.html

Some resources

Bayesian books

Highly recommended!

Uses his own „rethinking“ package, also based on Stan

Full course & online lectures:
https://github.com/rmcelreath/stat_rethinking_2024

brms translation by Solomon Kurz:
https://bookdown.org/content/4857/

3rd edition available soon?

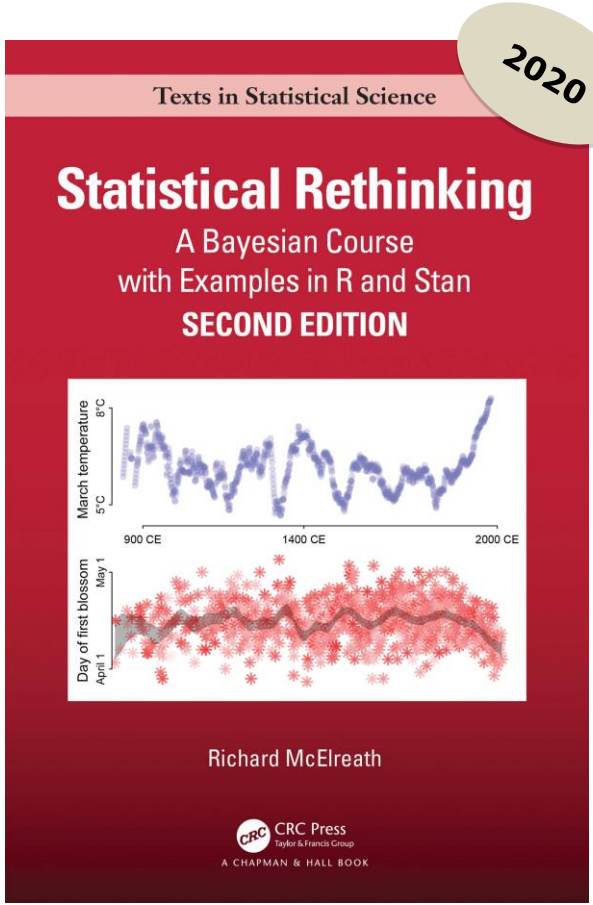

Bayesian books

Examples more focused on social & political sciences, but still general enough

Uses „rstanarm“ package

Update of their other book on mixed-effects models „Advanced Regression and Multilevel Models” will come out this year

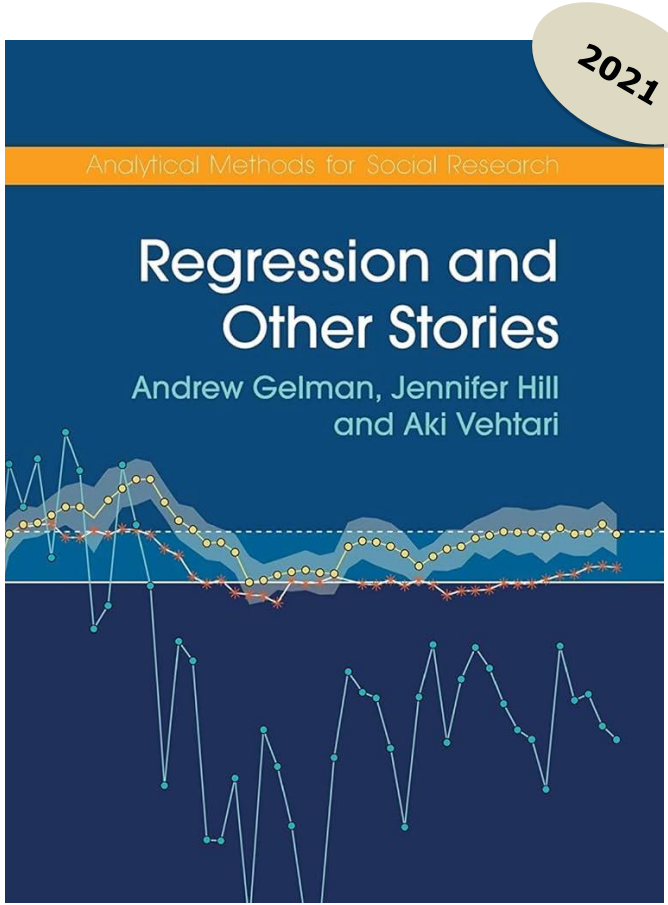

Bayesian books

Free online version available:
https://www.bayesrulesbook.com/

Uses „rstanarm“ and Stan

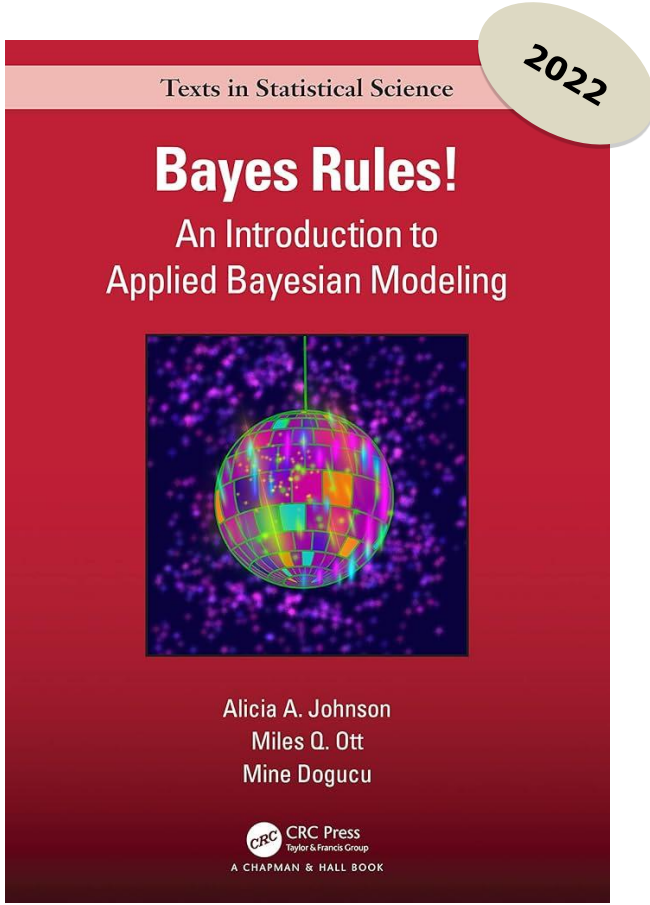

Bayesian books

Free online version available:
https://paulbuerkner.com/software/brms-book/

Work-in-progress, full version available this year

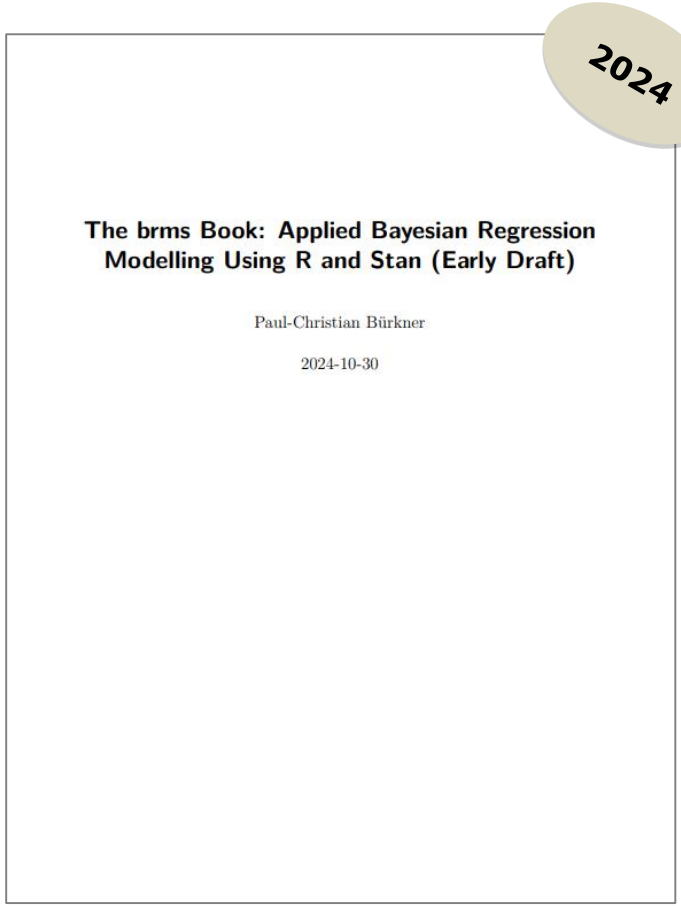

Ecology books

Purely frequentist

Still a classic for specific ecological problems
(mixed effects, autocorrelation, zero-inflation, …)

Many best-practice examples with step-by-step guide
of data analysis and statistical inference

Data package „AED“ discontinued, but use

```
> remotes::install_github("romunov/AED")
```

→ Try to replicate some analyses with brms

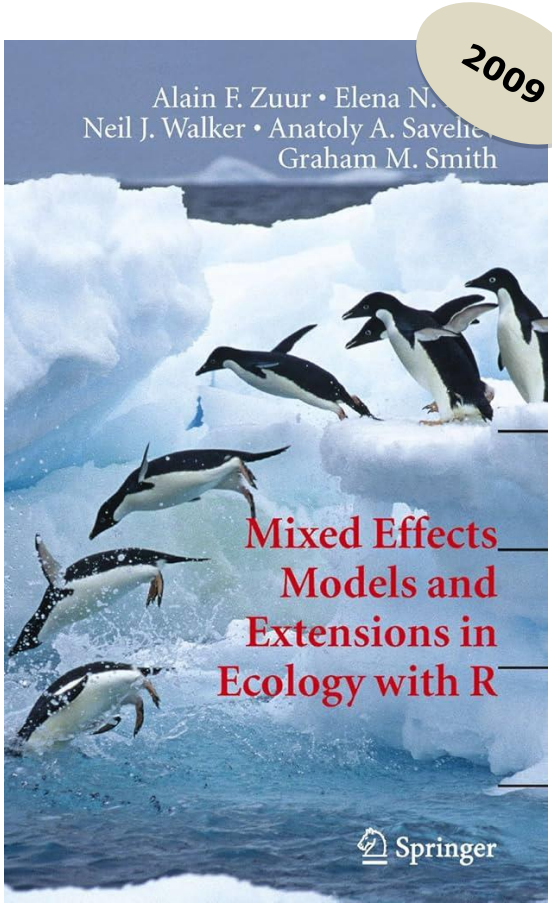

Ecology books

Purely frequentist

Half of the book dedicated to multivariate analysis!

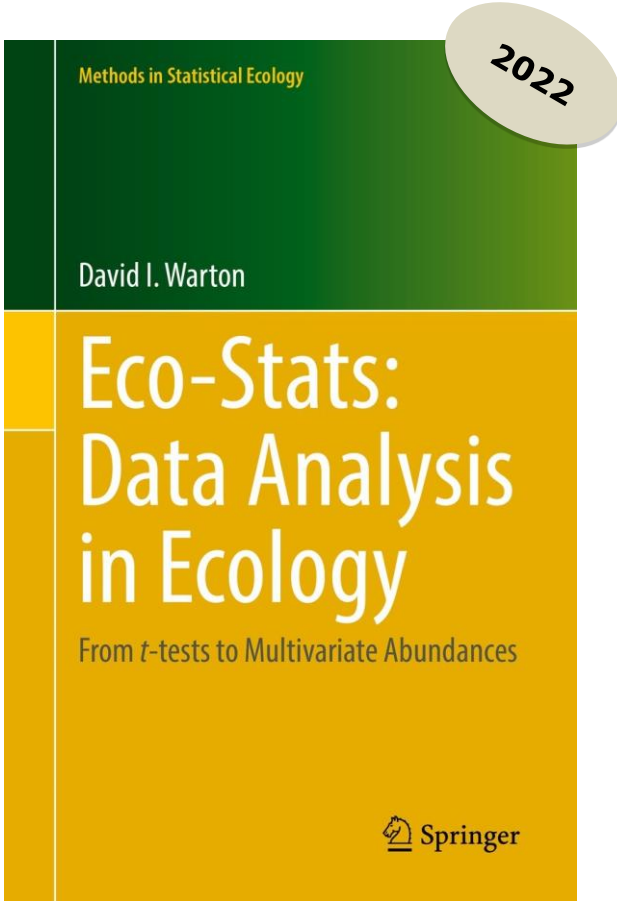

Ecology books

Free online version available:
https://statistics4ecologists-v3.netlify.app/

Mostly frequentist, just a brief Bayesian chapter

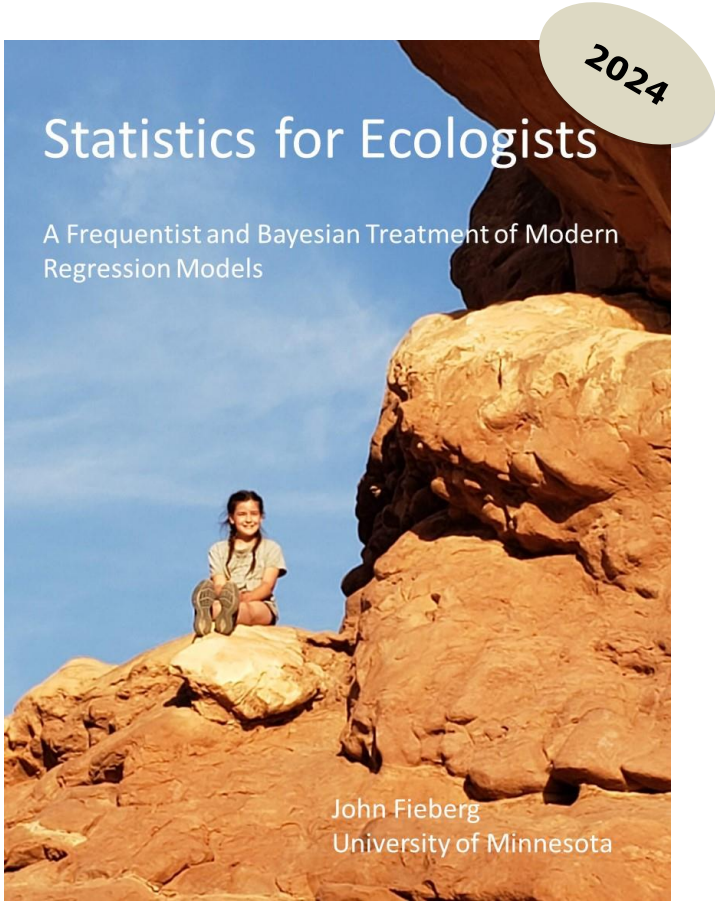

A Bayesian ecology-related book

A dual frequentist and Bayesian approach

Uses brms !

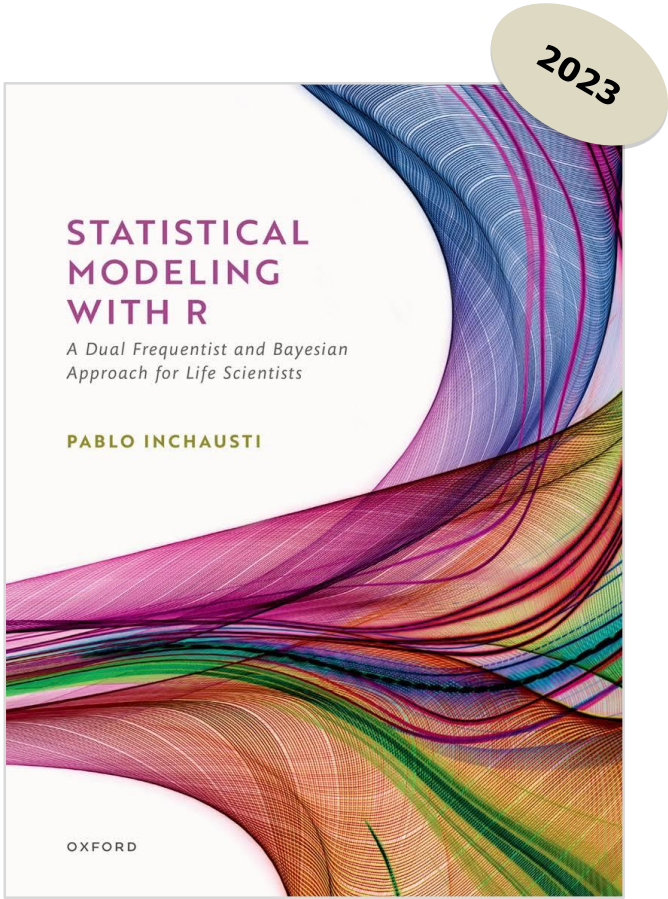

A Bayesian ecology book

Successor of Kery's 2010 WinBUGS book

Uses Stan, NIMBLE, etc, but no brms

If you're looking for an accessible & concise introduction to coding ecologically relevant models, this probably is the one.

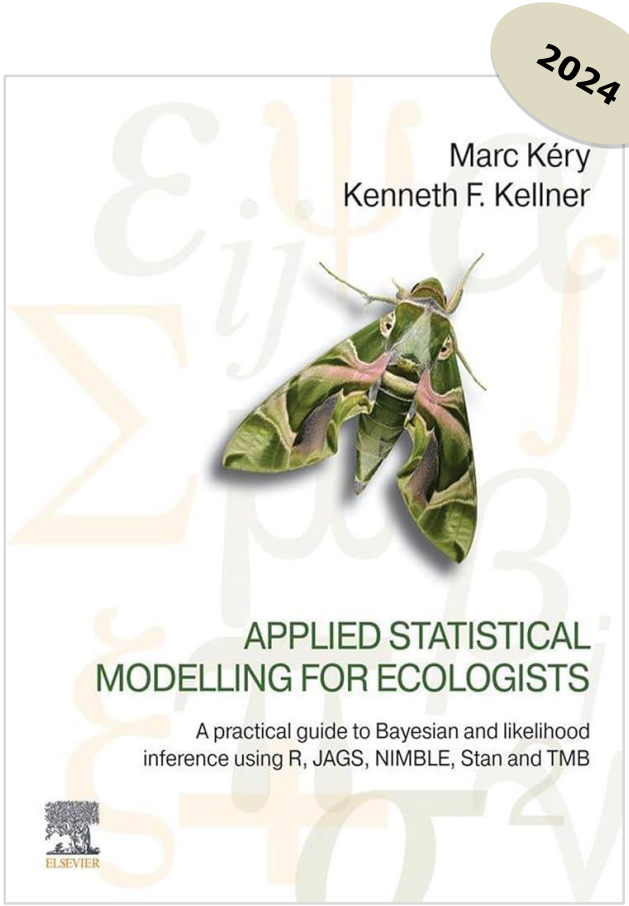

Online Resources

- Stan forums https://discourse.mc-stan.org/
- brms section https://discourse.mc-stan.org/c/interfaces/brms/36
- some tutorials: https://mc-stan.org/learn-stan/tutorials.html

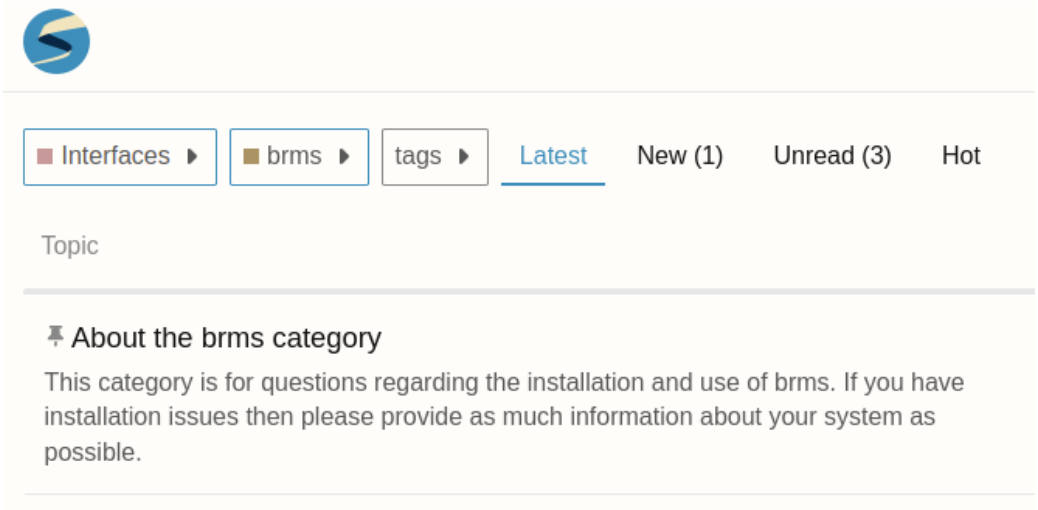

Take home:

Frequentist stats are often like:

Test *ABC* can do *XYZ*.

Bring hypotheses and data in alignment with

test‘s assumptions

Bayesian (& other!) stats should be like:

Model the data-generating process.

Derive something quantifiable (parameters or predictions)

from the model that describes you original research question.

The Bayesian 3D printer

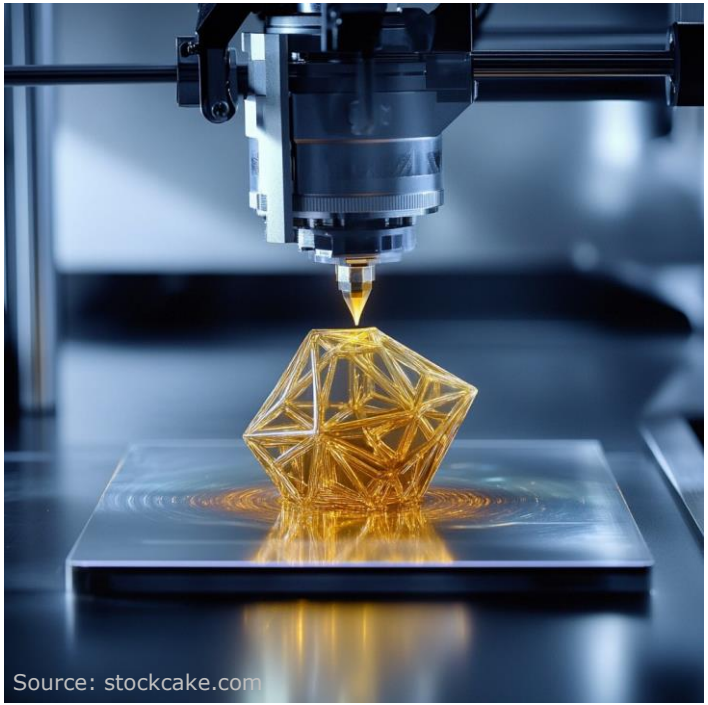

Good luck !