brms: Bayesian Multilevel Models using Stan

Paul Bürkner

2018-04-09

Why using Multilevel Models?

Example: Effects of Sleep Deprivation on Reaction Times

```
data("sleepstudy", package = "lme4")
head(sleepstudy, 10)
```

Reaction	Days	Subject
249.5600	0	308
258.7047	1	308
250.8006	2	308
321.4398	3	308
356.8519	4	308
414.6901	5	308
382.2038	6	308
290.1486	7	308
430.5853	8	308
466.3535	9	308

Linear Regression vs. Multilevel Regression

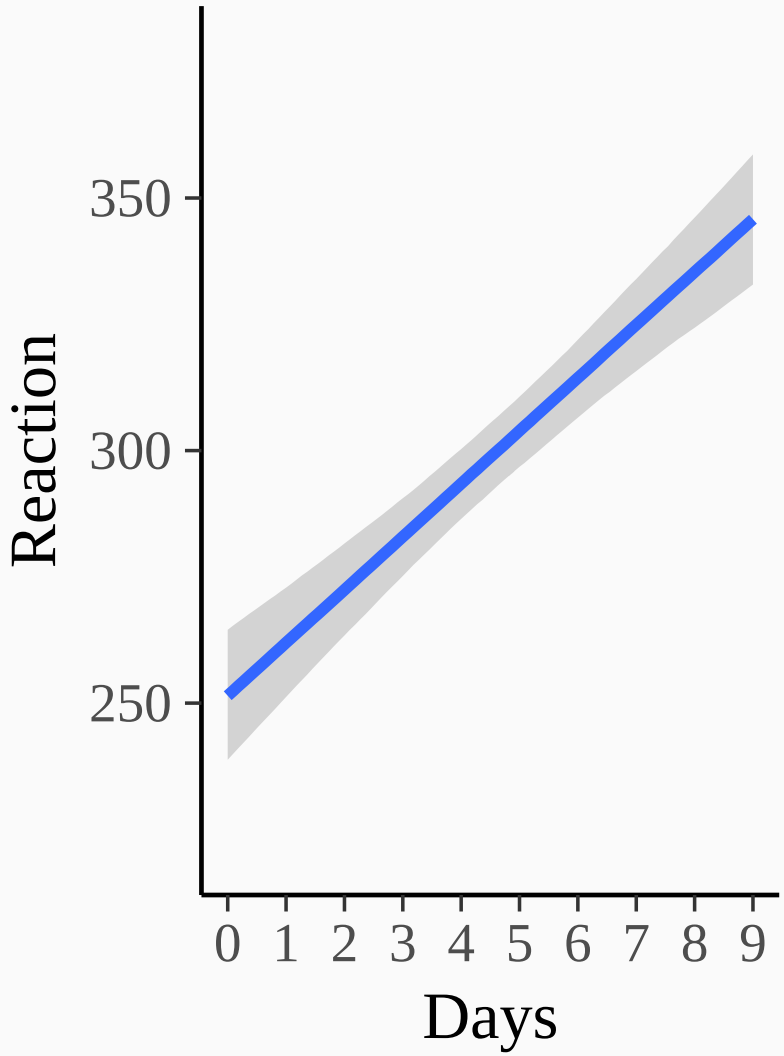

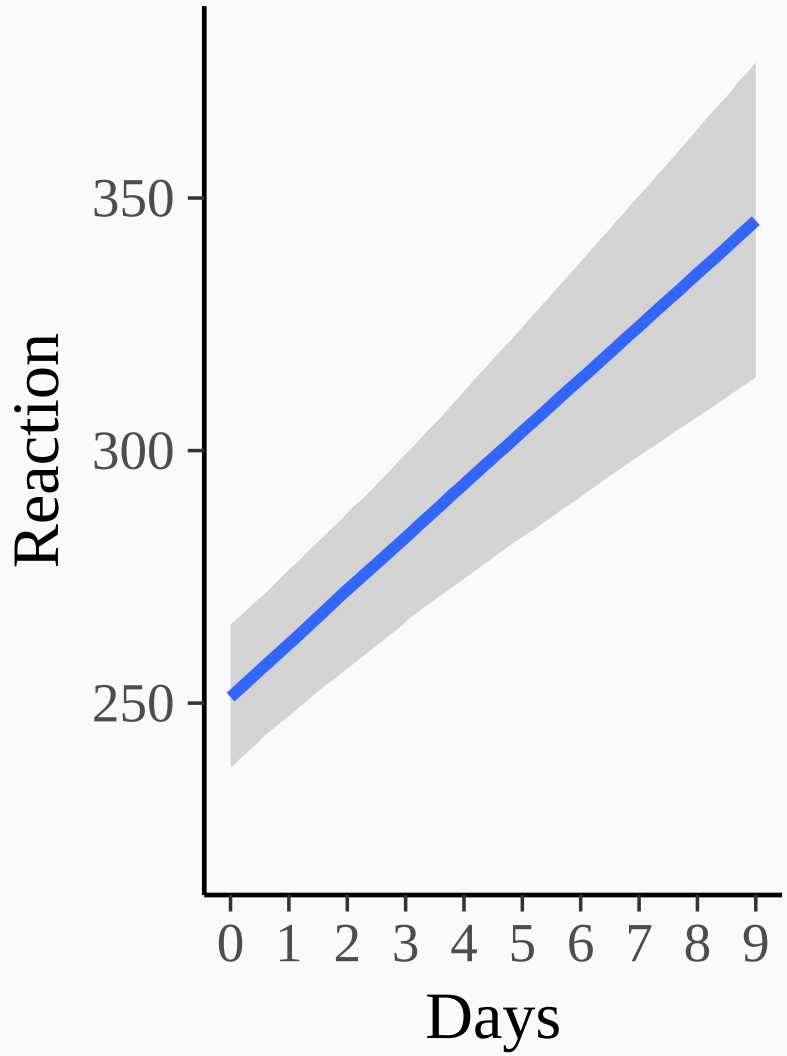

Regression Lines for Specific Subjects

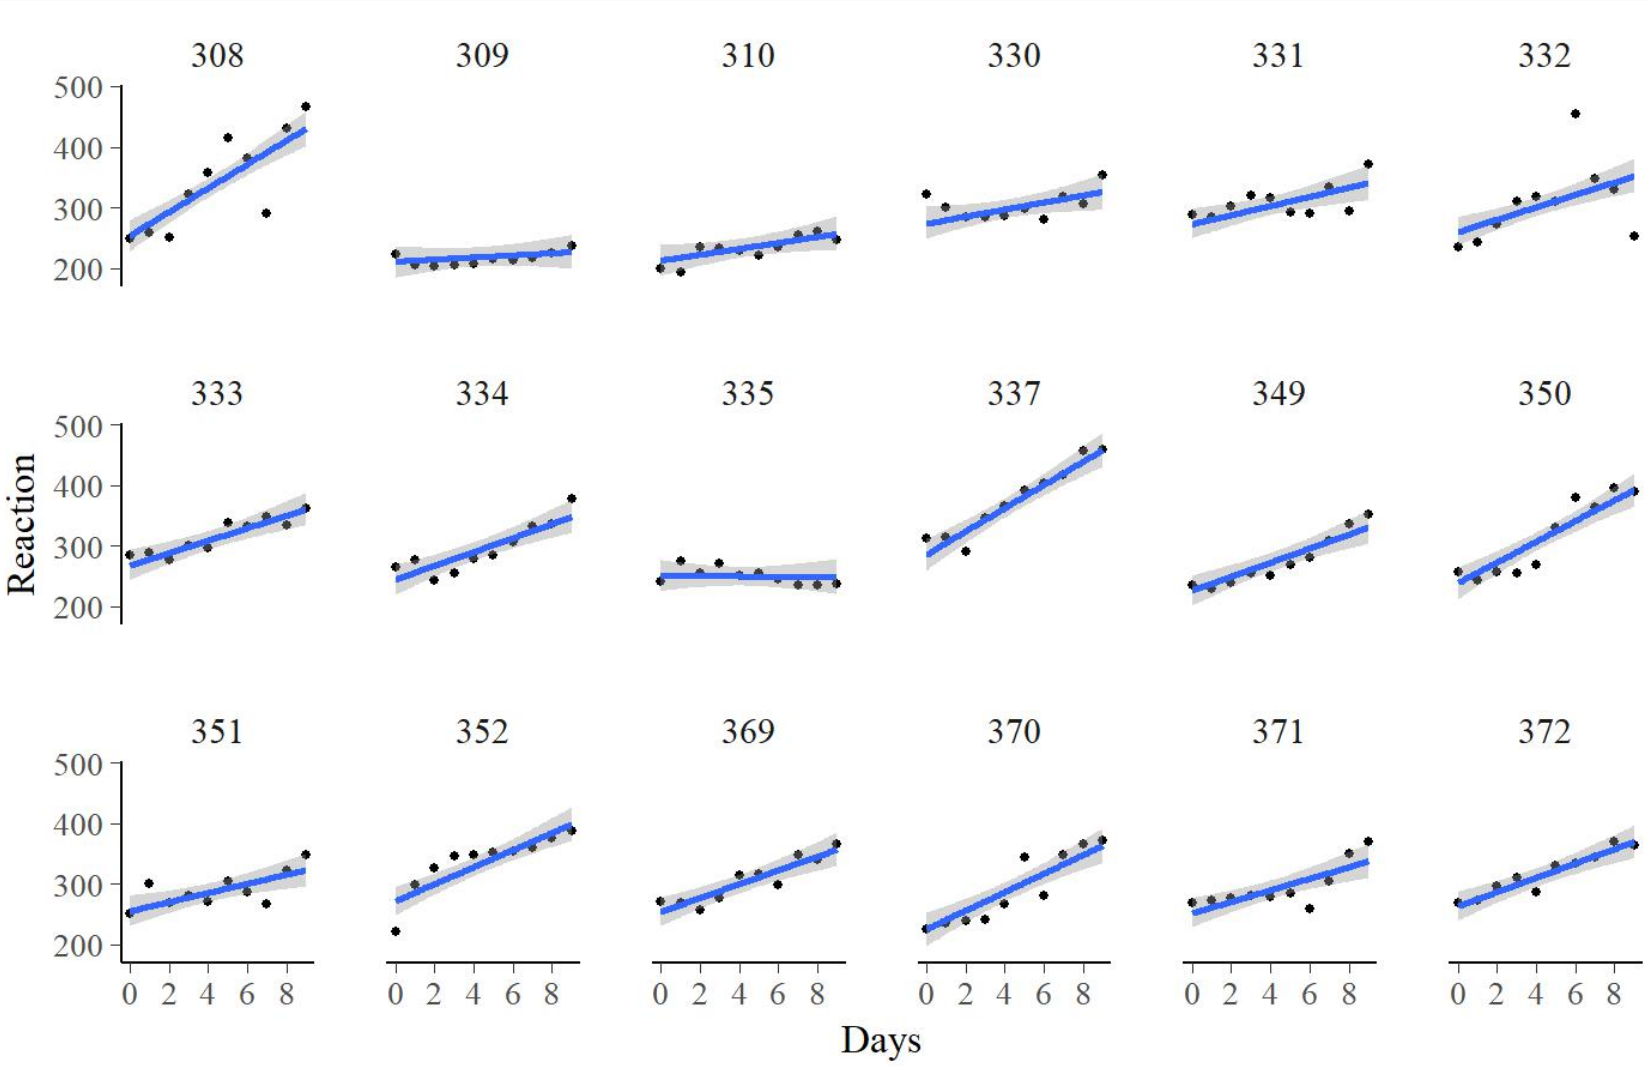

Why going Bayesian?

The Posterior Distribution

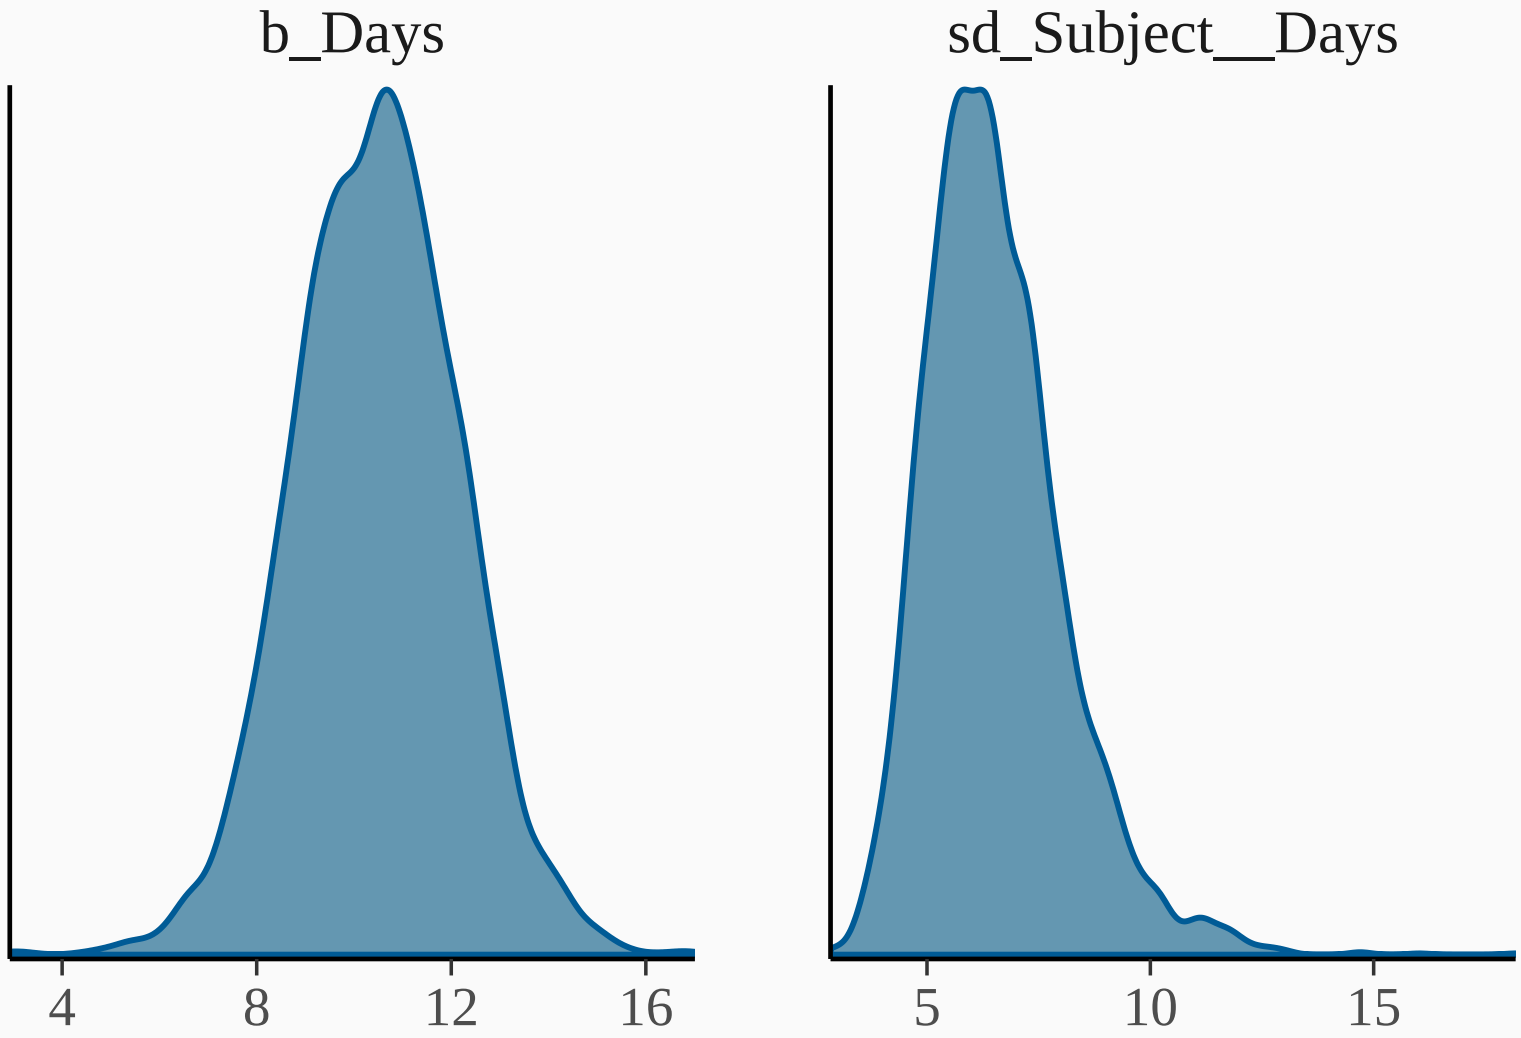

Why using Stan?

Multilevel Models in Classical Statistics with lme4

```
fit <- lmer(Reaction ~ 1 + Days + (1 + Days|Subject),
            data = sleepstudy)
```

```
fit <- brm(Reaction ~ 1 + Days + (1 + Days|Subject),
           data = sleepstudy)
```

Post-Processing

```
methods(class = "brmsfit")
```

```
##  [1] add_ic                  add_loo                 add_waic                as.array
##  [5] as.data.frame           as.matrix               as.mcmc                 coef
##  [9] control_params          expose_functions        family                  fitted
## [13] fixef                   formula                 hypothesis              kfold
## [17] launch_shiny            log_lik                 log_posterior           logLik
## [21] loo                     LOO                     loo_linpred             loo_predict
## [25] loo_predictive_interval marginal_effects        marginal_smooths        model.frame
## [29] neff_ratio              ngrps                   nobs                    nsamples
## [33] nuts_params             pairs                   parnames                plot
## [37] posterior_predict       posterior_samples       pp_check                pp_mixture
## [41] predict                 predictive_error        print                   prior_samples
## [45] prior_summary           ranef                   reloo                   residuals
## [49] rhat                    stancode                standata                stanplot
## [53] summary                 update                  VarCorr                 vcov
## [57] waic                    WAIC
## see '?methods' for accessing help and source code
```

The idea of **brms**: Fitting all kinds of regression models within one framework

Example: Censored Recurrance Times of Kidney Infections

```
fitk <- brm(time | cens(censored) ~
              age * sex + (1|patient),
            data = kidney, family = weibull())
```

```
marginal_effects(fitk, "age:sex")
```

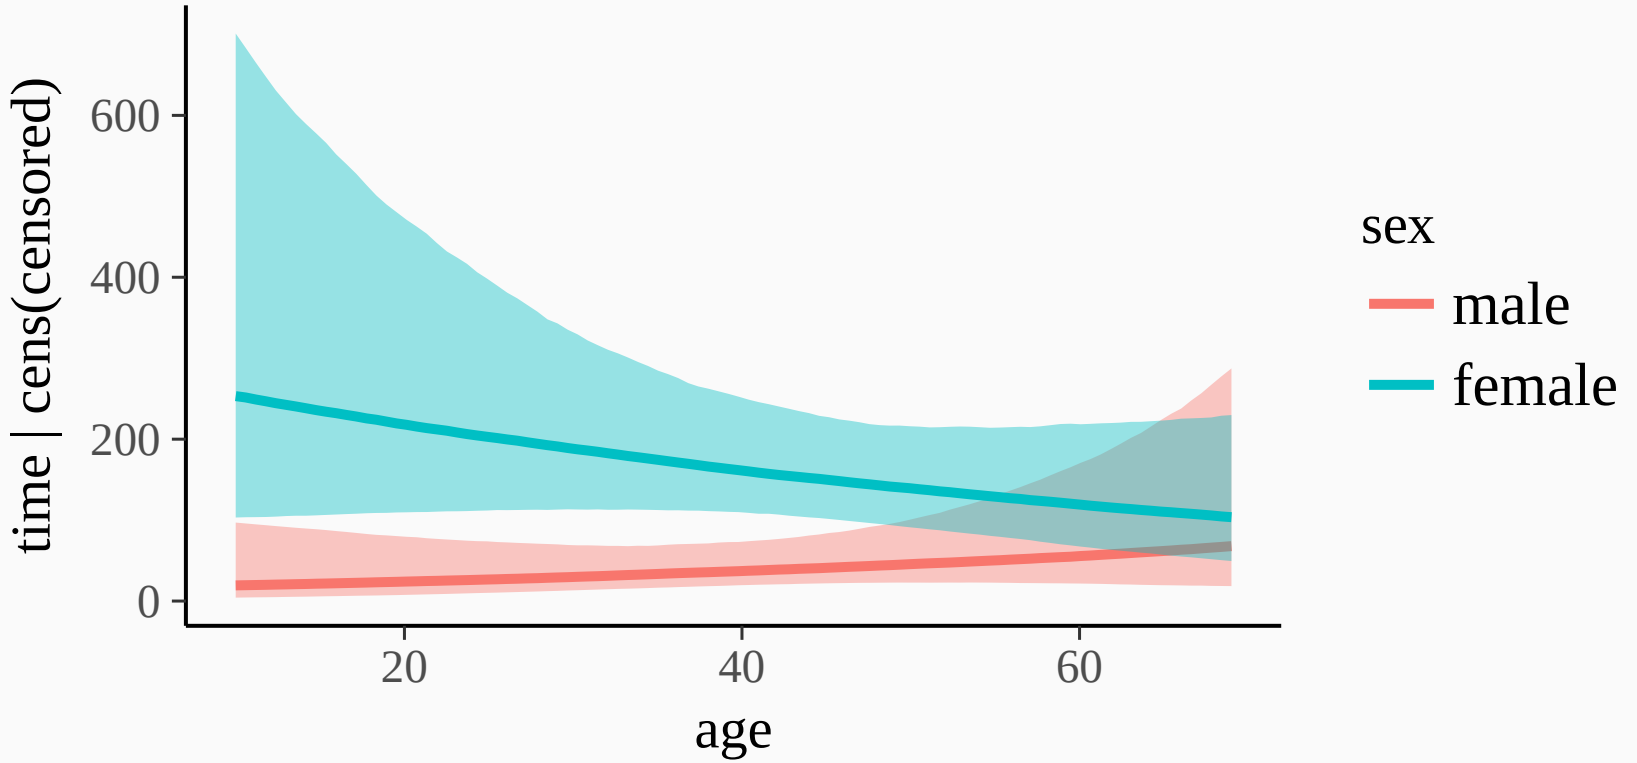

Example: Complex Non-Linear Relationships

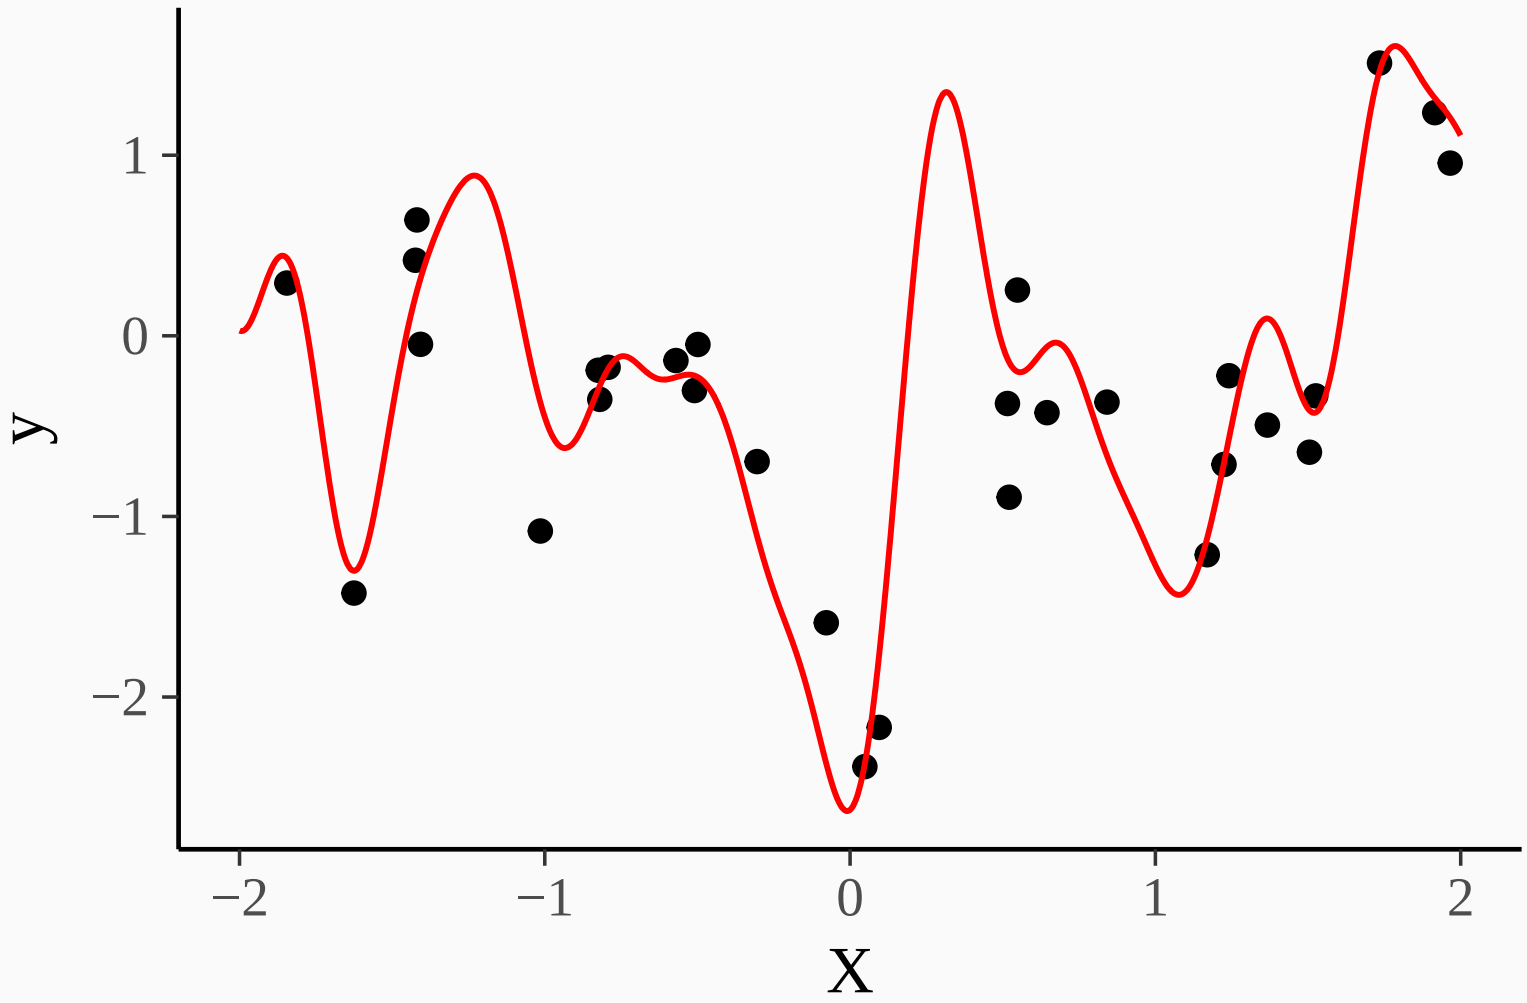

Modeling Non-Linear Relationships with Gaussian Processes

```
fitgp <- brm(y ~ gp(x), bdata)
```

```
marginal_effects(fitgp, nsamples = 100, spaghetti = TRUE)
```

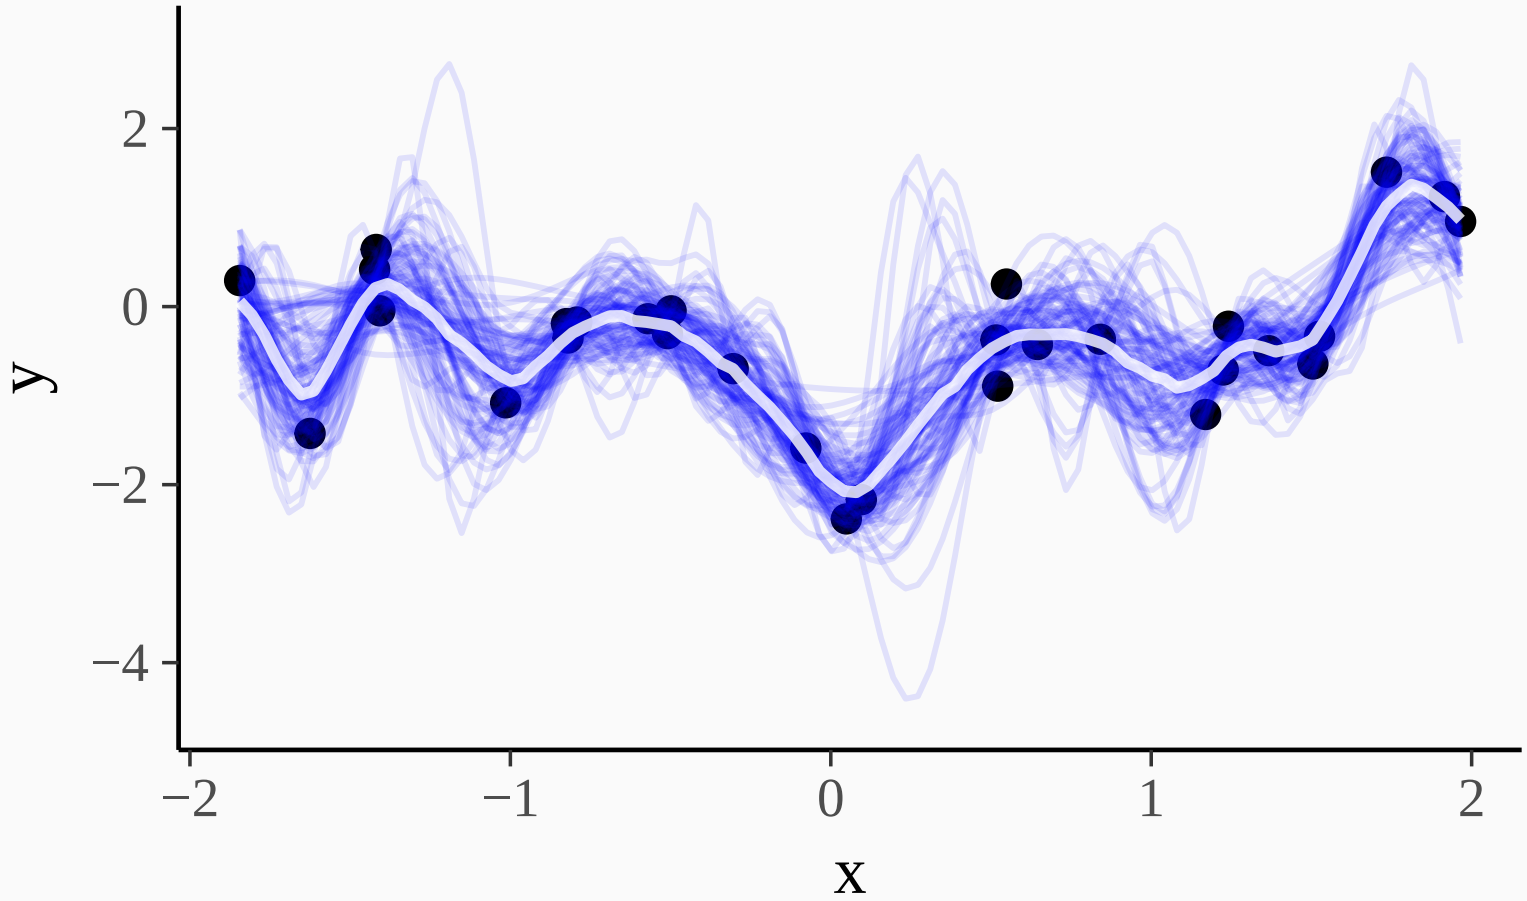

Modeling Non-Linear Relationships with Splines

```
fits <- brm(y ~ s(x), bdata)
```

```
marginal_effects(fits, nsamples = 100, spaghetti = TRUE)
```

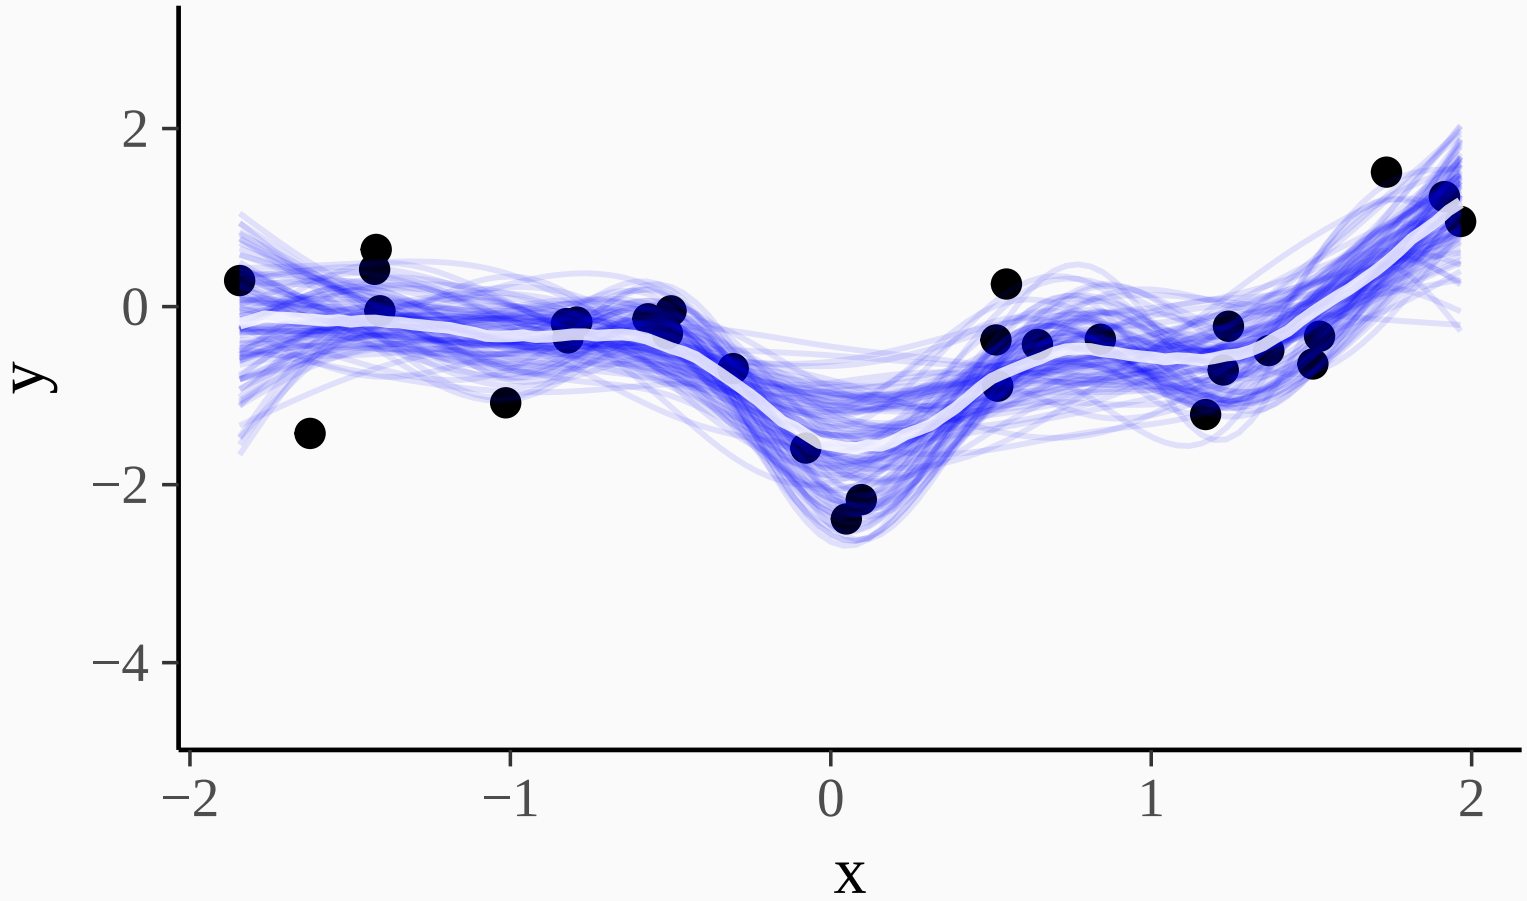

Example: Number of Fish Caught at a Camping Place

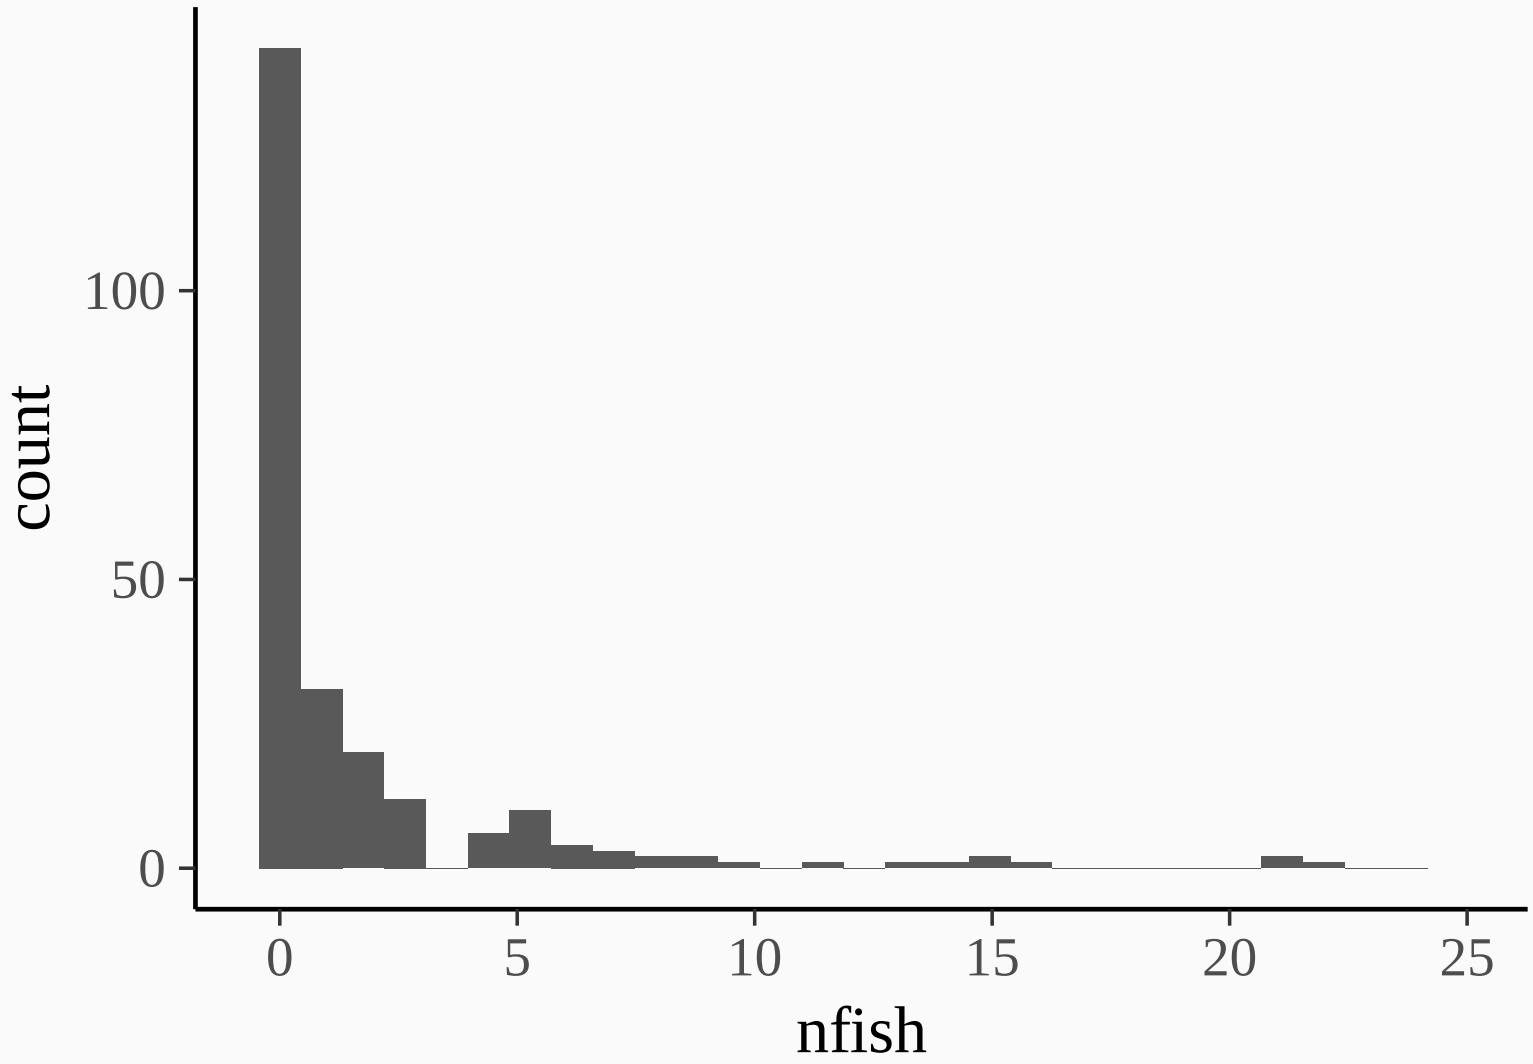

Modeling Zero-Inflation

```
form <- bf(nfish ~ persons + child + camper, zi ~ child)
fit_zinb <- brm(form, zinb, zero_inflated_poisson())
```

```
marginal_effects(fit_zinb, effects = "child")
```

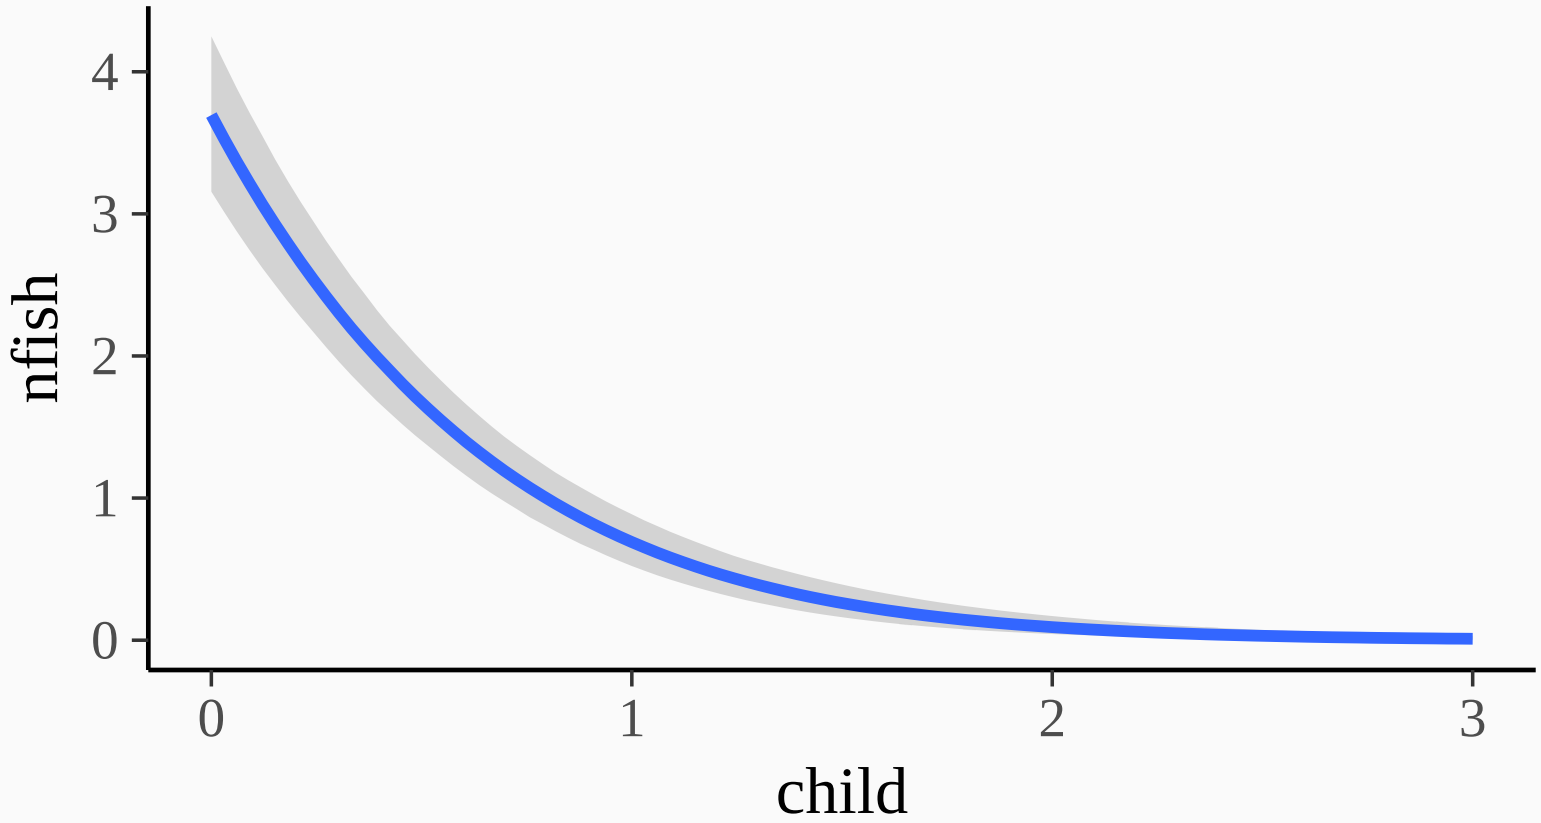

Learn More about brms and Stan

- Help within R: `help("brms")`
- Overview of vignettes: `vignette(package = "brms")`
- List of all methods: `methods(class = "brmsfit")`
- Website of brms: https://github.com/paul-buerkner/brms
- Website of Stan: http://mc-stan.org/
- Contact me: `paul.buerkner@gmail.com`
- Twitter: @paulbuerkner