Optimal Design and Bayesian Data Analysis

Paul-Christian Bürkner

2017-09-20

Optimal Design: Introduction

Simple example: Comparison of two independent groups

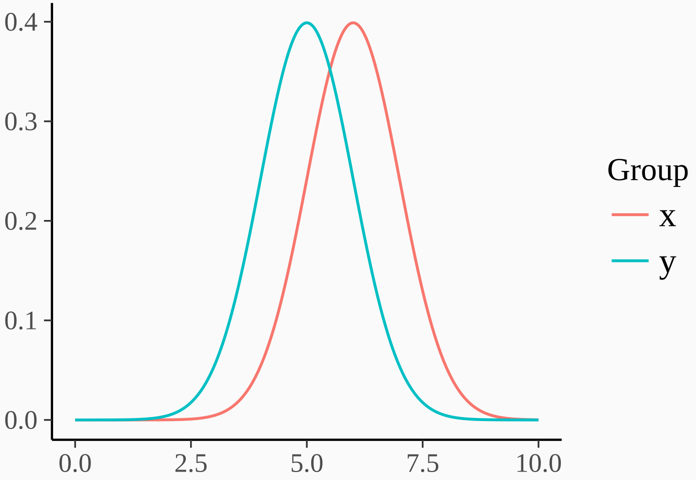

Optimal Design: Introduction

Comparison of two independent groups with unequal variances

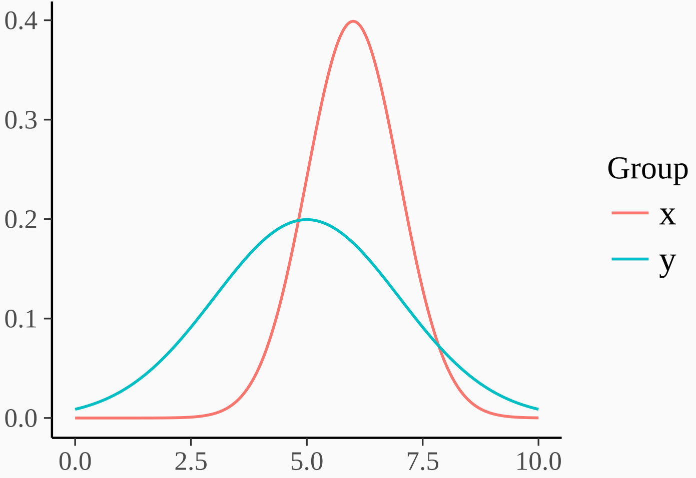

Optimal Design: Definition

Experimental Design:

- The experimental conditions
- The allocation of replications to the conditions

Optimal Experimental Design:

- The design ξ that optimizes a certain criterion function ψ
- For instance $\psi(\xi) := \text{Power}(T(\xi))$
- Alternatively $\psi(\xi) := \det(Cov(\theta, \xi))$

Optimal design of the Wilcoxon-Mann-Whitney-test

with Philipp Doebler and Heinz Holling

The Wilcoxon-Mann-Whitney-Test

- Assumption: Group x has continuous distribution F and group y has continuous distribution G.
- Hypotheses:

$$H_0 : G(x) = F(x) \quad \text{vs.} \quad H_1 : G(x) = F(x + a)$$

- Test Statistic:

$$U := \sum_{i=1}^{m} \sum_{j=1}^{n} \chi(x_i, y_j)$$

with $\chi(x_i, y_j) := 1$ if $x_i \geq y_j$ and $\chi(x_i, y_j) := 0$ if $x_i < y_j$.

- Experimental Design ξ: The proportion ω of replications allocated to group x.

Optimal Design of the U-Test: Symmetric Case

Theorem: Let the sample sizes be sufficiently large so that U is approximately normal. Then, for symmetric continuous distributions F and G with $G(x) = F(x + a)$ for some $a \neq 0$, the optimal design is given if $\omega^* = 0.5$.

Optimal Design of the U-Test: Asymmetric Case

Asymmetric distributions of x and y:

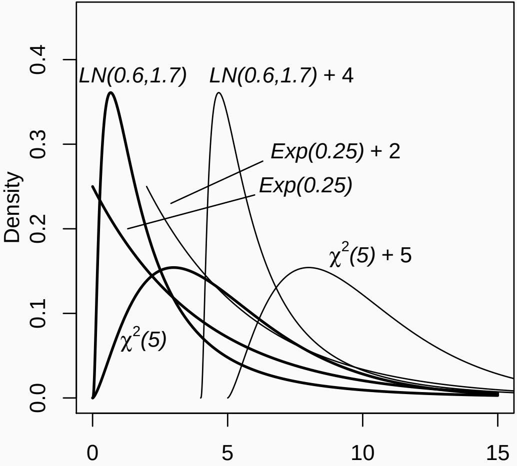

Optimal Design of the U-Test: Asymmetric Case

Power of the U-Test for asymmetric distribution:

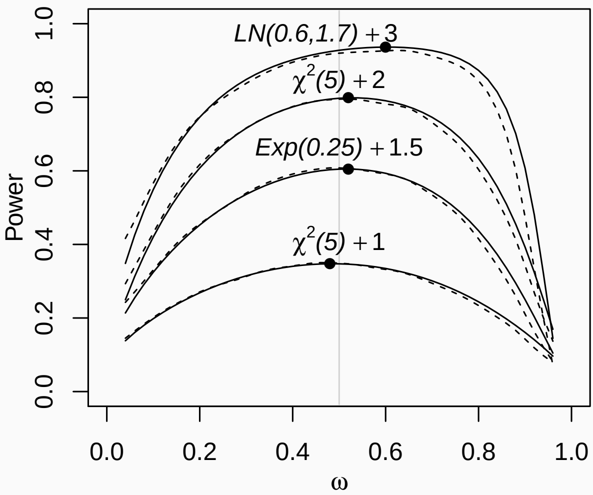

On the existence of optimal designs for the partial credits model

with Rainer Schwabe and Heinz Holling

The (Generalized) Partial Credits Model (PCM)

- A model for ordinal data in item response theory

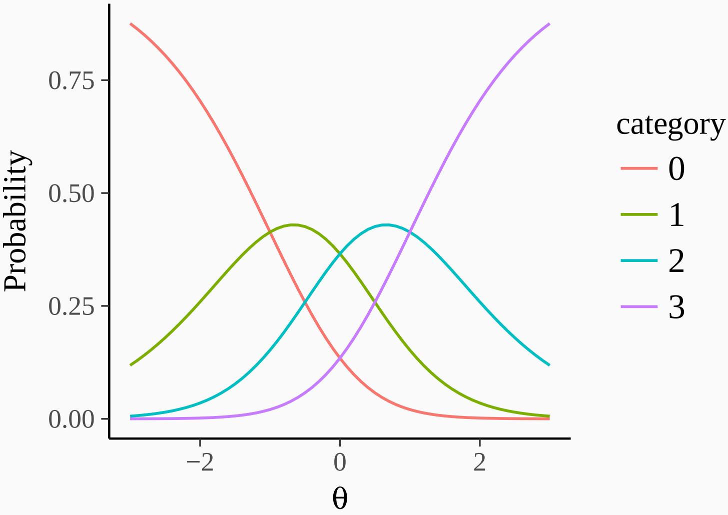

The (Generalized) Partial Credits Model (PCM)

- Assume an ordinal response with categories $j = 0, ..., J$

$$\pi_j := P(Y = j; \theta_p, \tau_i, \alpha_i) := \frac{\exp\left(\sum_{s=1}^{j} \alpha_{is}(\theta_p - \tau_{is})\right)}{\sum_{k=0}^{J} \exp\left(\sum_{s=1}^{k} \alpha_{is}(\theta_p - \tau_{is})\right)}$$

- θ_p = ability of person p
- τ_i = thresholds / difficulties of item i
- α_i = discriminations of item i

Locally optimal designs for the PCM

A design is (only) locally optimal if it is optimal for certain parameter values, but not for others.

Theorem: A local optimal design of the PCM must satisfy $\pi_0 = \pi_J = 1/2$ and $\pi_1 = ... = \pi_{J-1} = 0$. In this case, the PCM reduces to the ordinary 2PL model.

Theorem: The locally optimal design of the 2PL model is a one-point design for which $\tau_i = \theta_p$ and α_i as high as possible.

Bayes optimal designs

We assume a weight distribution Π over parameter values

A Bayesian design criterion could, for instance, look like

$$\psi(\xi) = \int \mathrm{det}(\mathrm{Cov}(\theta, \xi))\ \Pi(\theta)\ \mathrm{d}\theta$$

Π might also be called **prior** distribution

Bayes optimal designs for the PCM

Lemma: If the weight distribution Π is symmetric around some ability θ_0, α is fixed to any value, the Bayes optimal one-point design is the locally optimal design for θ_0.

Theorem: If the weight distribution Π is symmetric around some ability θ_0, α is fixed to any value, the Bayes optimal one-point design is Bayes optimal if the scale parameter s of Π does not exceed a certain value $s^*(\alpha) > 0$.

Visualization: Bayes optimal designs for the PCM

Values greater 0 indicate non-optimality of the one-point design.

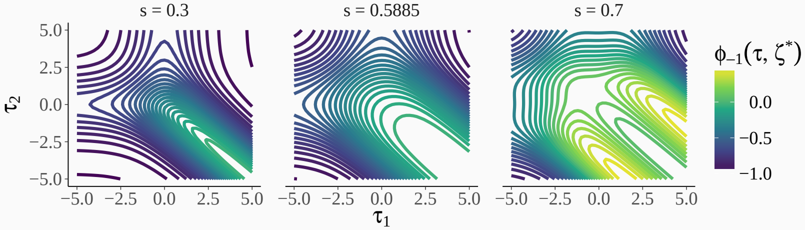

brms: An R package for Bayesian multilevel models using Stan

with a lot of coffee

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

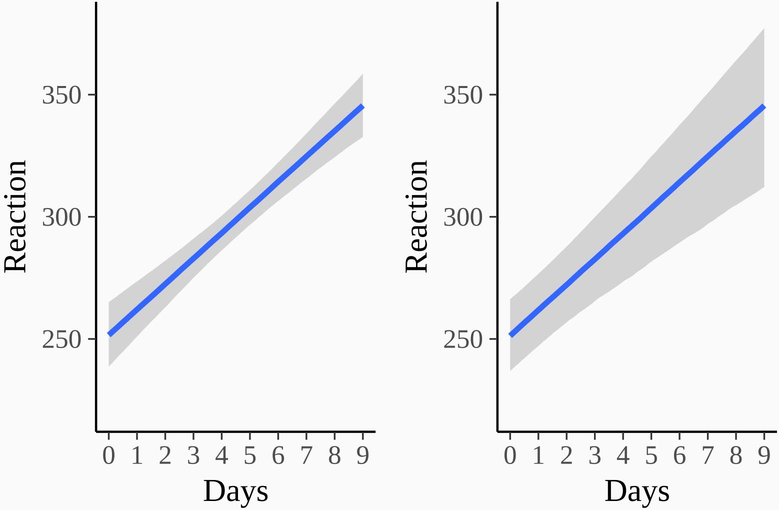

Regression Lines for Specific Subjects

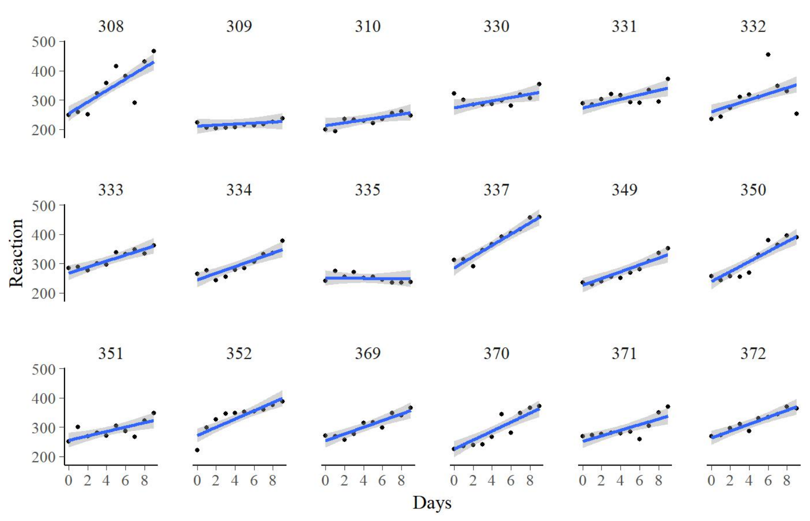

The Posterior Distribution

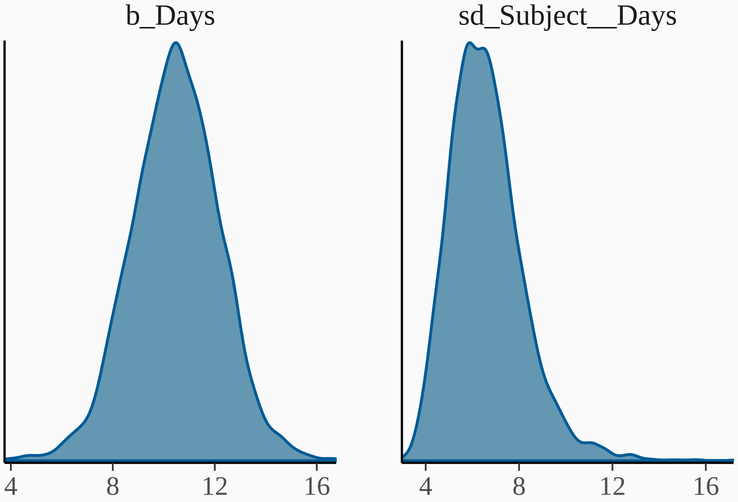

Multilevel Models in Bayesian Statistics with brms

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

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")
```

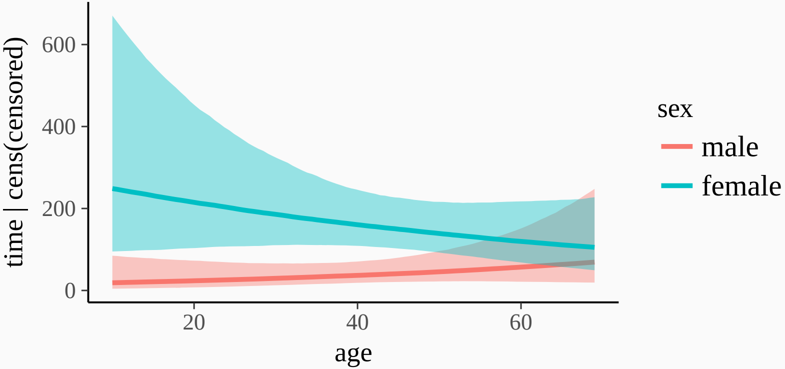

Example: Complex Non-Linear Relationships

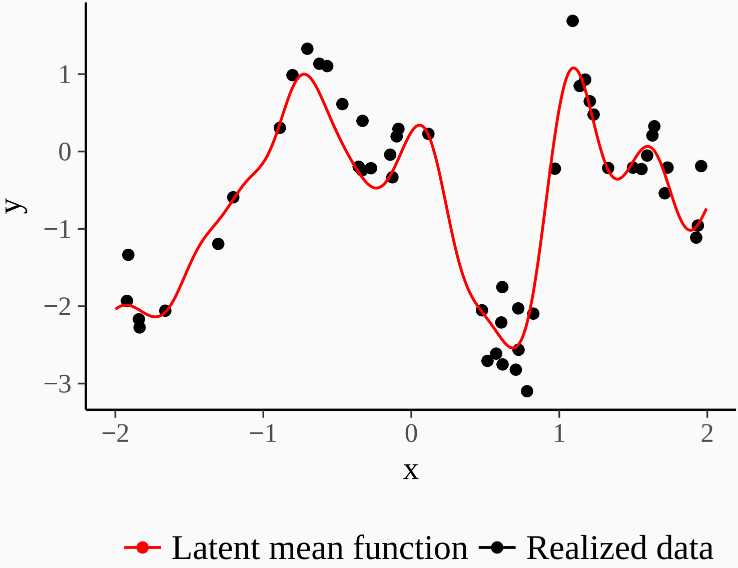

Modeling Non-Linear Relationships with Splines

```
fits <- brm(y ~ s(x), bdata, chains = 2)
```

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

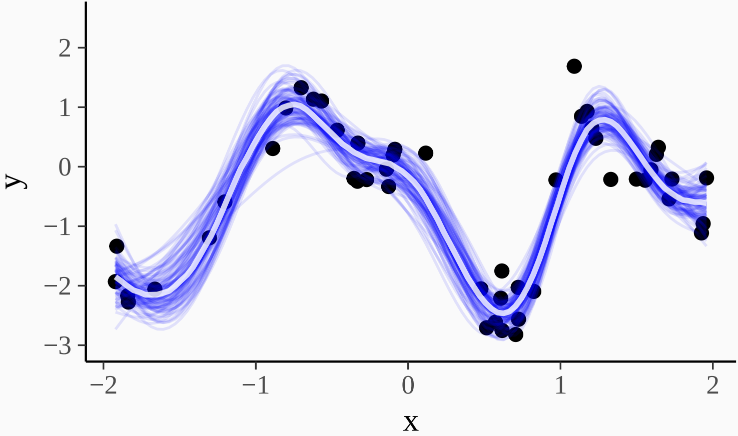

Example: Number of Fish Caught at a Camping Place

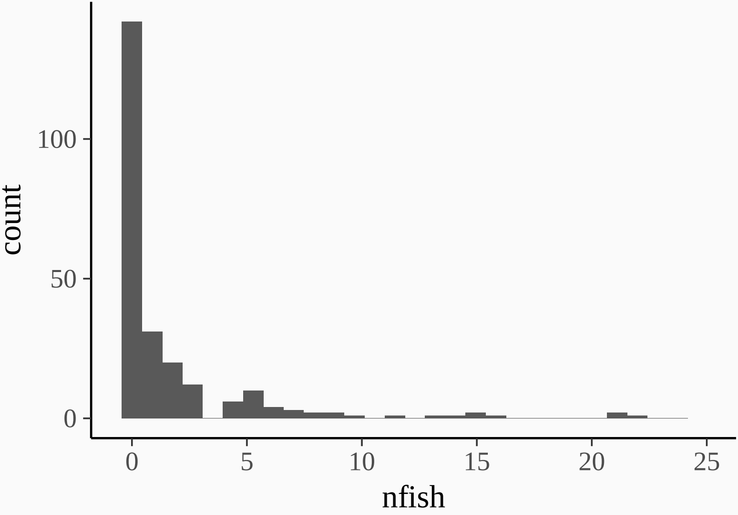

Modeling Zero-Inflation

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

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

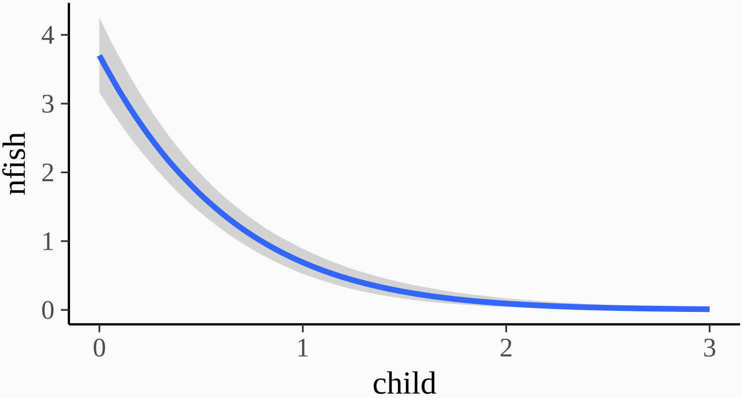

Thank you for your attention!