A Bayesian Feature-Sensitivity Analysis

Morten H. Falk, Håvard Rue and Andrea Riebler
Department of Mathematical Sciences, NTNU - Trondheim

NTNU – Trondheim
Norwegian University of
Science and Technology

Motivation

- In literature Bayesian sensitivity analysis has been highly advocated for decades with the rise of computing power.
- In practice Bayesian sensitivity analysis is rarely, if ever, done;
  - No "*plug 'n' play*" ™ sensitivity analysis available,
  - No globally applicable sensitivity measure for all models.
- We want to introduce a default sensitivity analysis for Bayesian Gaussian analysis within `INLA`.

Setup

- Suppose we consider iid. Gaussian data $X \sim \mathrm{N}(\mu, \sigma^2)$ with iid. skew-Gaussian residuals $\epsilon \in \mathrm{SN}(0, {\sigma_\epsilon}^2, \tau)$ with skewness $\tau \in [-1, 1]$.
- Consider two models: base model $\epsilon \in \mathrm{N}(0, {\sigma_\epsilon}^2)$ corresponding to $\tau = 0$ and flexible model $\epsilon \in \mathrm{SN}(0, {\sigma_\epsilon}^2, \tau)$.
- A PC-prior is an exponential distribution on distance $d(f||g) = \sqrt{2 \cdot \mathrm{KLD}(f||g)}$ scale, where $\mathrm{KLD}(f||g)$ denotes the Kullback-Leibler Divergence between flexible model f and base g.
- By default we elicit a PC-prior for the skewness τ. Depends on *user specification*:

$$\Pr(|\tau| > 0.5) = a.$$

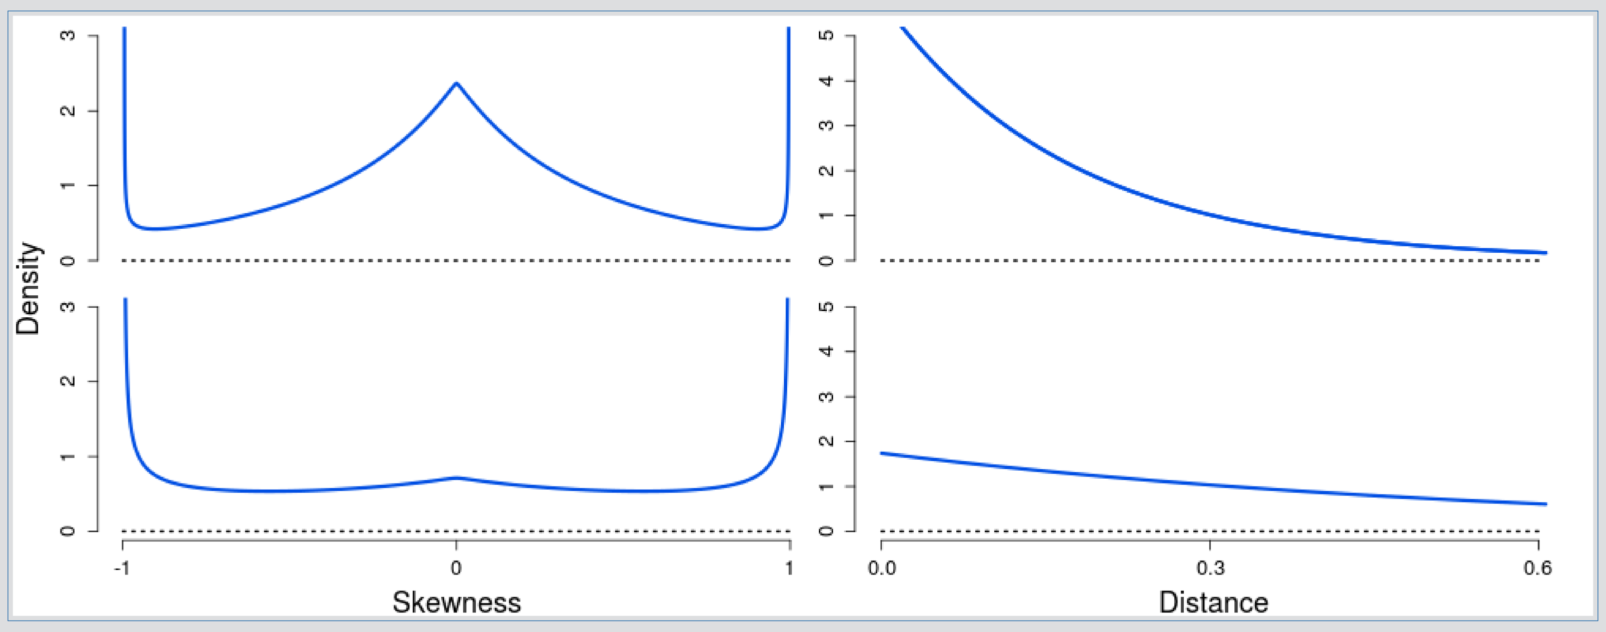

Figure 1: PC-priors. Top: $a = 0.3$. Bottom: $a = 0.7$.

A Metric of Sensitivity

- Assuming a linear model $Y_i = \beta_0 + \beta_1 \cdot X_i + \epsilon$, $i = 1, \ldots, n$, we use `R`-package `INLA` to compute posterior distributions of:
  - Intercept parameter β_0,
  - Slope parameter β_1,
  - Residual precision $\gamma = \frac{1}{{\sigma_\epsilon}^2}$ with assumed known prior,

  for both base and flexible models.
- For $\beta_0 = \mu = 0$, $\beta_1 = \sigma^2 = {\sigma_\epsilon}^2 = 1$, $\tau = 0.5$, $a = 0.3$ and $n = 500$:

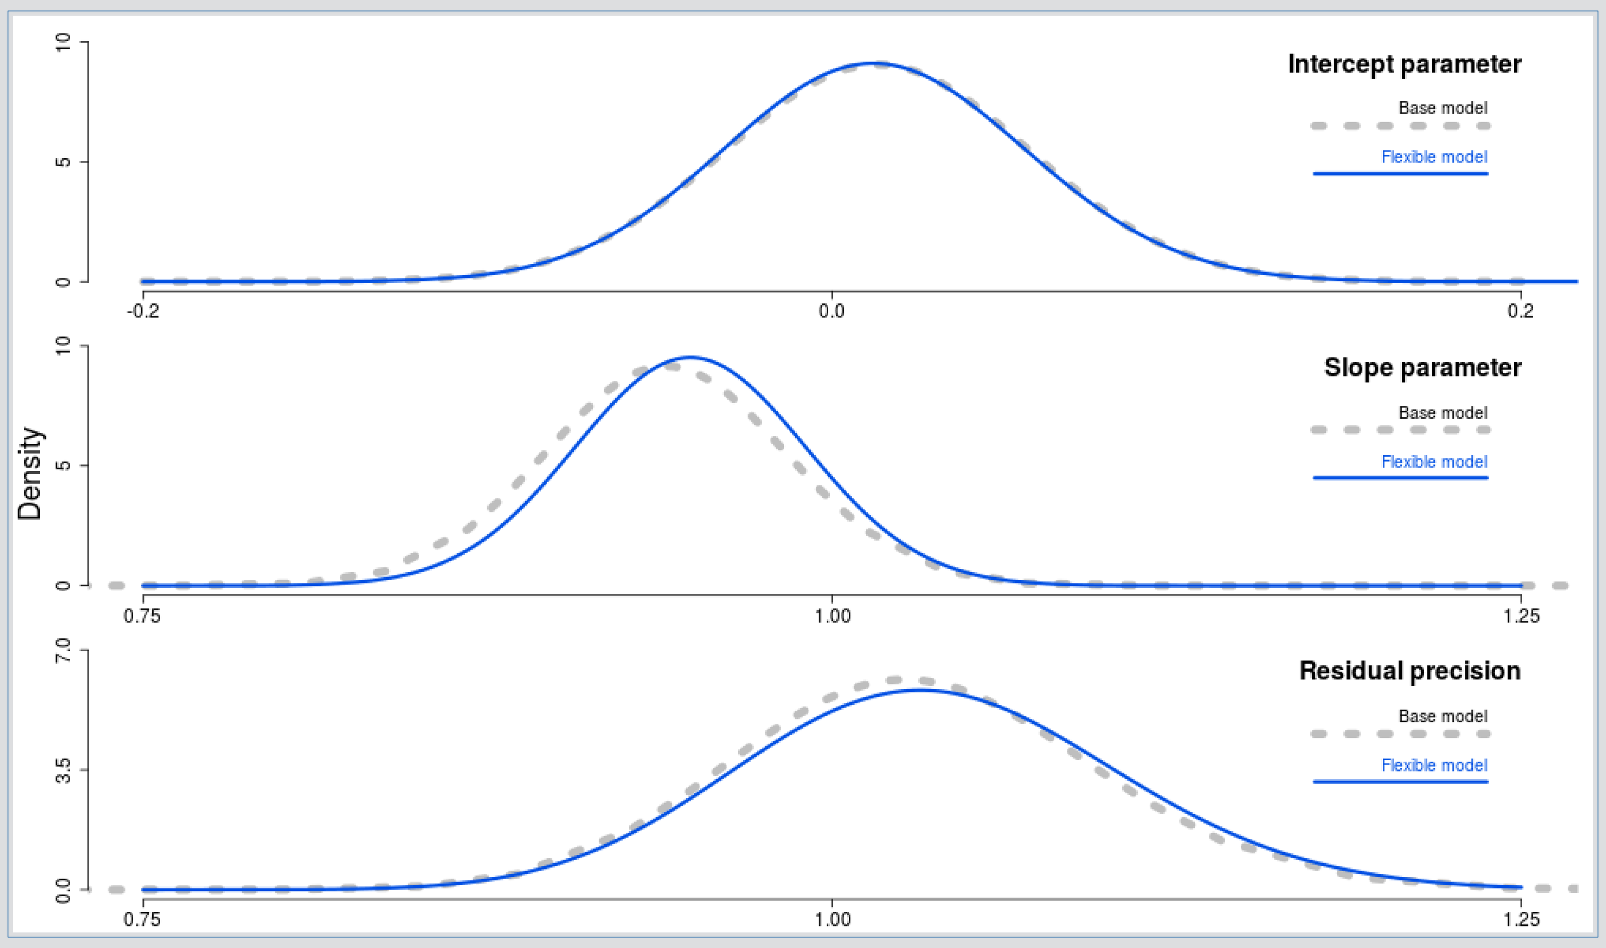

Figure 2: Parameter posteriors.

- Note that it can be shown that fixing a, fixes d(marginal likelihoods).
- Parameter sensitivity $\mathcal{F}_d(\cdot)$ is then computed as the quotient

$$\mathcal{F}_d(\cdot) = \frac{d(\text{posteriors})}{d(\text{marginal likelihoods})}.$$

Sensitivity Results

- For a cumulative set of standard Gaussian data X and residuals $\epsilon \sim \mathrm{SN}(0, 1, \tau)$ we observe the sensitivity $\mathcal{F}_d(\cdot)$ assuming τ is distributed according to the PC-prior with $a = 0.3$.

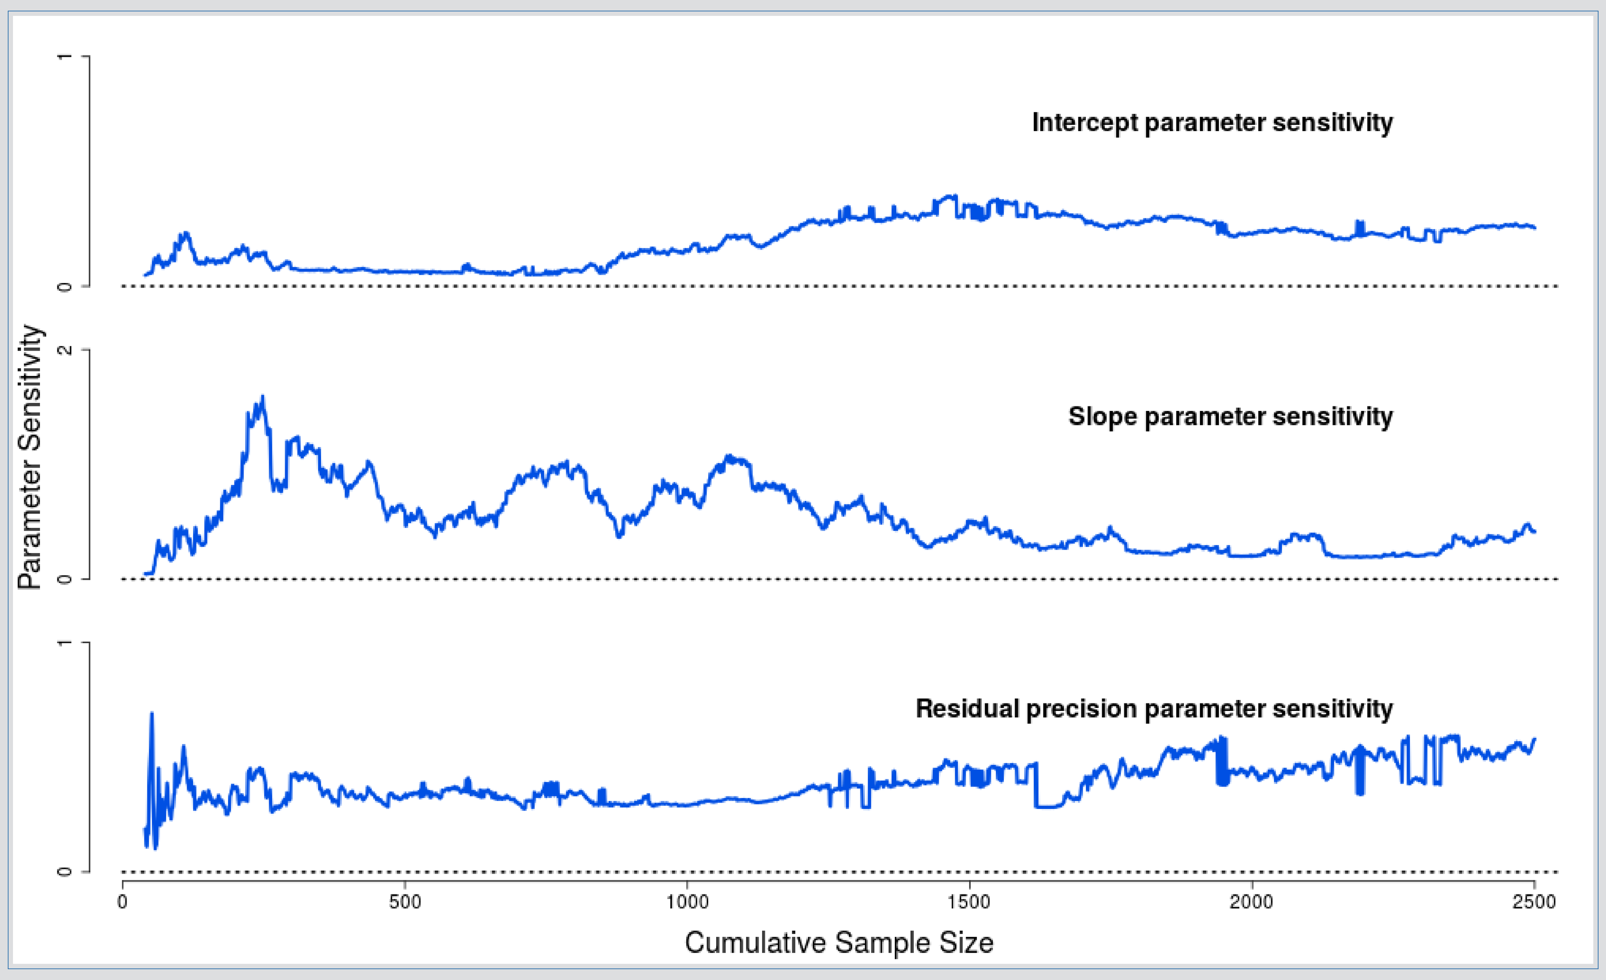

Figure 3: Sensitivity study. $\tau = 0.5$.

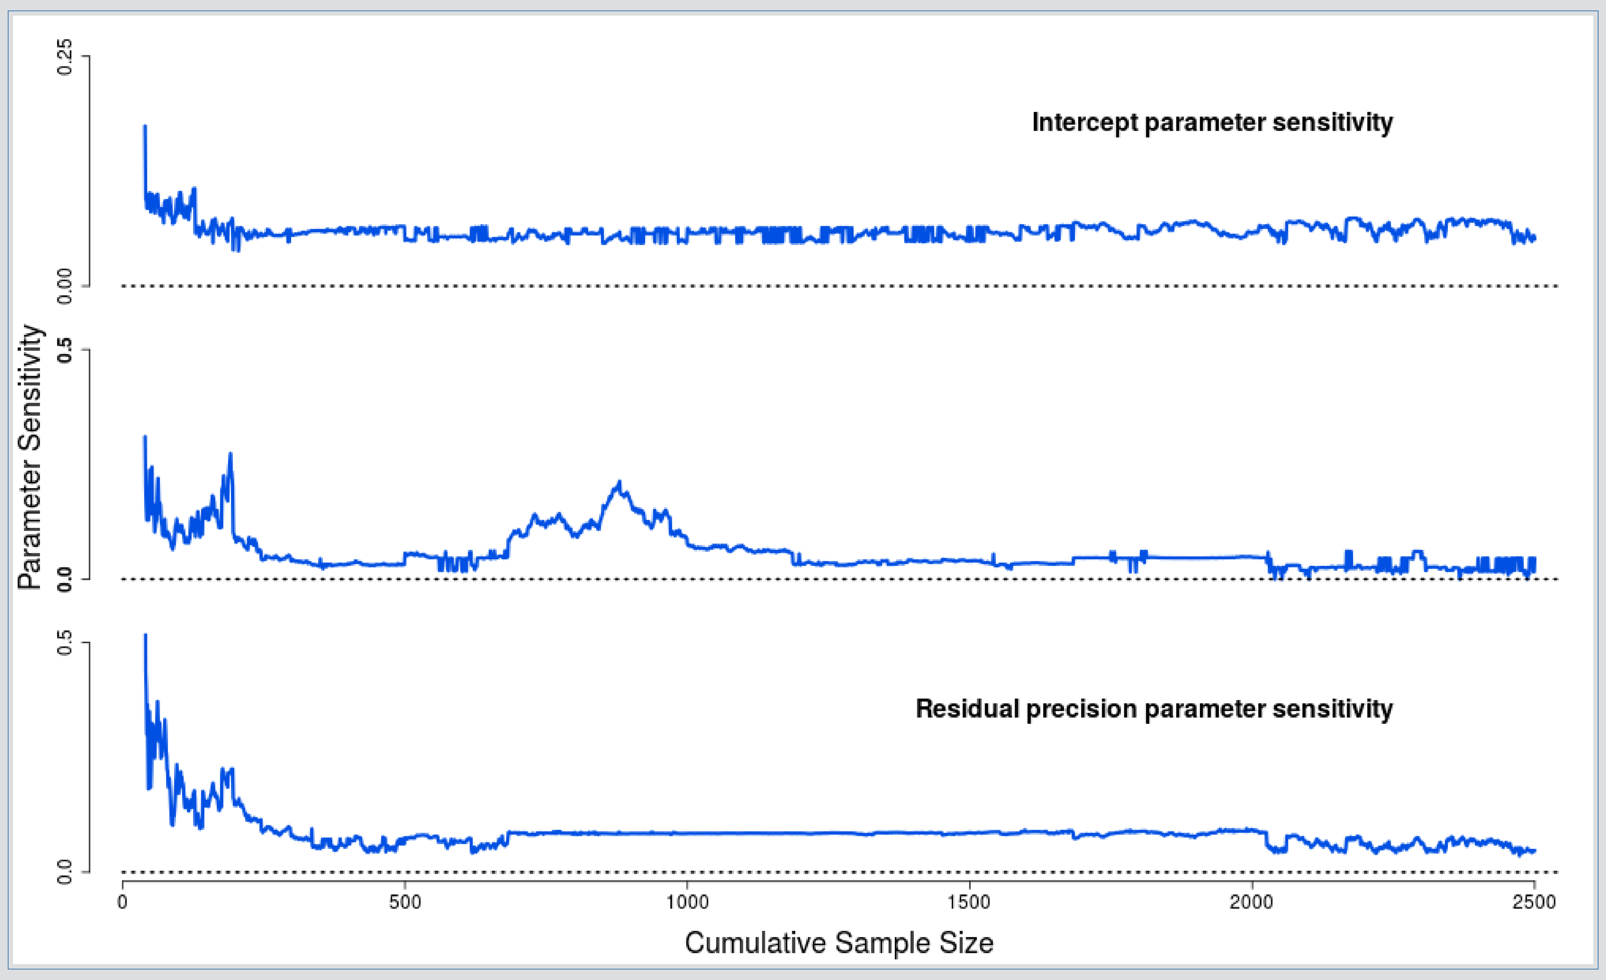

Figure 4: Sensitivity study. $\tau = 0$.

- The sensitivity measure appears to stabilise over sampling size.
- For $\tau = 0$, the sensitivity measures appear to stabilise close to 0.

Moving Forward

- Derive the rate of convergence of the sensitivity measure $\mathcal{F}_d(\cdot)$,
- Use a multivariate Gaussian flexible model to measure observation dependence sensitivity,
- Construct a flexible model to account for- and measure over-dispersion sensitivity,
- Generalise the skew-Gaussian flexible model to the skew-Student-t distribution.

References

[Berger (1994)] *An overview of robust Bayesian analysis [with Discussion].*, Berger, J. 1994, *Test*, vol. 3 pp. 5–124.

[Simpson et al. (2015)] *Penalising model component complexity: A principled, practical approach to constructing priors*, Simpson, D. P., Martins, T. G., Riebler, A., et al. 2015, arXiv:1403.4630.

Acknowledgments

▶ The author acknowledges the ongoing funding by '*Rejselegat for Matematikere*'.

morten.holm.falk@math.ntnu.no