Estimating the Extremal Index and Its Uncertainty

Jolan Bonny, Prof. Anthony Davison

Motivation

Extreme events, not averages, often determine risk in hydrology, finance and engineering.
In Switzerland, infrastructures such as the Grande Dixence dam must withstand rare, intense floods.
Extreme Value Theory provides a framework to estimate the probability of such extremes and to design robust protections.

Data & Clustering

To study extremes, we place a high threshold on a time series and focus on exceedances above it. If data are independent, exceedances also occur independently and standard methods apply. But real-world processes are often dependent: heavy rain today increases the chance of heavy rain tomorrow. Exceedances then form clusters. To describe and model this clustering, we introduce the extremal index θ, which can be interpreted as the inverse of the mean cluster size.

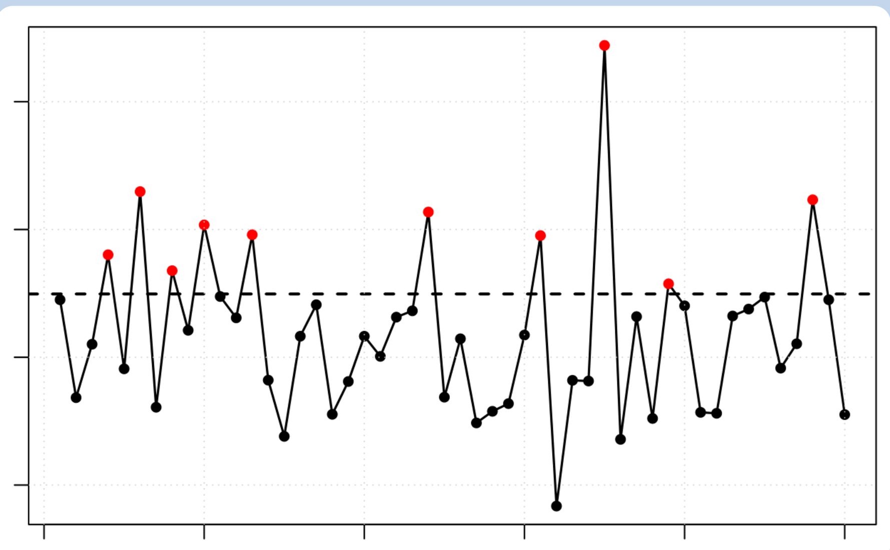

Artificially generated independent process

→ Exceedances occur randomly and separately above the threshold.

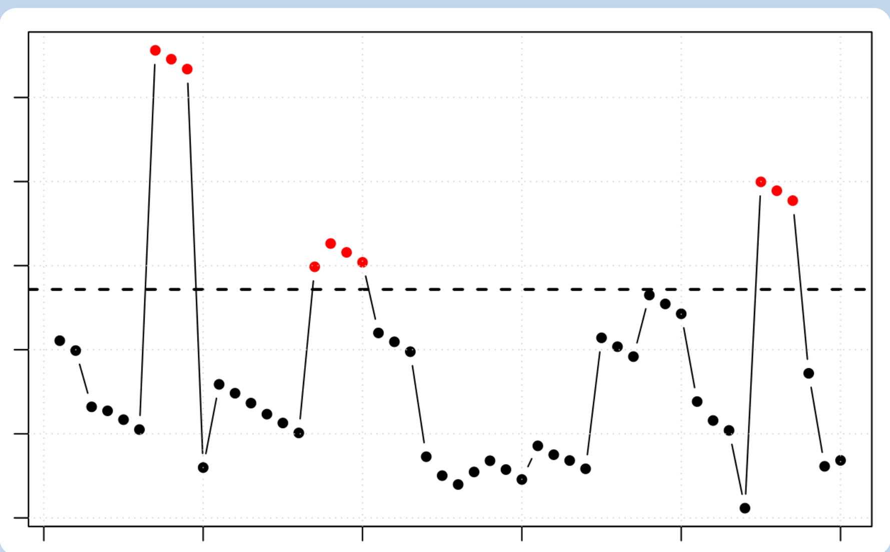

Artificially generated dependent process

→ Exceedances come in groups of 3–4 successive points, suggesting clustered extremes and $\theta \approx 1/3 – 1/4$

Here, the threshold is set at the 0.8 quantile, meaning that only the highest 20 % of values are considered. These exceedances are shown in red.

Estimation & Challenges

Estimators are formulas that convert data into an approximation of the extremal index and rely on asymptotic theory, valid as sample size tends to infinity.
With finite samples they can be biased or unstable, so **quantifying uncertainty is essential**.
In this work we use the 2022 estimator of Holešovský & Fusek.

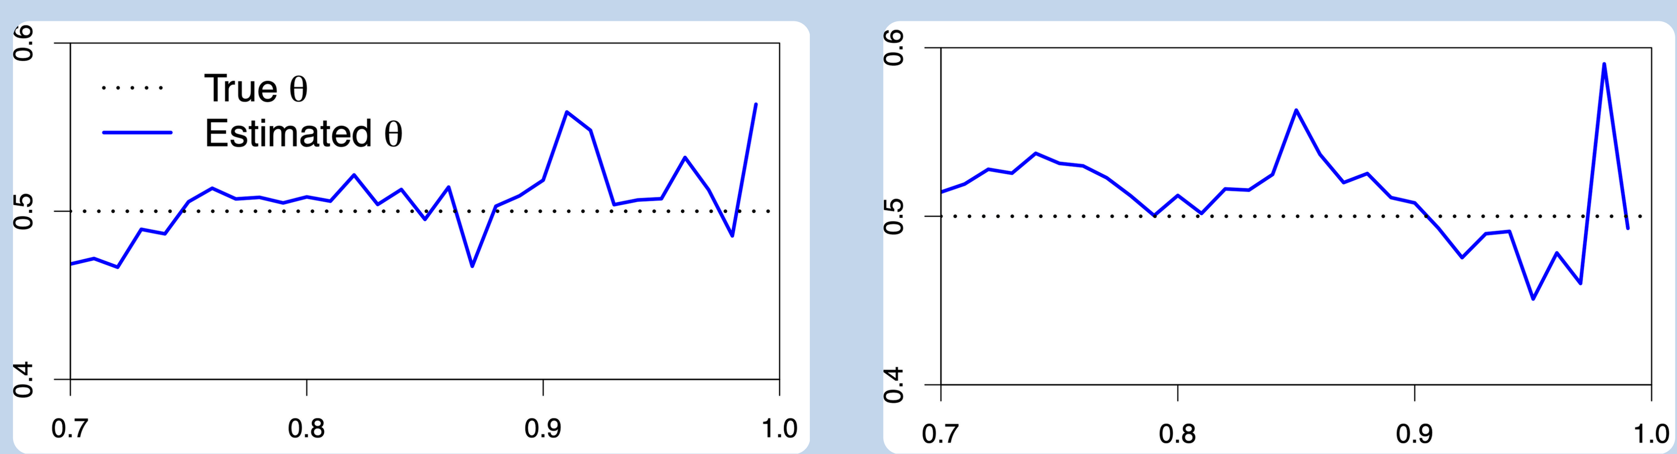

Two simulated series of 1,000 values from the same dependent process to show instability.
The X-axis shows the chosen threshold level, and the Y-axis gives the estimated extremal index θ.

Methodology

We want to assess how reliable our estimate of the extremal index θ is. But because our raw data is dependent, standard reliability methods that assume independence cannot be applied.

Instead of single events, we focus on whole clusters and the gaps between them, which can, under mild conditions, be treated as independent.

Since we are now treating independent data, we can apply and compare methods to estimate uncertainty.

The bootstrap method (established)

- Repeatedly build new data sets from our original, by picking random values from it (computationally heavy)

- Look at how our measured θ varies across these new datasets to estimate uncertainty.

Influence functions (newly proposed)

This analytical method examines how θ changes when the data are slightly “contaminated”, giving tiny extra weight to specific values. These infinitesimal changes yield a direct formula for θ’s variability, providing a fast, theory-driven measure of uncertainty.

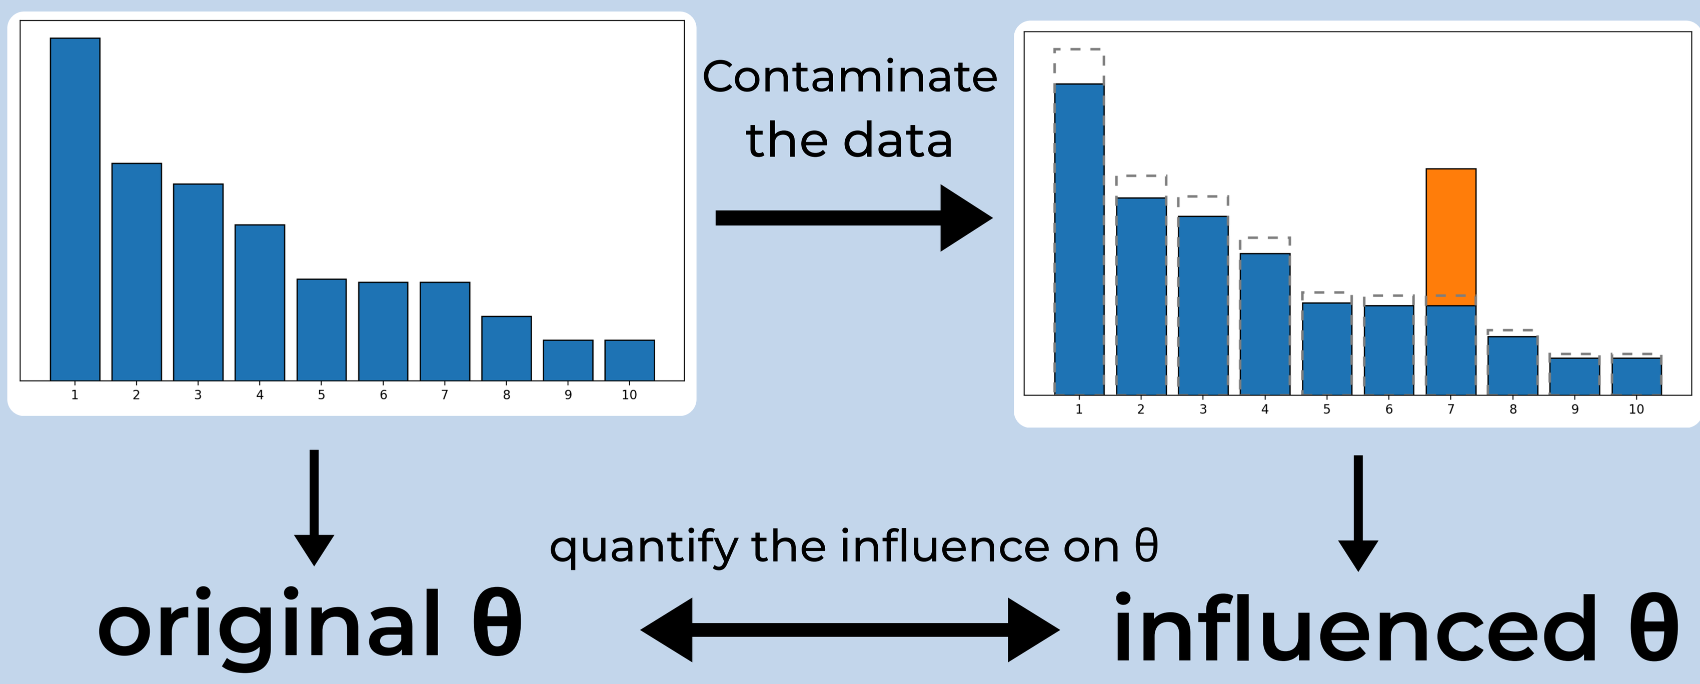

Results & Conclusion

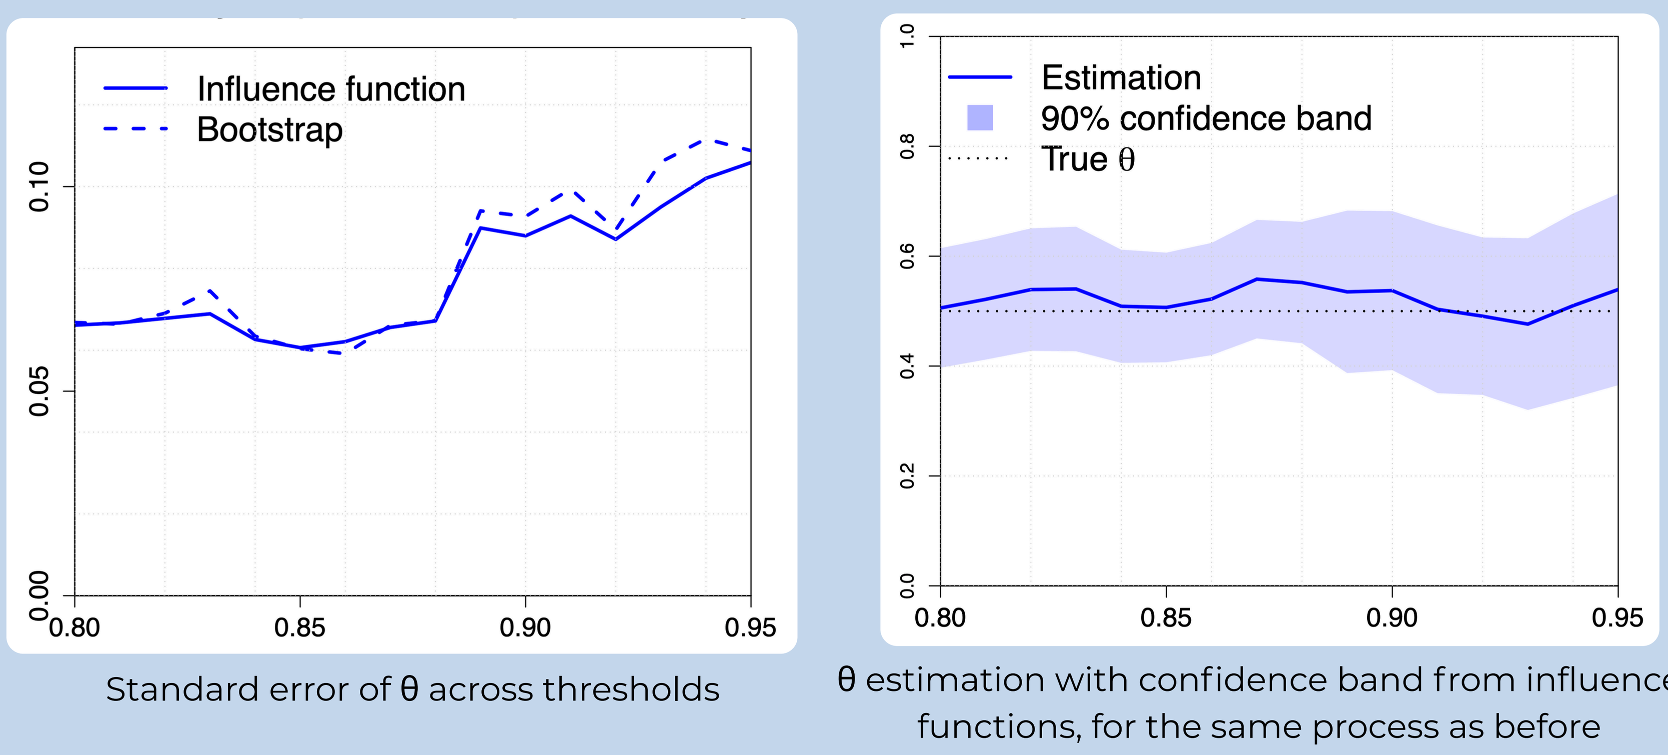

Standard error of θ across thresholds

θ estimation with confidence band from influence functions, for the same process as before

The bootstrap and influence function variances agree closely across thresholds, confirming the accuracy of the analytical formula. This allows fast, direct confidence intervals without heavy resampling and makes it easier to study the variance itself and assess its precision.