Did the incidence of high precipitation levels increase? Statistical evidence for the Netherlands

Alex J. Koning* Philip Hans Franses
Econometric Institute
Erasmus University Rotterdam

Econometric Institute Report 2003-13

Abstract

One of the possible consequences of global warming is that there will be more precipitation days throughout the year, and also that the level of precipitation will be higher. In this paper we provide a detailed statistical analysis of a century of daily precipitation levels for the Netherlands. We show that the often-considered gamma distribution does not fit well to the data. We argue that its incorrect use can lead to spuriously high probabilities of extreme precipitation levels. Relying on advanced nonparametric techniques, we first find that there is indeed more precipitation in the Netherlands, but that this involves only low levels, and second, that the probability of extremely high levels has not changed over time.

Keywords: Daily precipitation, gamma distribution, nonparametric analysis.

*Corresponding author: Econometric Institute, Erasmus University Rotterdam, P.O. Box 1738, NL-3000 DR Rotterdam, The Netherlands, `koning@few.eur.nl`.

1 Introduction

One of the possible consequences of global warming is that it changes precipitation patterns. More precise, when temperatures increase, it is postulated that there will be more precipitation days throughout the year, where also the level of precipitation will be higher. In other words, due to global warming, it rains not only more but also more heavily. It is the aim of this paper to statistically test this conjecture. The statistical analysis will focus on the observed levels of precipitation and on the probability of high levels, where it might occur that such high levels have not been observed. Indeed, we are interested in the likelihood of high precipitation levels, and whether this likelihood has changed over time.

In this paper we provide a detailed analysis of a century of daily precipitation levels observed at the ”de Bilt” weather station in the Netherlands. We performed similar analyses for the five other stations, but we did not find qualitatively different results, and hence our focus is on just a single series. In future work, we will repeat our analysis for many other stations and many other countries, but for the sake of exposition we choose to provide all possible details for a single series.

We have daily data on precipitation for 1906 until and including 2002, collected by the Royal Dutch Meteorological Institute. There were no measurements in April 1945. The measurement unit is 0.1 mm. The data are depicted in Figure 1, and in Figure 2 we display the log-transformed data.

A common way to summarize data as in Figure 1 is to consider the gamma distribution, see Groisman et al. (1999) and the may references cited therein. This distribution has only two parameters to be estimated, and it has been assumed to fit precipitation data rather well. Indeed, in many studies this fit has been assumed *a priori* and it has not been tested. As the estimated parameters of the gamma distribution immediately convey important information on the tail behaviour, and hence on the likelihood of high levels, we believe that a test of the empirical validity of the gamma distribution is of tantamount importance.

When we check the empirical validity of this distribution for the daily data, that is, for all considered years with daily data, we strongly reject the adequacy of the gamma distribution, see Section 2. Zooming in on year-specific data, we find in fact that in about one quarter of the years this distribution does not fit.

Before we turn to alternative methods to summarize the data, and to try to answer the question in the title, we should spend a few lines on the estimation procedure for the gamma distribution. In practice, its two parameters are estimated using the two sufficient statistics. As we will outline in Section 2, these sufficient statistics put heavy weight on small-valued observations. Hence, if, over time, something has changed for these observations, one would find other estimated parameters. As these parameters have a one-to-one link with the probabilities of high-valued observations, it is useful to see how sensitive the estimated parameters are to possible changes in small-valued observations. In terms of our research question, if it had occurred that there are now more days with low precipitation, a straightforward use of the gamma distribution would automatically lead to the suggestion that the probability of high levels had increased too.

One solution could now be to estimate parameters for a left-truncated gamma distribution. Of course, the truncation point is arbitrary, but as such, this kind of analysis can provide some first insights. Another solution, which we will outline in detail in Section 3, is to rely on nonparametric techniques. In that section, we first use nonparametric monitoring techniques to examine whether properties of the empirical distributions of the annual data have changed over time. Our main finding is that there are noticeable changes around 1970. Next, we use nonparametric methods to examine if these changes have affected the probabilities of high levels of precipitation. Our main finding here is that there are no significant changes in these probabilities, whichever formal or graphical test we use.

To summarize, our detailed empirical analysis of a century long series of daily data on precipitation in the Netherlands reveals that there is indeed more precipitation, but that this involves only low levels. And, more importantly, the probability of having extremely high levels has not changed over time.

2 Parametric analysis

In this section we start off analyzing the daily precipitation levels using the same approach as is typically followed in literature, that is, we consider fitting the gamma distribution to the annual data. Next, and this seems to be less common in literature, we examine its empirical validity.

2.1 The gamma distribution

Our intention is to use the daily precipitation data to test the null hypothesis

$$H_0 \ : \ F_{1906} = F_{1907} = \ldots = F_{2002}, \tag{1}$$

where F_i is the cumulative distribution function specifying the distribution of precipitation data $X_{i,1}, \ldots, X_{i,m_i}$ in year i. We shall refer to $X_{i,1}, \ldots, X_{i,m_i}$ as the subsample corresponding to year i. Together, the subsamples form the full sample. The size $m_{1906} + m_{1907} + \ldots + m_{2002}$ of the full sample is denoted by m.

It seems common practice to describe precipitation levels by means of a two-parameter gamma distribution, see, among others, Das (1955), Kotz and Neumann (1963), Thom (1951), Thom (1958), Buishand (1978), Guttman et al. (1993), Groisman et al. (1999). Accordingly, one can consider a parametric test of (1) by assuming that

$$F_i(x) = \int_0^x f(s; \alpha_i, \beta_i) ds \quad \text{for each } i = 1906, \ldots, 2002, \tag{2}$$

with

$$f(x; \alpha, \beta) = \frac{x^{\alpha-1} e^{-x/\beta}}{x^{\alpha} \Gamma(\alpha)}, \tag{3}$$

where Γ denotes the gamma function. The parameter α describes the shape of the gamma distribution, and the parameter β its scale.

A parametric approach requires parameter estimates, which may be obtained by maximum likelihood. If the maximum likelihood estimator of the parameter vector is unique, then it depends only on the complete and sufficient statistic, provided the latter exists, see Arnold (1990, p. 338). In case of the two-parameter gamma distribution, the complete and sufficient statistic belonging to the i^{th} subsample $X_{i,1}, \ldots, X_{i,m_i}$ is the random vector with components $(m_i)^{-1} \sum_{j=1}^{m_i} X_{i,j}$ and $(m_i)^{-1} \sum_{j=1}^{m_i} \log X_{i,j}$. Loosely speaking, the maximum likelihood estimators of α_i and β_i are derived by comparing the mean of the original data to the mean of the log-transformed data, see (17.48a) and (17.48b) in Johnson et al. (1997, p. 361).

The maximum likelihood estimators of the parameters of a gamma distribution are thus largely determined by observations in the left-hand tail of this distribution. Hence, inference for the right-hand tail of the precipitation levels involving "gamma" maximum likelihood estimators is based on extrapolating behaviour of the left-hand tail of the sample to behaviour of the right-hand tail of the distribution. Notice by the way the marked lines for low levels in Figure 2, which emphasize that the statistics $(m_i)^{-1}\sum_{j=1}^{m_i}\log X_{i,j}$ are extremely sensitive to rounding of the observations close to zero.

A possible way of overcoming this sensitivity is by fitting a left-truncated gamma distribution to daily precipitation levels above a threshold value, τ say. The left-truncated gamma probability density function with threshold τ is defined as

$$f_\tau(x;\alpha,\beta) = \frac{f(x;\alpha,\beta)}{1 - I(\tau/\beta,\alpha)}, \quad x > \tau,$$

where

$$I(\eta,\alpha) = \frac{1}{\Gamma(\alpha)}\int_0^\eta u^{\alpha-1}\exp\{-u\}\,du$$

is the incomplete Gamma integral. Applying the maximum likelihood procedure to those observations in the i^{th} subsample which exceed the threshold τ, yields an estimator $\hat{\theta}_i^{(\tau)} = (\hat{\alpha}_i^{(\tau)}, \hat{\beta}_i^{(\tau)})$ of $\theta_i = (\alpha_i, \beta_i)$ from solving the equations

$$\frac{\sum_{j=1}^{m_i} 1_{\{X_{i,j}>\tau\}}\ln X_{i,j}}{\sum_{j=1}^{m_i} 1_{\{X_{i,j}>\tau\}}} = \ln\hat{\beta}_i^{(\tau)} + \Psi(\hat{\alpha}_i^{(\tau)}) - \frac{I_\alpha(\tau/\hat{\beta}_i^{(\tau)}, \hat{\alpha}_i^{(\tau)})}{1 - I(\tau/\hat{\beta}_i^{(\tau)}, \hat{\alpha}_i^{(\tau)})},$$

$$\frac{\sum_{j=1}^{m_i} 1_{\{X_{i,j}>\tau\}}X_{i,j}}{\sum_{j=1}^{m_i} 1_{\{X_{i,j}>\tau\}}} = \hat{\alpha}_i^{(\tau)}\hat{\beta}_i^{(\tau)} + \tau\frac{I_\eta(\tau/\hat{\beta}_i^{(\tau)}, \hat{\alpha}_i^{(\tau)})}{1 - I(\tau/\hat{\beta}_i^{(\tau)}, \hat{\alpha}_i^{(\tau)})}.$$

Here $I_\eta(\eta,\alpha)$ and $I_\alpha(\eta,\alpha)$ denote the first order derivatives of $I(\eta,\alpha)$ with respect to η and α. These derivatives may be numerically evaluated with the algorithm in Moore (1982). And, $\Psi(\alpha)$ is the digamma function.

It is shown in Appendix A that if the i^{th} subsample is a random sample from the gamma distribution with shape parameter α_i and scale parameter β_i, then the distribution of

$$T_i^{(\tau)} = n \begin{pmatrix} \hat{\alpha}_i^{(\tau)} - \hat{\alpha}_i \\ \hat{\beta}_i^{(\tau)} - \hat{\beta}_i \end{pmatrix}^T \left(\mathbf{\Sigma}_\tau^{-1} - \mathbf{\Sigma}^{-1} \right)^{-1} \begin{pmatrix} \hat{\alpha}_i^{(\tau)} - \hat{\alpha}_i \\ \hat{\beta}_i^{(\tau)} - \hat{\beta}_i \end{pmatrix}$$

tends to a chi-square distibution with 2 degrees of freedom, as the subsample size m_i tends to infinity. The matrices $\mathbf{\Sigma}$ and $\mathbf{\Sigma}_\tau$ are Fisher information matrices belonging to $\hat{\theta}_i$ and $\hat{\theta}_i^{(\tau)}$. One may view $T_i^{(\tau)}$ as a statistic for testing the goodness-of-fit of the gamma distribution.

2.2 Results

For each year $i = 1906, \ldots, 2002$, we evaluate the statistic $T_i^{(\tau)}$ after computing $\hat{\theta}_i = (\hat{\alpha}_i, \hat{\beta}_i)^T$ and $\hat{\theta}_i^{(\tau)} = (\hat{\alpha}_i^{(\tau)}, \hat{\beta}_i^{(\tau)})^T$, with threshold $\tau = 5$mm. This yields a sequence $T_{1906}^{(\tau)}, \ldots, T_{2002}^{(\tau)}$ of test statistics, which approximately behaves as a random sample from the chi-squared distribution with 2 degrees of freedom under the joint null hypothesis that the precipitation levels within each year $i = 1906, \ldots, 2002$ follow a gamma distribution with shape parameter α_i and scale parameter β_i. In particular, the "sample sum" $\sum_{i=1906}^{2002} T_i^{(\tau)}$ approximately has a chi-squared distribution with $2 * (2002 - 1906 + 1) = 194$ degrees of freedom under this joint null hypothesis. The test statistic $\sum_{i=1906}^{2002} T_i^{(\tau)}$ takes the value 463.399 [P-value 0.000], and hence the null hypothesis should be rejected. The gamma distribution apparently gives an inadequate description of the daily precipitation levels throughout the whole period 1906–2002. Hence, there are one or more years in which the precipitation levels do not follow a gamma distribution.

To gain further insight into the nature of the rejection of the null hypothesis, Figure 3 displays a QQ-plot of the "sample" $T_{1906}^{(\tau)}, \ldots, T_{2002}^{(\tau)}$ versus the chi-square distribution with 2 degrees of freedom. This plot may be interpreted in a similar way as the Daniel's plot in industrial statistics, see Daniel (1959), Box et al. (1978, p. 329). Points clearly deviating from the line with intercept 0 and slope 1 correspond to years in which the precipitation levels do not follow a gamma distribution. The 22 points marked with a dot in Figure 3 are considered to correspond with "non-gamma" years. Hence, in about 1 out of 4 years, the gamma distribution does not seem to fit the daily precipitation levels.

In Figure 4 the difference $\hat{\beta}_i^{(\tau)} - \hat{\beta}_i$ between scale estimators is plotted versus the difference $\hat{\alpha}_i^{(\tau)} - \hat{\alpha}_i$ between shape estimators. It is interesting to observe that $\hat{\beta}_i^{(\tau)} < \hat{\beta}_i$ and $\hat{\alpha}_i^{(\tau)} > \hat{\alpha}_i$ for all 22 non-gamma years.

Using the approximation of Gray et al. (1969), see also (17.26) in Johnson et al. (1997, p. 347), one can show that

$$\lim_{x \to \infty} \frac{1}{x} \log \left(1 - F(x; \alpha, \beta)\right) = -\frac{1}{\beta} \tag{4}$$

holds for any gamma cumulative distribution function $F(x; \alpha, \beta)$. In words, for large values of x, the gamma cumulative distribution function is primarily determined by the scale parameter β. Hence, the observation that $\hat{\beta}_i^{(\tau)} < \hat{\beta}_i$ for non-gamma years in Figure 4 implies that in those years the probability of exceeding large precipitation levels is overestimated when using the estimator $\hat{\theta}_i = (\hat{\alpha}_i, \hat{\beta}_i)^T$. The "threshold" estimator $\hat{\theta}_i^{(\tau)} = (\hat{\alpha}_i^{(\tau)}, \hat{\beta}_i^{(\tau)})^T$ suffers from the same drawback as the "full" estimator, and hence estimating the probabilities of extremely high precipitation values for left-truncated data does not provide a satisfactory solution.

3 Nonparametric analysis

The empirical results in the previous sesction conveyed that for many years the gamma distribution does not fit the data. Additionally, and perhaps due to this, this poor fit was seen to imply large probabiblities of high precipitation levels. To be able to provide a better answer to the question in the title, we now resort to nonparametric methods. Again, we first deal with the method, and next we present our results.

3.1 Methods

Let

$$\hat{F}_i(x) = \frac{1}{m_i} \sum_{j=1}^{m_i} 1_{\{X_{i,j} \le x\}}$$

be the empirical estimator of $F(x)$ in the subsample corresponding to year i, and let

$$\bar{F}_{1906:2002}(x) = \frac{1}{m} \sum_{i=1906}^{2002} m_i \hat{F}_i(x) = \frac{1}{m} \sum_{i=1906}^{2002} \sum_{j=1}^{m_i} 1_{\{X_{i,j} \le x\}}$$

be the empirical estimator of $F(x)$ in the full sample. Lemma 1 in Hjort and Koning (2001) implies that under the null hypothesis (1), the process

$$B_{1906:2002}(u,x) = \frac{1}{\sqrt{m}} \sum_{i=1906}^{1905+[97u]} m_i \left(\hat{F}_i(x) - \bar{F}_{1906:2002}(x)\right), \quad u \in [0,1],\ x \in I\!R, \tag{5}$$

approximately behaves as a zero mean Gaussian process with covariance function $\{(u \wedge u') - uu'\}\,\{F_0\,(x \wedge x') - F_0(x)F_0(x')\}$. Observe that we may view u as "coded" time. As u progresses from 0 to 1, the "uncoded" time progresses from the year 1906 to 2002. In Section 2.6 in Csörgő and Horváth (1997), a multivariate version of $B_{1906:2002}(u,x)$ is used to detect change point alternatives.

We can test the constancy hypothesis (1) by evaluating the supremum Kolmogorov statistic

$$T_{\text{Kol}} = \sup_{x \in \boldsymbol{R}} \sup_{u \in [0,1]} \left| B_{1906:2002}(u,x) \right|,$$

the supremum Kuiper statistic

$$T_{\text{Kui}} = \sup_{x \in \boldsymbol{R}} \left\{ \sup_{u \in [0,1]} B_{1906:2002}(u,x) - \inf_{u \in [0,1]} B_{1906:2002}(u,x) \right\},$$

the supremum Cramér-von Mises statistic

$$T_{\text{CvM}} = \sup_{x \in \boldsymbol{R}} \left\{ \int_{u \in [0,1]} \left(B_{1906:2002}(u,x)\right)^2 d\mu_{1906:2002}(u) \right\}^{1/2},$$

and the supremum Andersen-Darling statistic

$$T_{\text{AD}} = \sup_{x \in \boldsymbol{R}} \left\{ \int_{u \in [0,1]} \frac{\left(B_{1906:2002}(u,x)\right)^2}{\mu_{1906:2002}(u)\,\left(1 - \mu_{1906:2002}(u)\right)} d\mu_{1906:2002}(u) \right\}^{1/2}.$$

The definition of T_{CvM} and T_{AD} involves the random distribution function $\mu_{1906:2002}(u)$, defined by

$$\mu_{1906:2002}(u) = m^{-1} \sum_{i=1906}^{1905+[97u]} m_i, \quad u \in [0,1],$$

and this may be considered as a convenient representation of the subsample sizes m_i occurring in the period from year 1906 to year 2002.

The null hypothesis distributions of T_{Kol}, T_{Kui}, T_{CvM} and T_{AD} are intricate and unknown. Fortunately, the bootstrap "works" for these test statistics [cf. paragraph 3.4 in Hjort and Koning (2001)], and therefore we also resort to bootstrap tests.

Let T be a test statistic, and T^* be a bootstrap replication of T. The bootstrap test based on T employs the distribution of T^* to evaluate the achieved significance level (ASL) of T. The usual way of implementing the bootstrap test is to generate a number of bootstrap replications, and count the replications greater than or equal to the achieved value of T, see (Efron and Tibshirani, 1993, p. 232). However, for test statistics such as T_{Kol}, T_{Kui}, T_{CvM} and T_{AD} it is known that its distribution under the null hypothesis (1) approximately has a normal right hand tail, see Koning and Protasov (2003). Thus, a normal probability plot of the bootstrap replications should become linear for large values of the normal score. One may interpret the location where the normal probability plot exceeds the attained value of the test statistic as a "z-score" corresponding to the ASL. Determining the ASL of a bootstrap test via a normal probability plot has the advantage that the number of bootstrap replications can be kept relatively low, like for example in accordance with the rule of thumb (6.4.2) in in Efron and Tibshirani (1993, p. 52).

3.2 Results

We present the empirical results in three parts. First, we consider all years, then only days with substantial precipitation, and finally the annual maximum levels.

3.2.1 Daily precipitation levels

The time series plot in Figure 1 leads to the monitoring plot in Figure 5. Observe that we have $m_i = 365$ for ordinary years, and $m_i = 366$ for leap years. From the monitoring plot we can derive statistics T_{Kol}, T_{Kui}, T_{CvM} and T_{AD}. ASL's belonging to these statistics are obtained from 200 bootstrap simulations, see also Figures 6–9.

Statistic	*Value*	x_{opt}	*ASL*
T_{Kol}	3.46284	0.3115	0.000
T_{Kui}	3.70821	0.3115	0.000
Cramer-von Mises	1.67673	0.3115	0.000
T_{AD}	3.73182	0.3115	0.000

All tests indicate clearly that the constancy hypothesis (1) should be rejected for the daily data. The monitoring plot in Figure 5 suggests the existence of a change point around 1970. Note that $x_{\text{opt}} = 0.3115$ distinguishes between "dry" (or no precipitation) and "wet" (or positive amount of precipitation) days. In fact, we may interpret one minus the cumulative distribution function evaluated in 0.3115 as the probability of a wet day. Hence, it follows from Figure 5 that the number of wet days suddenly decreases around 1970.

The sudden decrease in the number of wet days around 1970 is confirmed by Figure 10. A closer look reveals that of the 23741 days in the period 1906:1970 there are 14080 wet days. Moreover, of the 11658 days in the period 1971:2002, there are 5938 wet days. Hence, the estimated probability of a wet day is 59.31 percent for the period 1906:1970, and 50.93 percent for the years 1971:2002.

3.2.2 "Wet" daily precipitation levels

The monitoring plot in Figure 5 immediately draws attention to the incidence of wet days, but it does not seem to indicate that there is strong non-constancy for positive amounts of precipitation. This suggests that the decrease in incidence of wet days is partially compensated by changes in the amount of precipitation on wet days. To investigate whether these changes indeed occur, we exclude the dry days from the sample. After reconstructing the monitoring plot in Figure 11, where the subsample sizes m_i now range between 147 and 255, the statistics T_{Kol}, T_{Kui}, T_{CvM} and T_{AD} are derived. ASL's belonging to these test statistics are again obtained from 200 bootstrap simulations, see also Figures 12–15.

Statistic	*Value*	x_{opt}	*ASL*
T_{Kol}	2.87243	11.263	0.000
T_{Kui}	3.23032	5.043	0.000
Cramer-von Mises	1.55713	11.263	0.000
T_{AD}	3.58396	11.263	0.000

All tests indicate clearly that the constancy hypothesis (1) should be rejected also for the "wet" days. The upward pointing triangular shapes in the monitoring plot suggest that the amount of precipitation on a wet day suddenly becomes larger around and after 1970.

The sudden increase of the amount of precipitation on a wet day around 1970 is confirmed by Figure 16 and Figure 17. The cumulative distribution function of daily precipitation levels in 1971–2002 is larger than the cumulative distribution function of daily precipitation levels in 1906–1970, when the latter takes a value less than 0.65, which corresponds to precipitation levels lower than 3.25mm per day.

3.2.3 Annual maxima of daily precipitation levels

Figure 16 and Figure 17 do not show a clear difference (at the top right corner) between the cumulative distribution function of daily precipitation levels in 1971–2002 and the cumulative distribution function of daily precipitation levels in 1906–1970 for the higher precipitation levels. To pursue this matter further, we compute the maximum of daily precipitation levels for each year between 1906 and 2002. The time series plot of these annual maxima in Figure 18 leads to the monitoring plot in Figure 19. Observe that subsample sizes m_i are all equal to 1.

The relevant test statistics T_{Kol}, T_{Kui}, T_{CvM} and T_{AD} and corresponding ASL obtained from 200 bootstrap simulations are as follows.

Statistic	*Value*	x_{opt}	*ASL*
T_{Kol}	0.567338	270.114	0.425
T_{Kui}	0.849960	372.406	0.135
T_{CvM}	0.249818	372.406	0.605
T_{AD}	0.581932	271.398	0.660

All tests indicate clearly, see also Figures 20 through 23, that the constancy hypothesis (1) cannot be rejected for the maxima data. Hence, there is no evidence that the distribution of the annual maximum of daily precipitation levels is not constant.

4 Conclusion

This paper relied on parametric and nonparametric techniques to examine whether the incidence of high precipitation levels in The Netherlands could have increased over time. Relying on the parametric techniques, one would be tempted to say yes. However, we showed that the gamma distribution did not fit the data well. Hence, relying on this distribution would lead to the spurious suggestion that extremes can occur more frequently.

Instead, the use of more robust nonparametric techniques led to the suggestion that around 1970 there is a sudden change in the incidence of lower precipitation levels, which consists of two opposite effects. First, the number of wet days decreases, and second, the amount of precipitation on a wet day increases. More importantly, using the same techniques we documented that the incidence of higher precipitation levels does not seem to be affected. In particular, the cumulative distribution function of the annual maximum precipitation levels remains constant throughout the period 1906–2002.

In sum, we conclude that the incidence of high precipitation levels did not increase. As said, in our future work we aim to analyze daily data for various other counties.

References

S. F. Arnold. *Mathematical statistics*. Prentice-Hall, Englewood Cliffs, New Jersey, 1990.

Ørnulf Borgan. Correction: "Maximum likelihood estimation in parametric counting process models, with applications to censored failure time data" [Scand. J. Statist. **11** (1984), no. 1, 1–16]. *Scandinavian Journal of Statistics. Theory and Applications*, 11(4):275, 1984a.

Ørnulf Borgan. Maximum likelihood estimation in parametric counting process models, with applications to censored failure time data. *Scandinavian Journal of Statistics. Theory and Applications*, 11(1):1–16, 1984b.

G. E. P. Box, W. G. Hunter, and J. S. Hunter. *Statistics for experimenters. An introduction to design, data analysis and model building*. Wiley, New York, 1978.

T.A. Buishand. Some remarks on the use of daily rainfall models. *Journal of Hydrology*, 36:295–308, 1978.

M. Csörgő and L. Horváth. *Limit theorems in change-point analysis*. Wiley, New York, 1997.

C. Daniel. Use of half-normal plot in interpreting factorial two-level experiments. *Technometrics*, 1:149, 1959.

S. C. Das. The fitting of truncated type III curves to daily rainfall data. *Australian Journal of Physics*, 7:298–304, 1955.

Bradley Efron and Robert J. Tibshirani. *An introduction to the bootstrap*, volume 57 of *Monographs on Statistics and Applied Probability*. Chapman and Hall, New York, 1993.

H. L. Gray, R. W. Thompson, and G. V. McWilliams. A new approximation for the chi-square integral. *Mathematics of Computation*, 23:85–89, 1969.

P.Ya. Groisman, T.R. Karl, D.R. Easterling, R.W. Knight, Hennessy K.J. Jamason, P.F., R. Suppiah, C.M. Page, J. Wibig, K. Fortuniak, V.N. Razuvaev, A. Douglas, E. Førland, and P.-M. Zhai. Changes in the probability of heavy precipitation: important indicators of climate change. *Climatic Change*, 42:243–283, 1999.

N.B. Guttman, J.R.M. Hosking, and J.R. Wallis. Regional precipitation quantile values for the continental united states computed from l-moments. *Journal of Climate*, 6: 2326–2340, 1993.

N. L. Hjort and A. J. Koning. Constancy of distributions: nonparametric monitoring of probability distributions over time. Technical Report EI 2001-50, Econometric Institute, Erasmus University Rotterdam, 2001.

N. L. Johnson, S. Kotz, and N. Balakrishnan. *Continuous univariate distributions*, volume 1. Wiley, New York, second edition edition, 1997.

A. J. Koning and N. L. Hjort. Constancy of distributions: asymptotic efficiency of certain nonparametric tests of constancy. Technical Report EI 2002-33, Econometric Institute, Erasmus University Rotterdam, 2002.

A. J. Koning and V. Protasov. Tail behaviour of Gaussian processes with applications to the Brownian pillow. *Journal of Multivariate Analysis*, page in print, 2003.

S. Kotz and J. Neumann. On distribution of precipitation amounts for the periods of increasing length. *Journal of Geophysical Research*, 68:3635–3641, 1963.

R. J. Moore. Algorithm as 187: derivatives of the incomplete gamma integral. *Applied Statistics*, 31(3):330–335, 1982.

H. C. S. Thom. A frequency distribution for precipitation (abstract). *Bulletin of the American Meteorological Society*, 32(10):397, 1951.

H. C. S. Thom. A note on the gamma distribution. *Monthly Weather Review*, 86(4): 117–122, 1958.

A On truncated distributions

Let $X_1, \ldots, X_n$ be a random sample from a distribution with probability density function $f_0(x;\theta)$, cumulative distribution function $F_0(x;\theta)$, hazard function $\lambda_0(x;\theta) = f_0(x;\theta)/(1 - F_0(x;\theta))$ and cumulative hazard function $\Lambda_0(x;\theta) = \int_{-\infty}^{x} \lambda_0(s;\theta)ds$, where θ is a q-dimensional parameter vector. That is, q denotes the number of unknown parameters.

According to the proof of Theorem 2 in Borgan (1984b) there exists a q-dimensional mean zero normal random vector $\mathbf{U}$ such that $\sqrt{n}\left(\hat{\theta} - \theta\right)$ tends in distribution to $\mathbf{\Sigma}^{-1}\mathbf{U}$ under regularity conditions, see also Borgan (1984a). Here $\mathbf{\Sigma}$ is the covariancematrix of $\mathbf{U}$.

It follows in a similar way that under the same regularity conditions there exists a q-dimensional mean zero normal random vector $\mathbf{U}_\tau$ such that $\sqrt{n}\left(\hat{\theta} - \theta\right)$ and $\sqrt{n}\left(\hat{\theta}^{(\tau)} - \theta\right)$ tend jointly in distribution to $\mathbf{\Sigma}^{-1}\mathbf{U}$ and $\mathbf{\Sigma}_\tau^{-1}\mathbf{U}_\tau$, where $\mathbf{\Sigma}_\tau$ denotes the covariancematrix of $\mathbf{U}_\tau$. In addition, we have $\mathcal{E}\mathbf{U}^T\mathbf{U}_\tau^T = \mathcal{E}\mathbf{U}_\tau\mathbf{U}^T = \mathbf{\Sigma}_\tau$.

Hence, $\sqrt{n}\left(\hat{\theta}^{(\tau)} - \hat{\theta}\right) = \sqrt{n}\left(\hat{\theta}^{(\tau)} - \theta\right) - \sqrt{n}\left(\hat{\theta} - \theta\right)$ tends in distribution to the q-dimensional normal random vector $\mathbf{\Sigma}_\tau^{-1}\mathbf{U}_\tau - \mathbf{\Sigma}^{-1}\mathbf{U}$, which has expectation $\mathbf{0}$ and covariance matrix

$$\mathcal{E}\left(\mathbf{\Sigma}^{-1}\mathbf{U} - \mathbf{\Sigma}_\tau^{-1}\mathbf{U}_\tau\right)\left(\mathbf{\Sigma}^{-1}\mathbf{U} - \mathbf{\Sigma}_\tau^{-1}\mathbf{U}_\tau\right)^T = \mathbf{\Sigma}_\tau^{-1} - \mathbf{\Sigma}^{-1}.$$

This implies that the limit distribution of

$$T^{(\tau)} = n\left(\hat{\theta}^{(\tau)} - \hat{\theta}\right)^T\left(\mathbf{\Sigma}_\tau^{-1} - \mathbf{\Sigma}^{-1}\right)^{-1}\left(\hat{\theta}^{(\tau)} - \hat{\theta}\right)$$

is chi-square with q degrees of freedom. One may estimate $\mathbf{\Sigma}_\tau$ and $\mathbf{\Sigma}$ consistently by means of their respective estimators

$$\hat{\mathbf{\Sigma}}_\tau = \frac{1}{n}\sum_{\{i:X_i>t\}} \phi(X_i;\hat{\theta})\phi(X_i;\hat{\theta})^T, \quad \hat{\mathbf{\Sigma}} = \frac{1}{n}\sum_{i=1}^{n} \phi(X_i;\hat{\theta})\phi(X_i;\hat{\theta})^T,$$

see Theorem 2 in Borgan (1984b).

B Figures

List of Figures

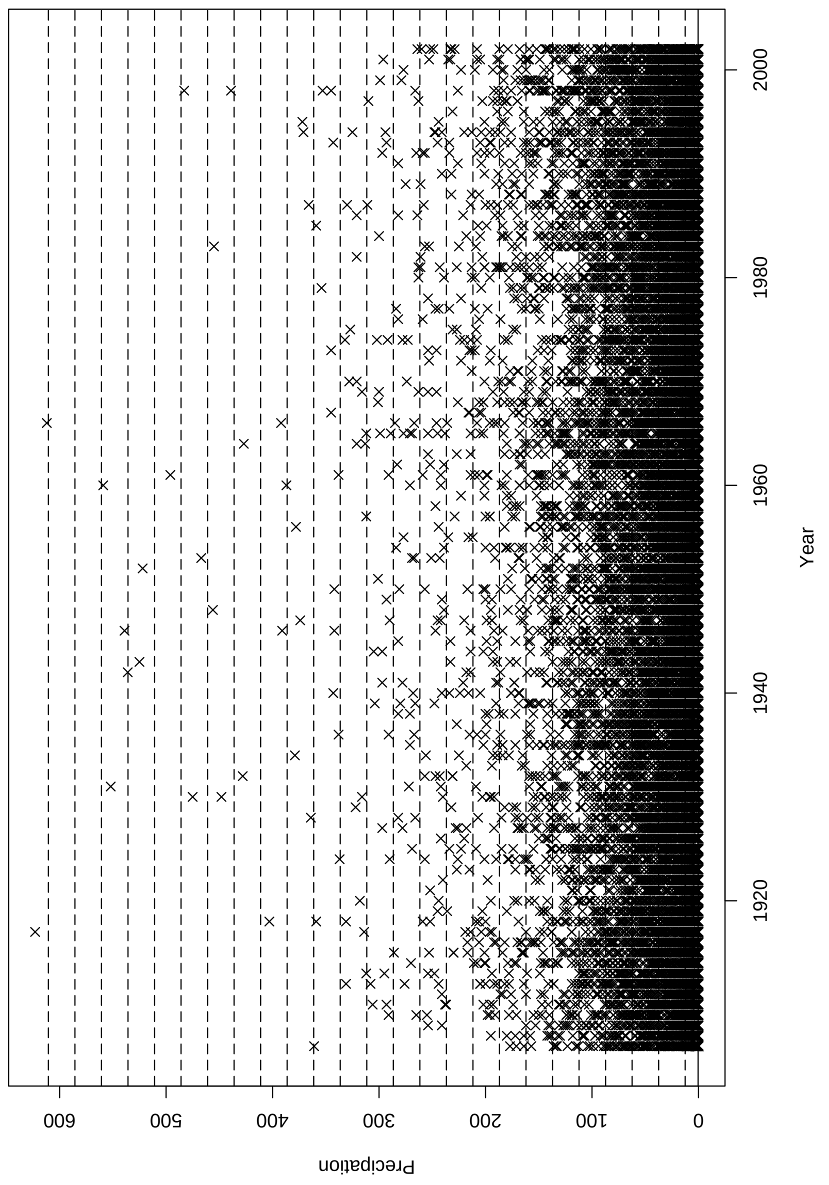

Figure 1: Time series plot of daily precipitation levels, measured at “de Bilt”, The Netherlands. A total number of 35399 daily precipitation levels were recorded during the measurement period starting at January 1, 1906 and ending at December 31, 2002. The data are grouped in 97 subsamples, each covering a calendar year.

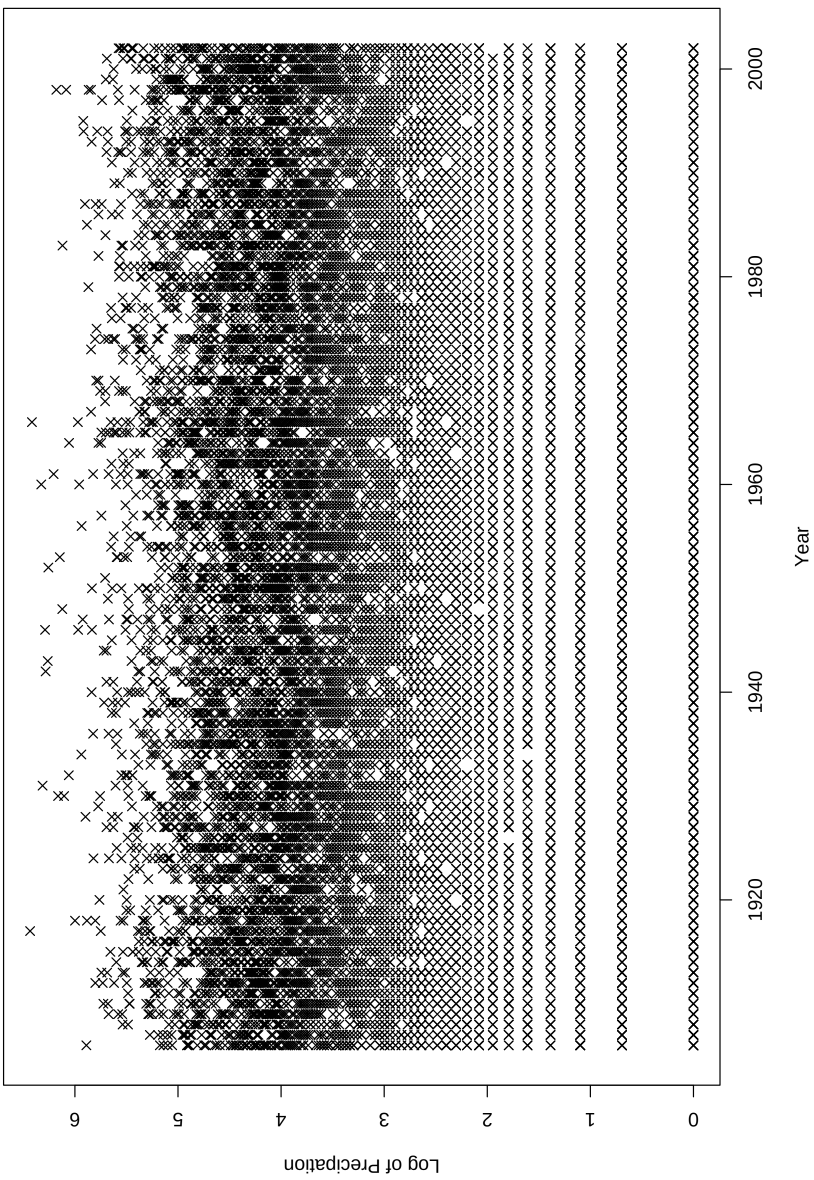

Figure 2: Time series plot of log-transformed daily precipitation levels.

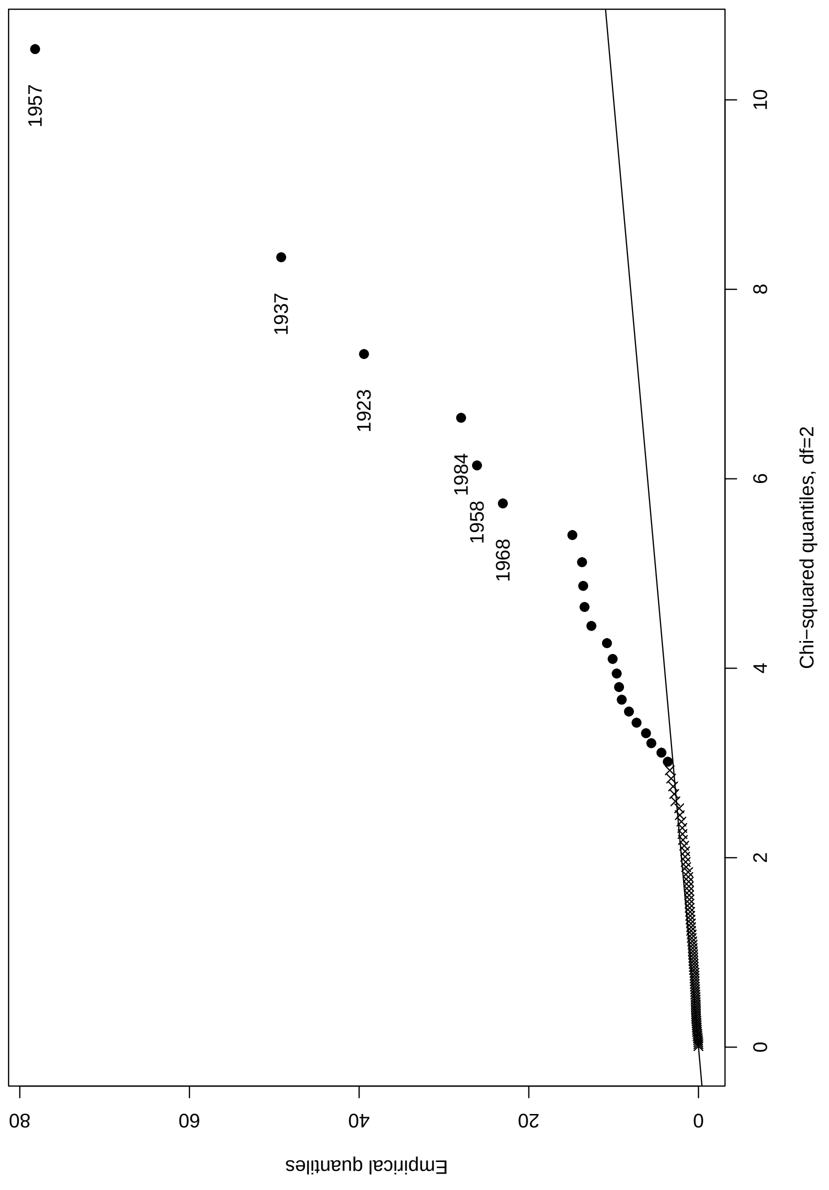

Figure 3: QQ-plot of the test statistics $T_i^{(\tau)}$ versus a chi-squared distribution with 2 degrees of freedom. This plot suggests that precipitation levels in the 22 (dotted) years 1957, 1937, 1923, 1984, 1958, 1968, 1970, 1996, 1945, 1927, 1914, 1994, 1911, 1999, 1988, 1935, 1992, 1947, 1916, 1949, 2002 and 1993 do not follow a gamma distribution.

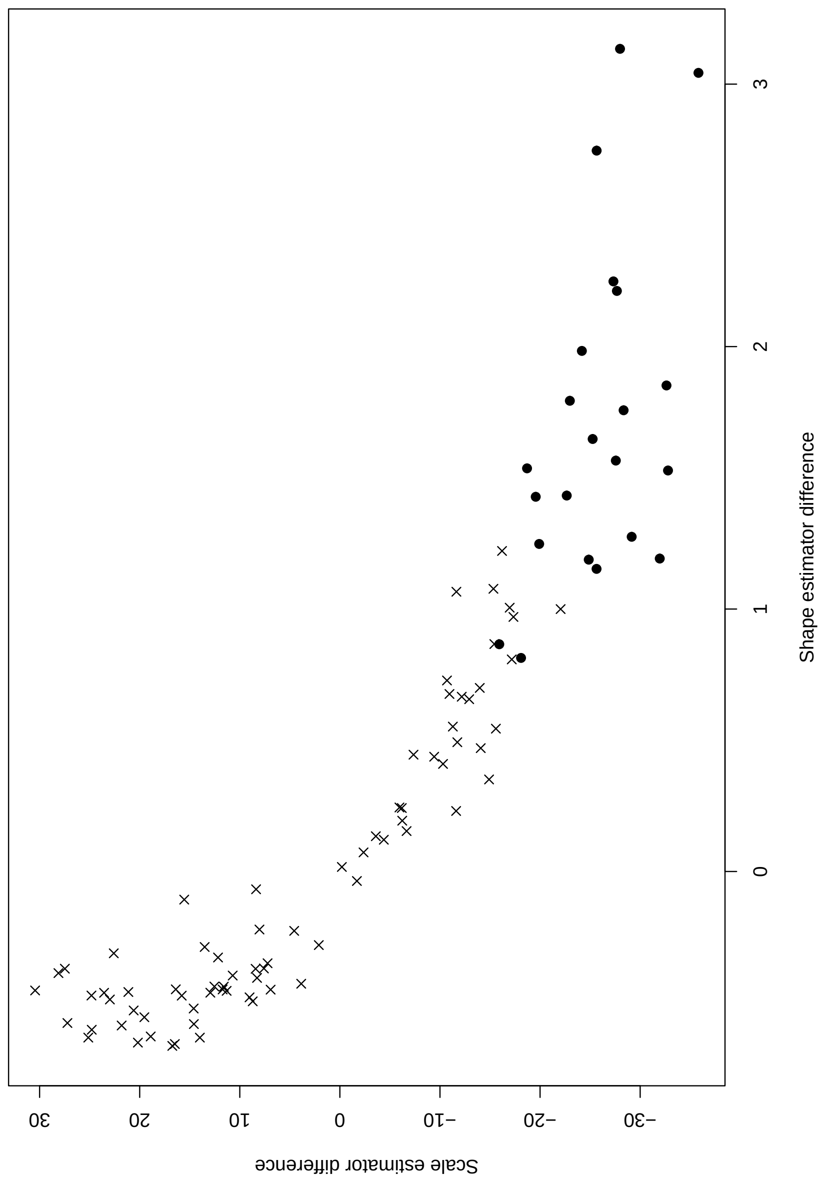

Figure 4: The difference between scale estimators $\hat{\beta}^{(\tau)}$ and $\hat{\beta}$ versus the difference between shape estimators $\hat{\alpha}$ and $\hat{\alpha}^{(\tau)}$, where the dots mark the non-gamma years.

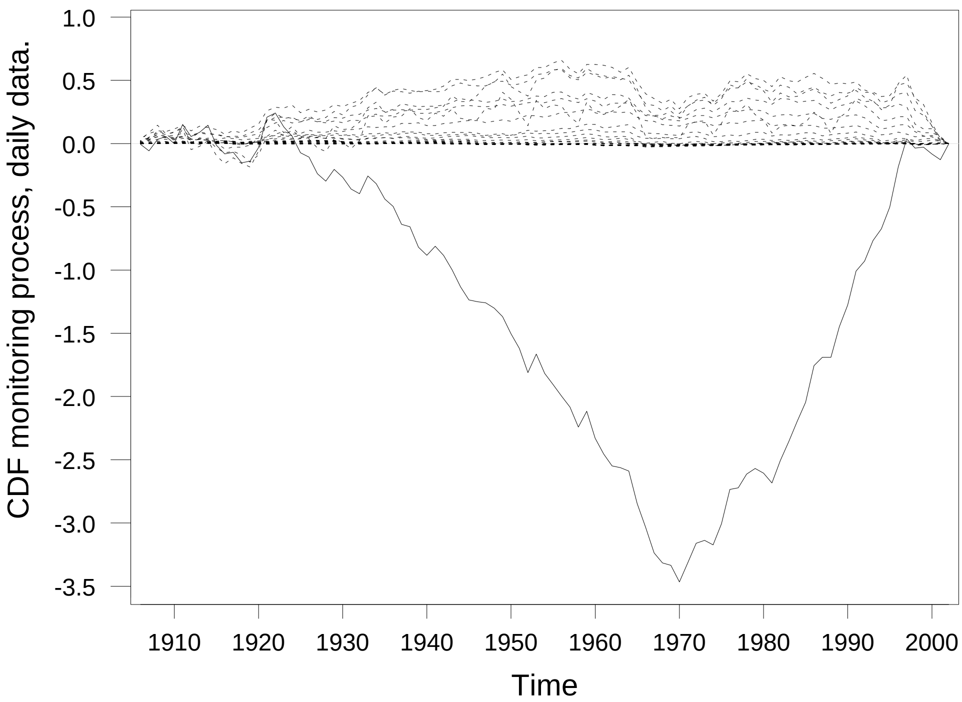

Figure 5: The monitoring process $B_{1906:2002}(u, x)$ for fixed x versus "un-coded" time [instead of versus "coded" time u], for daily precipitation. The dotted lines and the solid line are the results of "scanning" the monitoring process $B_{1906:2002}(u, x)$ along the dotted lines in Figure 1 and the line $x = x_{\text{opt}} = 0.3115$, respectively. The triangular shapes suggest the existence of a rather abrupt change in the cumulative distribution function of the daily precipitation levels around 1970.

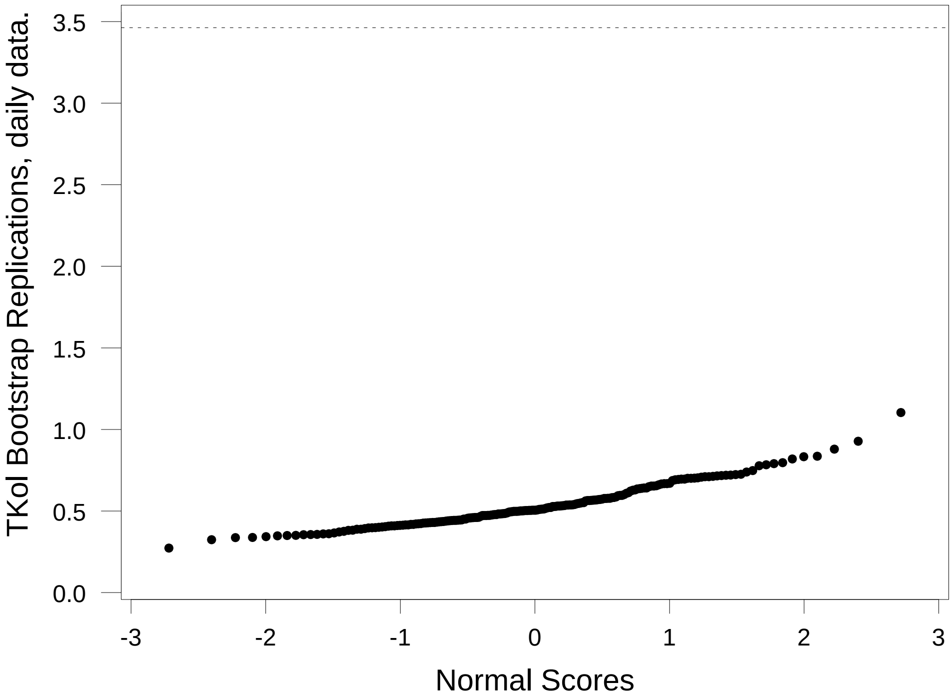

Figure 6: Normal probability plot of 200 bootstrap replications of T_{Kol}, daily precipitation. The dotted line indicates the value 3.463 taken by the test statistic T_{Kol}. According to the theory in Koning and Hjort (2002), the normal probability plot should become linear for larger values of the normal score. As one may interpret the location where the normal probability plot exceeds 3.463 as an estimate of the "z-score" corresponding to the ASL, the plot shows that 3.463 is indeed a highly significant value of T_{Kol}.

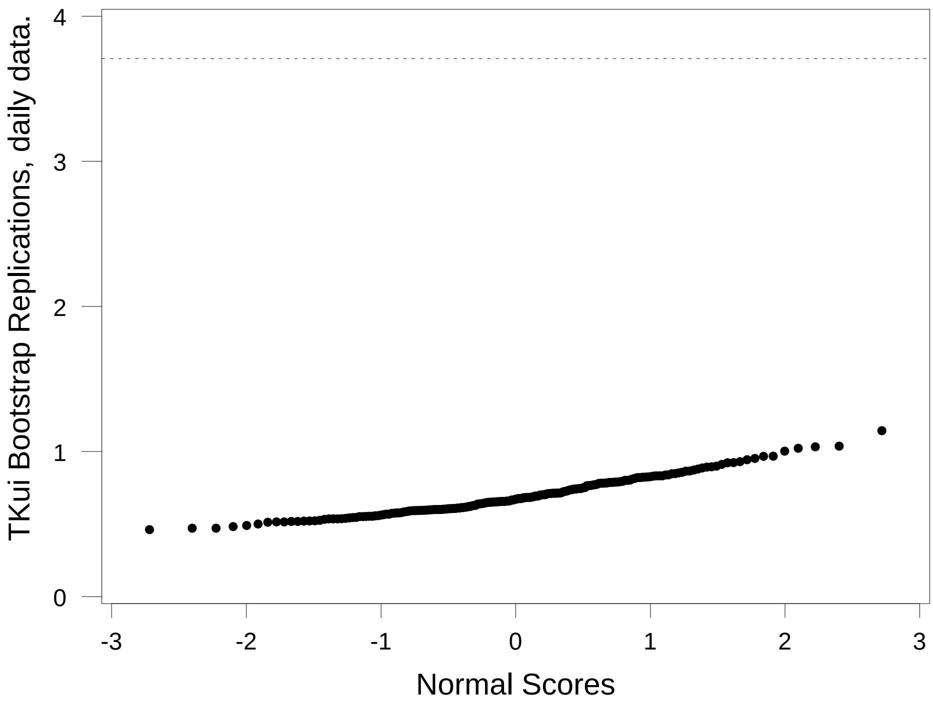

Figure 7: Normal probability plot of 200 bootstrap replications of T_{Kui}, daily precipitation. The dotted line indicates the value 3.708 taken by the test statistic T_{Kui}.

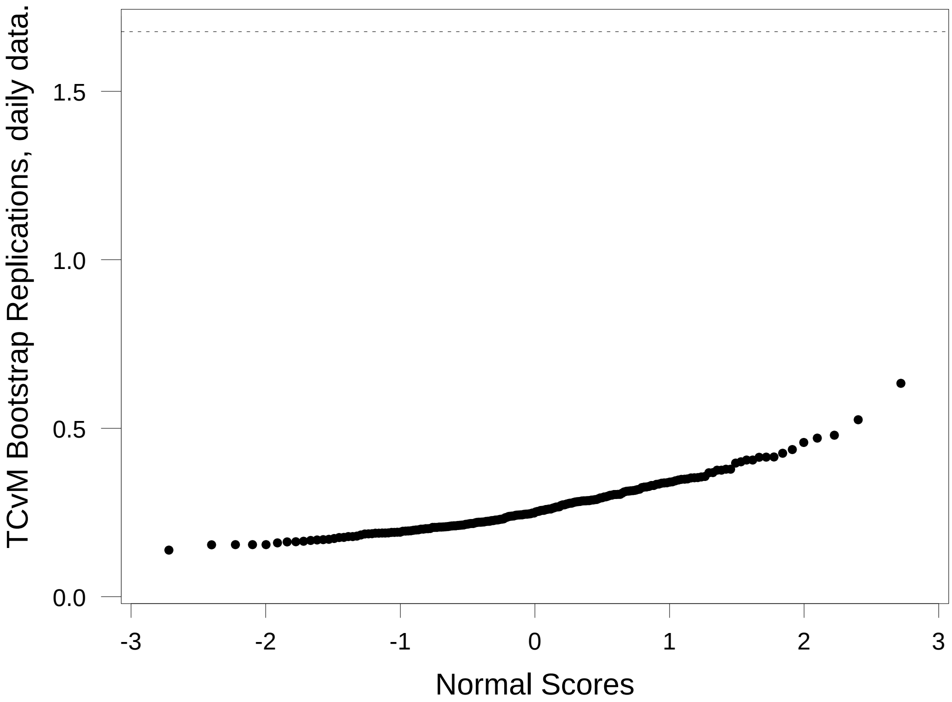

Figure 8: Normal probability plot of 200 bootstrap replications of T_{CvM}, daily precipitation. The dotted line indicates the value 1.677 taken by the test statistic T_{CvM}.

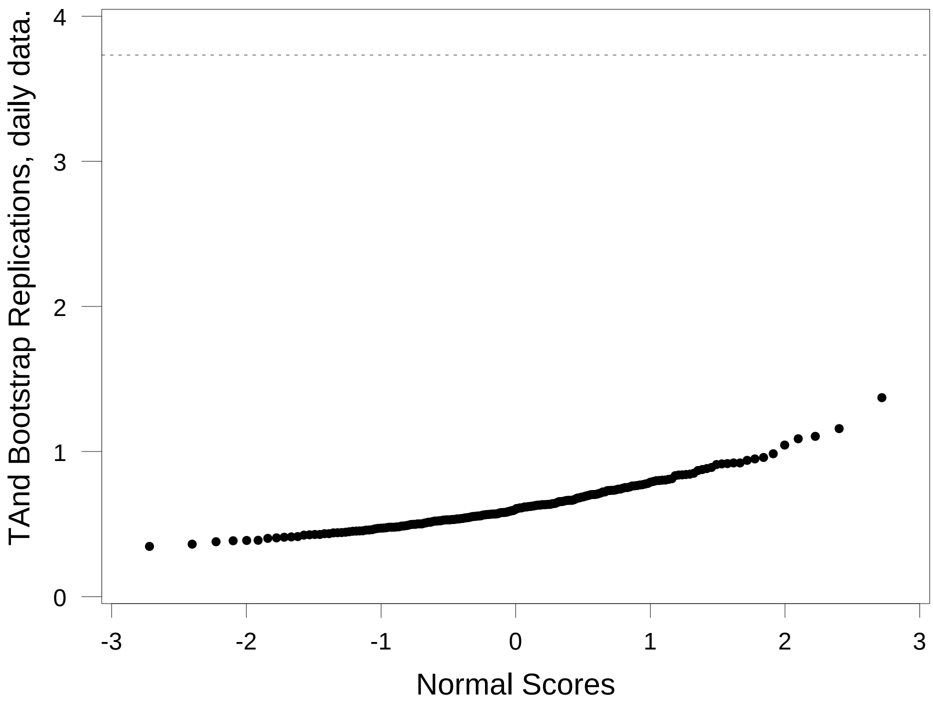

Figure 9: Normal probability plot of 200 bootstrap replications of T_{AD}, daily precipitation. The dotted line indicates the value 3.732 taken by the test statistic T_{AD}.

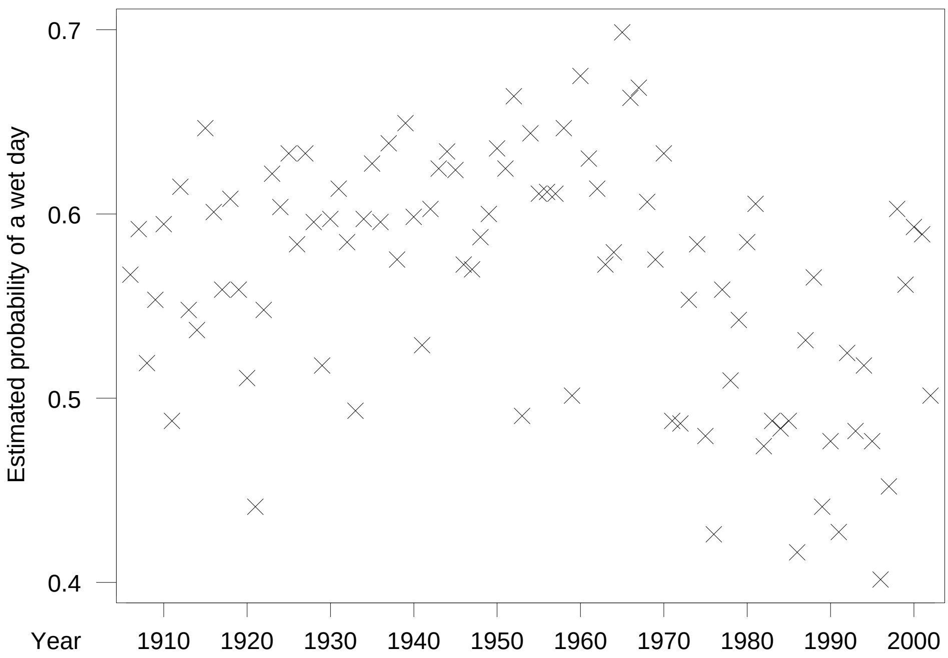

Figure 10: Time series plot of the estimated probability of a wet day , The Netherlands.

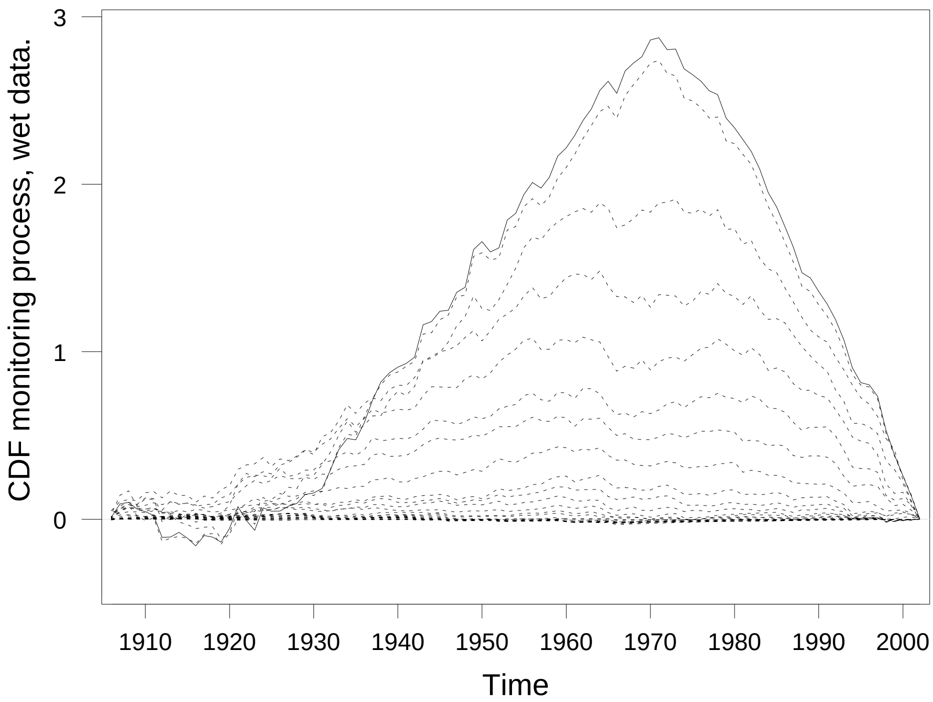

Figure 11: The monitoring process $B_{1906:2002}(u, x)$ for fixed x versus "uncoded" time, for wet daily precipitation data. The dotted lines and the solid line [corresponding to $x_{\text{opt}} = 11.263$] are the results of "scanning" the monitoring process $B_{1906:2002}(u, x)$. The triangular shapes suggest the existence of an abrupt change in the cumulative distribution function of the daily precipitation levels around 1970.

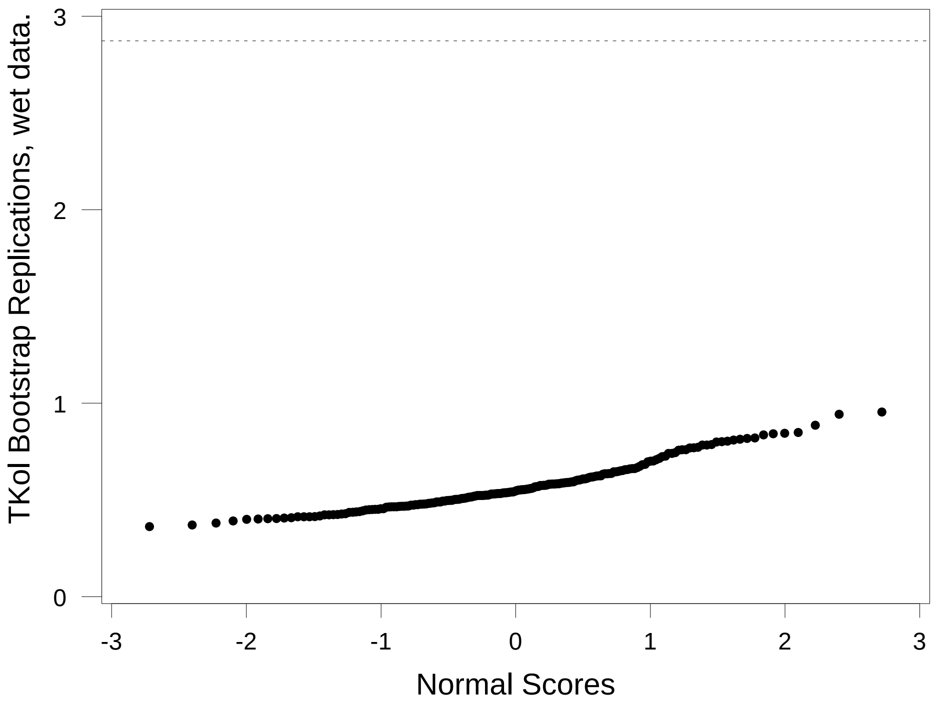

Figure 12: Normal probability plot of 200 bootstrap replications of T_{Kol}, wet daily precipitation. The dotted line indicates the value 2.872 taken by the test statistic T_{Kol}.

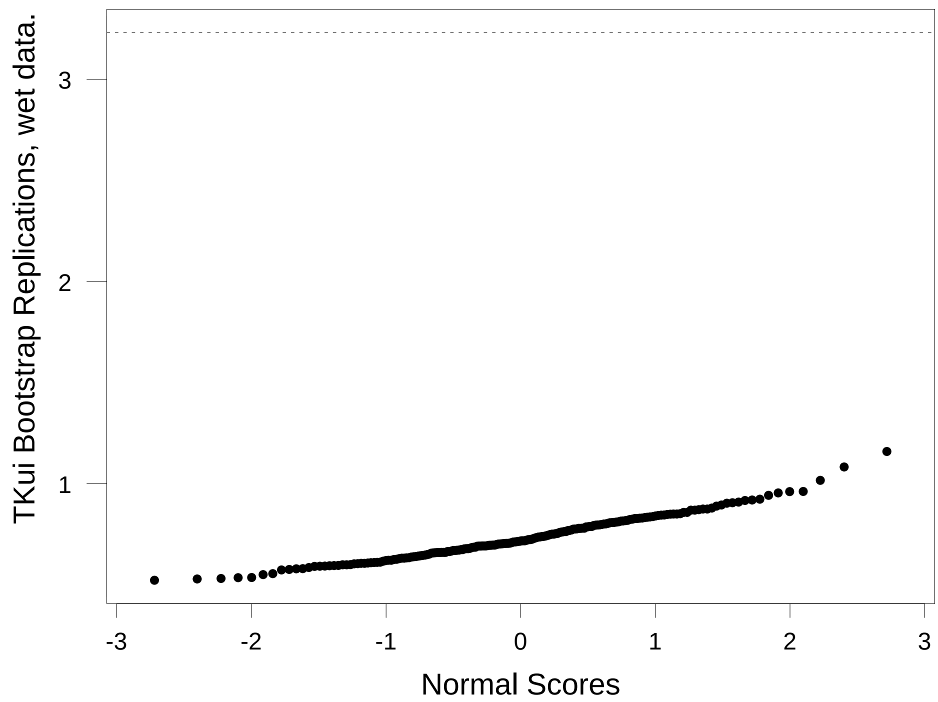

Figure 13: Normal probability plot of 200 bootstrap replications of T_{Kui}, wet daily precipitation. The dotted line indicates the value 3.230 taken by the test statistic T_{Kui}.

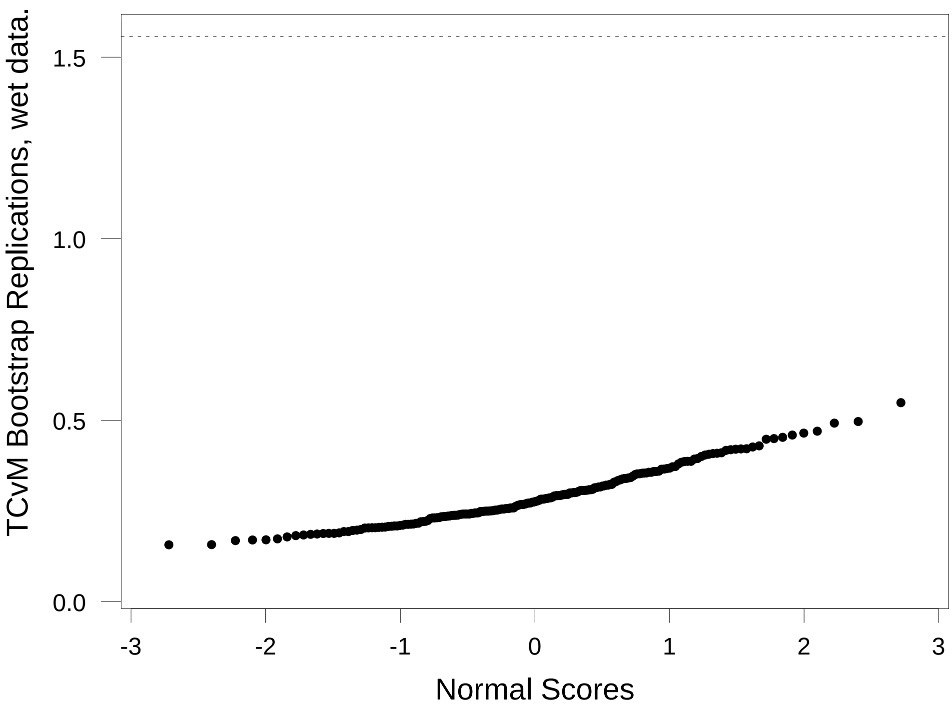

Figure 14: Normal probability plot of 200 bootstrap replications of T_{CvM}, wet daily precipitation. The dotted line indicates the value 1.557 taken by the test statistic T_{CvM}.

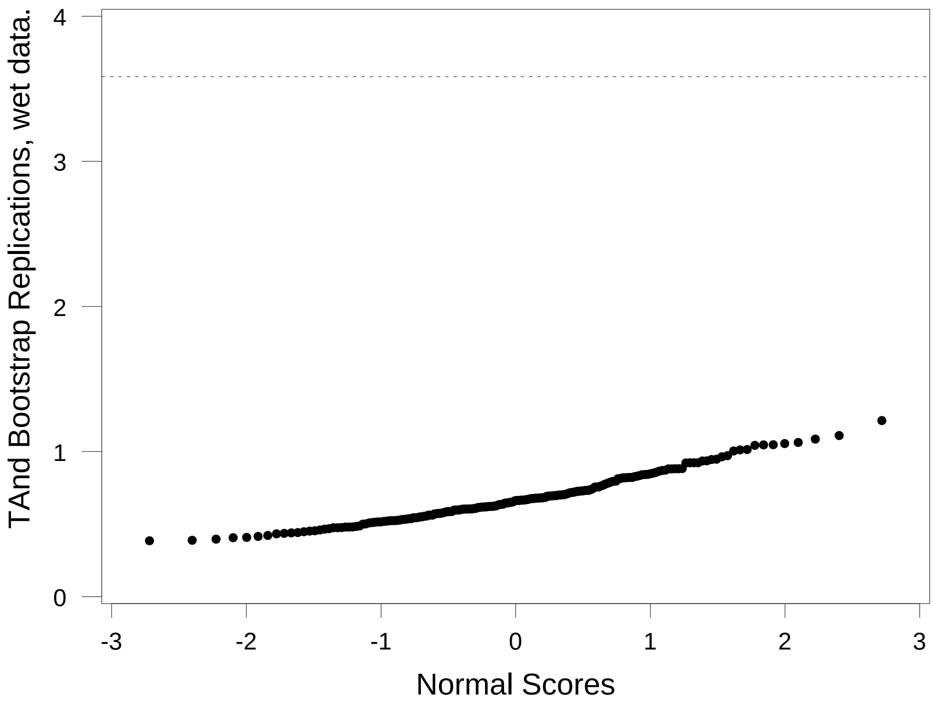

Figure 15: Normal probability plot of 200 bootstrap replications of T_{AD}, wet daily precipitation. The dotted line indicates the value 3.584 taken by the test statistic T_{AD}.

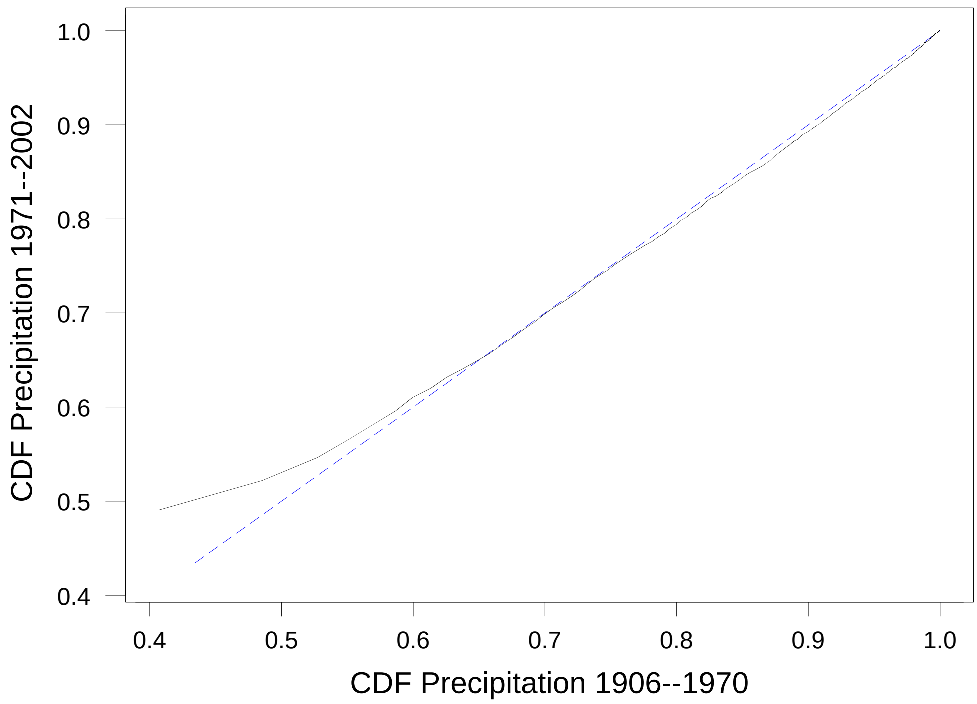

Figure 16: PP-plot of “wet” precipitation levels 1906–1970 versus “wet” precipitation levels 1971–2002. The dashed line represents equality, and the difference between the solid and the dashed line is depicted in Figure 17. Observe that for the lower precipitation levels, the cumulative distribution function of precipitation levels over the period 1971–2002 exceeds the cumulative distribution function of precipitation levels over the period 1906–1970, indicating that lower precipitation levels are relatively more frequent in the period 1971–2002.

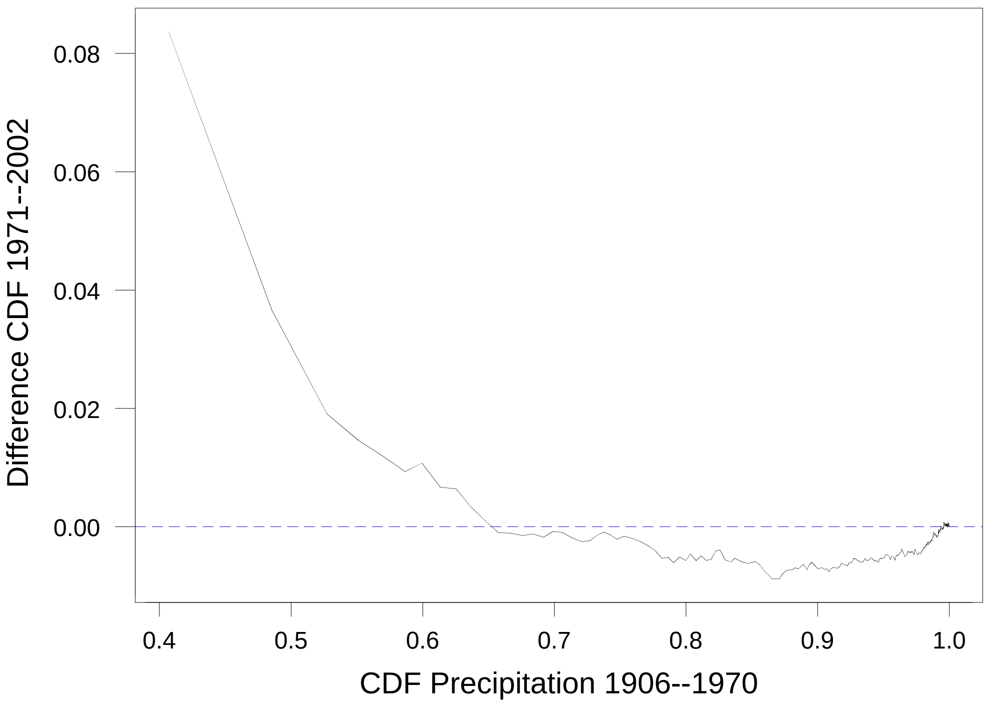

Figure 17: “Detrended” PP-plot of “wet” precipitation levels 1906–1970 versus “wet” precipitation levels 1971–2002. The solid line represents the difference between the solid and the dashed lines in Figure 16. Observe that for the lower precipitation levels, the cumulative distribution function of precipitation levels over the period 1971–2002 exceeds the cumulative distribution function of precipitation levels over the period 1906–1970, indicating that lower precipitation levels were relatively more frequent in the period 1971–2002.

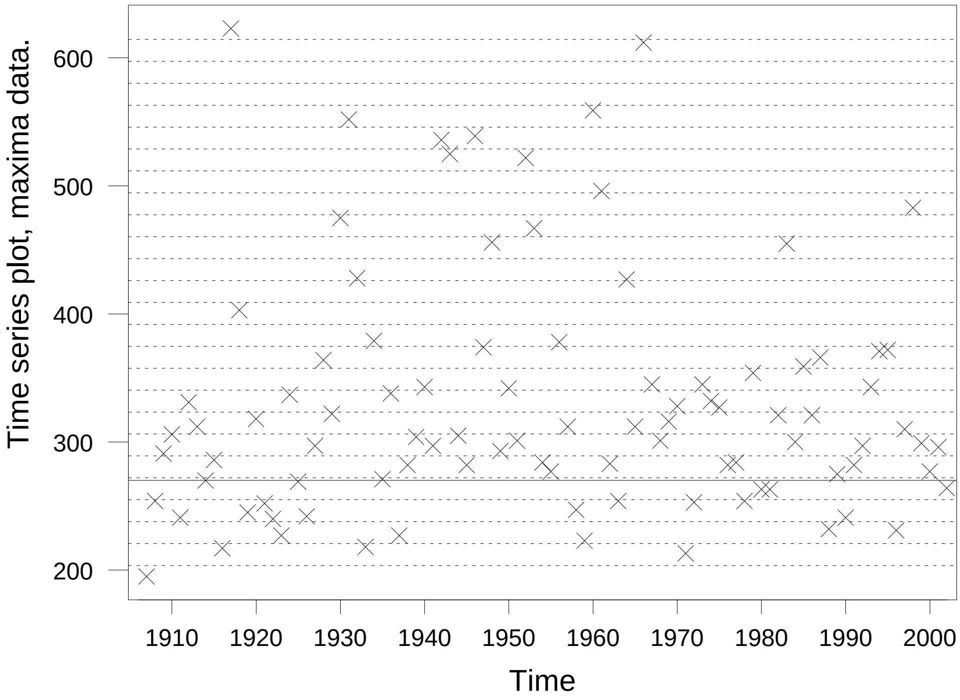

Figure 18: Time series plot of annual maximum precipitation levels, The Netherlands. A total number of 97 annual maxima of daily precipitation levels are recorded during the measurement period starting at January 1, 1906 and ending at December 31, 2002.

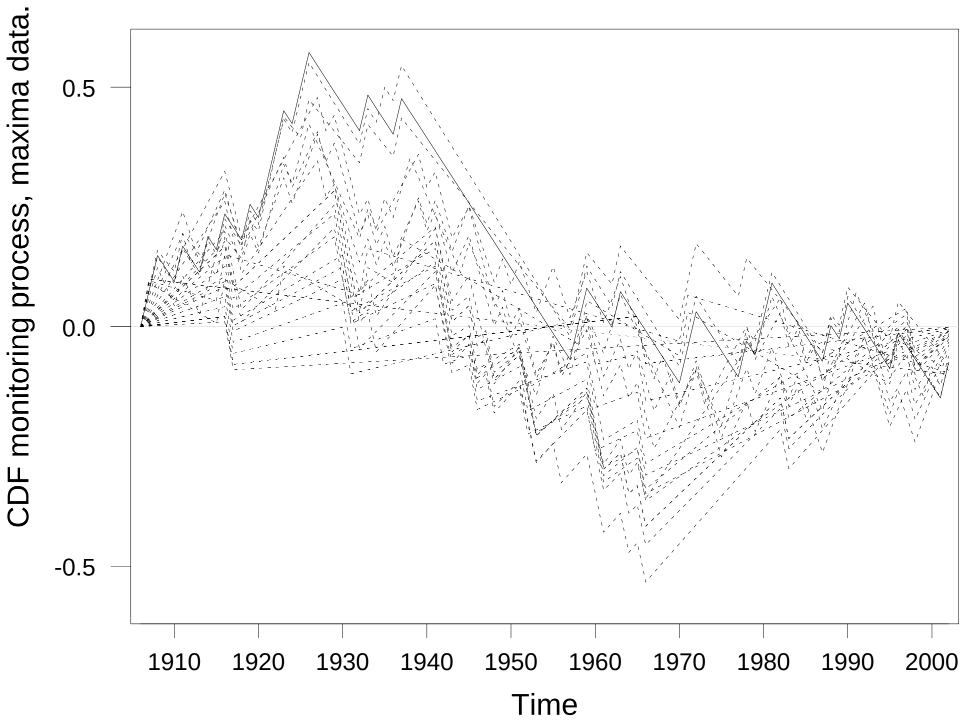

Figure 19: The monitoring process $B_{1906:2002}(u, x)$ for fixed x versus “uncoded” time, annual maxima. The dotted lines and the solid line are the results of “scanning” the monitoring process $B_{1906:2002}(u, x)$ along the dotted lines in Figure 18.

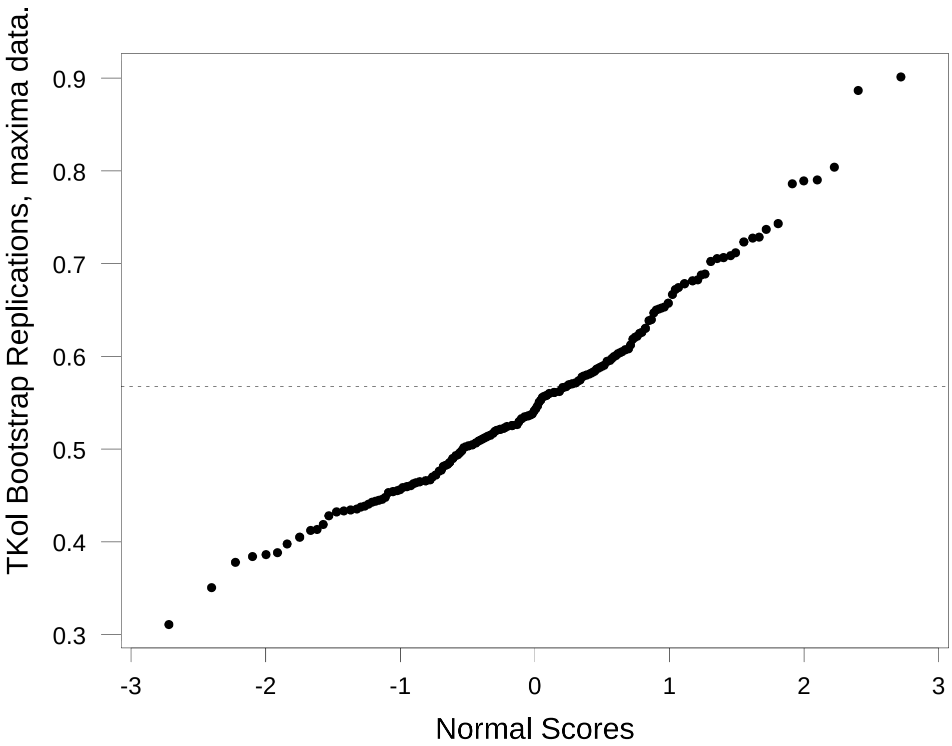

Figure 20: Normal probability plot of 200 bootstrap replications of T_{Kol}, annual maxima. The dotted line indicates the value 0.567 taken by the test statistic T_{Kol}. Counting the number of replications above the dotted line yields an ASL of 0.425.

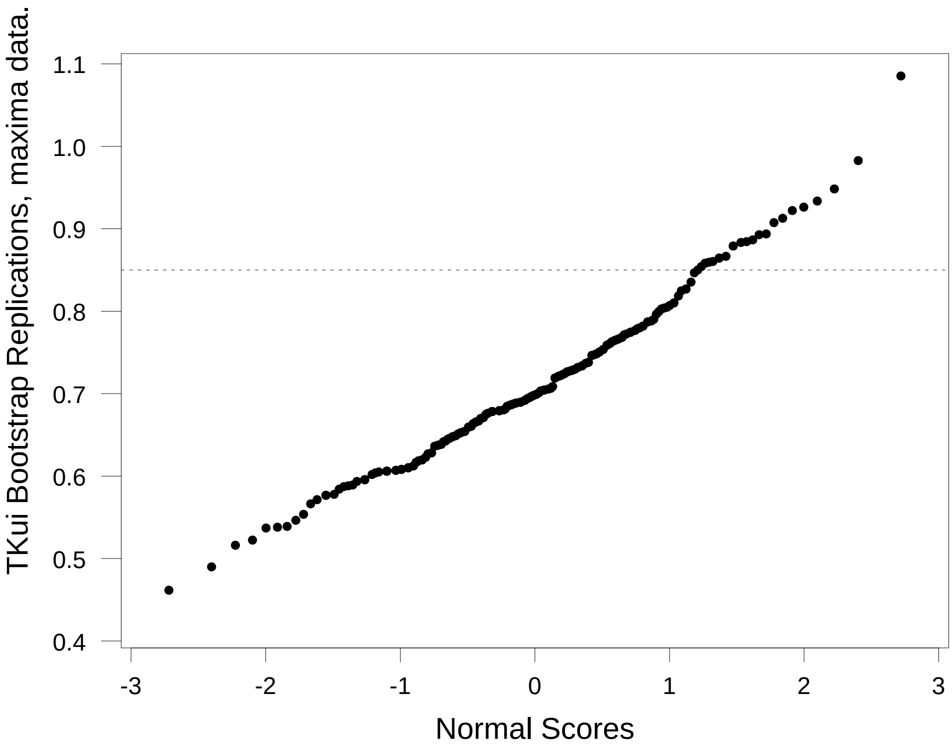

Figure 21: Normal probability plot of 200 bootstrap replications of T_{Kui}, annual maxima. The dotted line indicates the value 0.850 taken by the test statistic T_{Kui}. Counting the number of replications above the dotted line yields an ASL of 0.135.

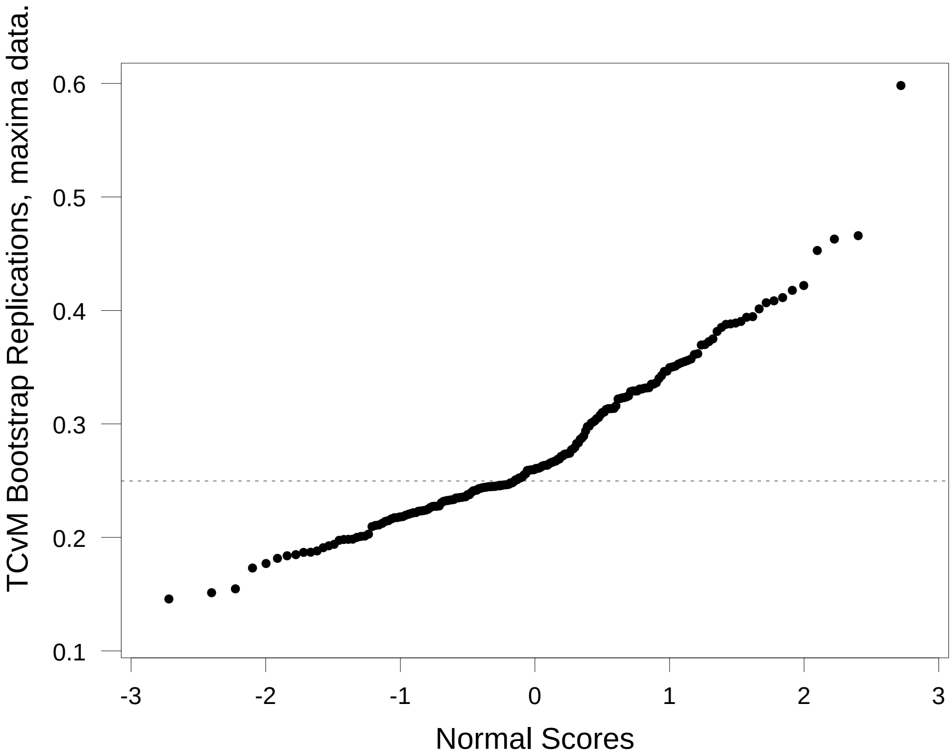

Figure 22: Normal probability plot of 200 bootstrap replications of T_{CvM}, annual maxima. The dotted line indicates the value 0.250 taken by the test statistic T_{CvM}. Counting the number of replications above the dotted line yields an ASL of 0.605.

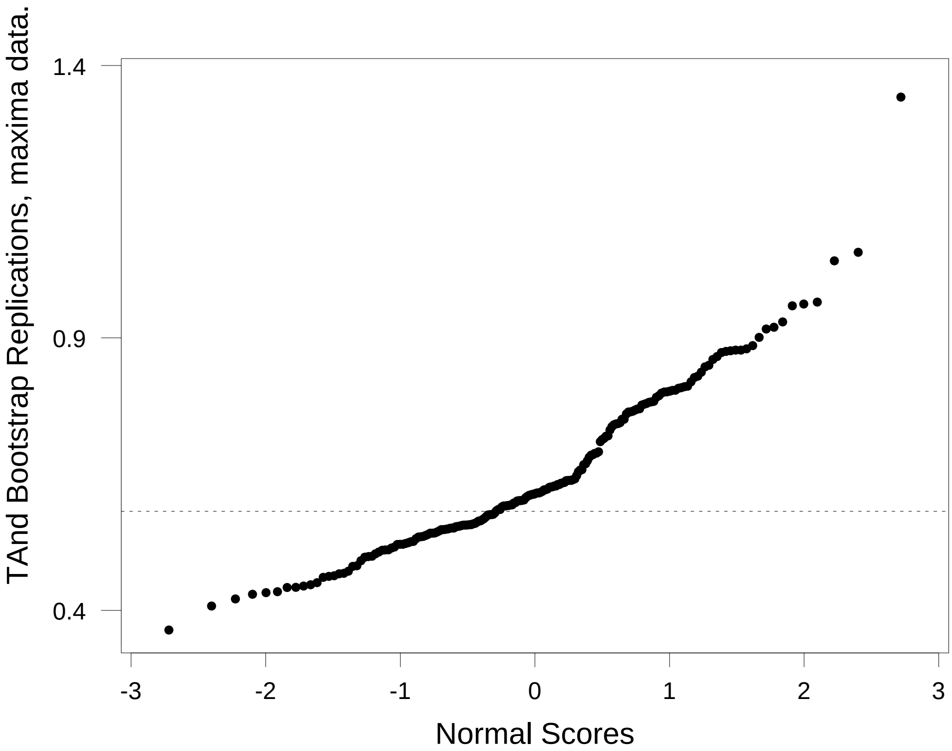

Figure 23: Normal probability plot of 200 bootstrap replications of T_{AD}, annual maxima. The dotted line indicates the value 0.582 taken by the test statistic T_{AD}. Counting the number of replications above the dotted line yields an ASL of 0.660.