Technische Universität Wien

Vienna University of Technology

Testing for Structural Change in GLMs
Implementation in R and Application

Achim Zeileis Kurt Hornik

Contents

- What is a structural change?
- What is R?
- F tests
- Generalized fluctuation tests, with extensions:
  - Quantile residual-based fluctuation tests
  - Score-based fluctuation tests
- Applications

What is a structural change?

Consider the linear regression model

$$y_i = x_i^\top \beta_i + u_i \qquad (i = 1, \dots, n),$$

where at time i:

- y_i — dependent variable,
- x_i — vector of k regressors,
- β_i — vector of k unknown regression coefficients,
- u_i — error term.

What is a structural change?

The null hypothesis is

$$H_0: \quad \beta_i = \beta_0 \qquad (i = 1, \dots, n),$$

which will be tested against the alternatives

$$H_1: \quad \text{not } H_0$$

or

$$H_1^*: \quad \beta_i = \begin{cases} \beta_A & (1 \leq i \leq i_0) \\ \beta_B & (i_0 < i \leq n) \end{cases},$$

respectively.

What is R?

R is a software package for statistical computing—the GNU implementation of the programming language S.

`http://www.R-project.org/`

All the functions and methods for testing for structural change introduced here are implemented in the package `strucchange` availabe from the Comprehensive R Archive Network (CRAN):

`http://cran.R-project.org/`

Authors: Achim Zeileis, Friedrich Leisch, Bruce Hansen, Kurt Hornik, Christian Kleiber, Andrea Peters.

F tests

F tests are designed for the single shift alternatvie H_1^*.

(i) for known (potential) shift point: Chow-Test
Two separate regressions are fitted for the subsamples defined by i_0 and the resulting residuals

$$\hat{e} = (\hat{u}_A, \hat{u}_B)^\top$$

are compared with the residuals from the model without a shift by an F test:

$$F_{i_0} = \frac{(\hat{u}^\top \hat{u} - \hat{e}^\top \hat{e})}{\hat{e}^\top \hat{e}/(n-2k)}.$$

F tests

(ii) for unknown shift point: sup$F$ test
Compute the test statistics F_i for all possible shift points in an interval $[\underline{i}, \bar{\imath}]$ and reject the null hypothesis if any of these F statistics is improbably large.

alternatively: ave$F$ and exp$F$ test
Reject the null hypothesis if the average or the exp functional of the F statistics is improbably large. These tests have certain optimality properties (Andrews & Ploberger, 1994).

F tests

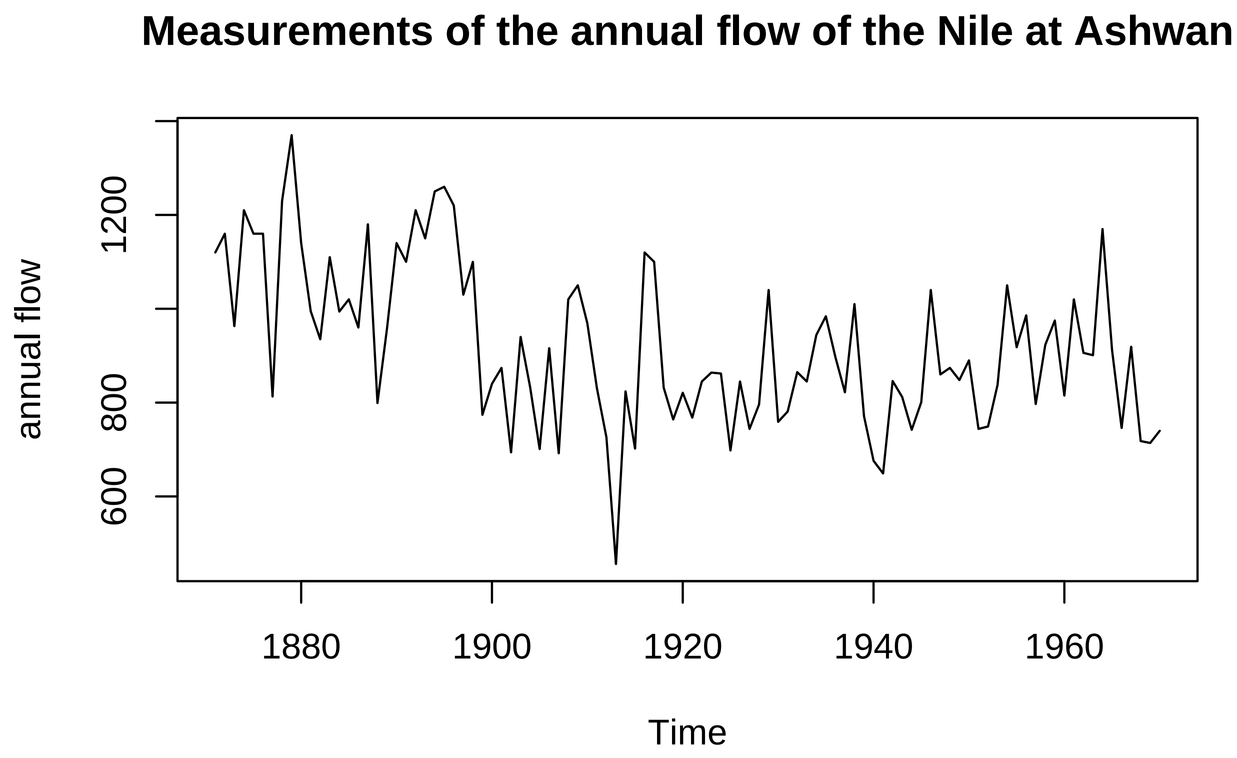
Measurements of the annual flow of the Nile at Ashwan
annual flow
600
800
1200
1880
1900
1920
1940
1960
Time

F tests

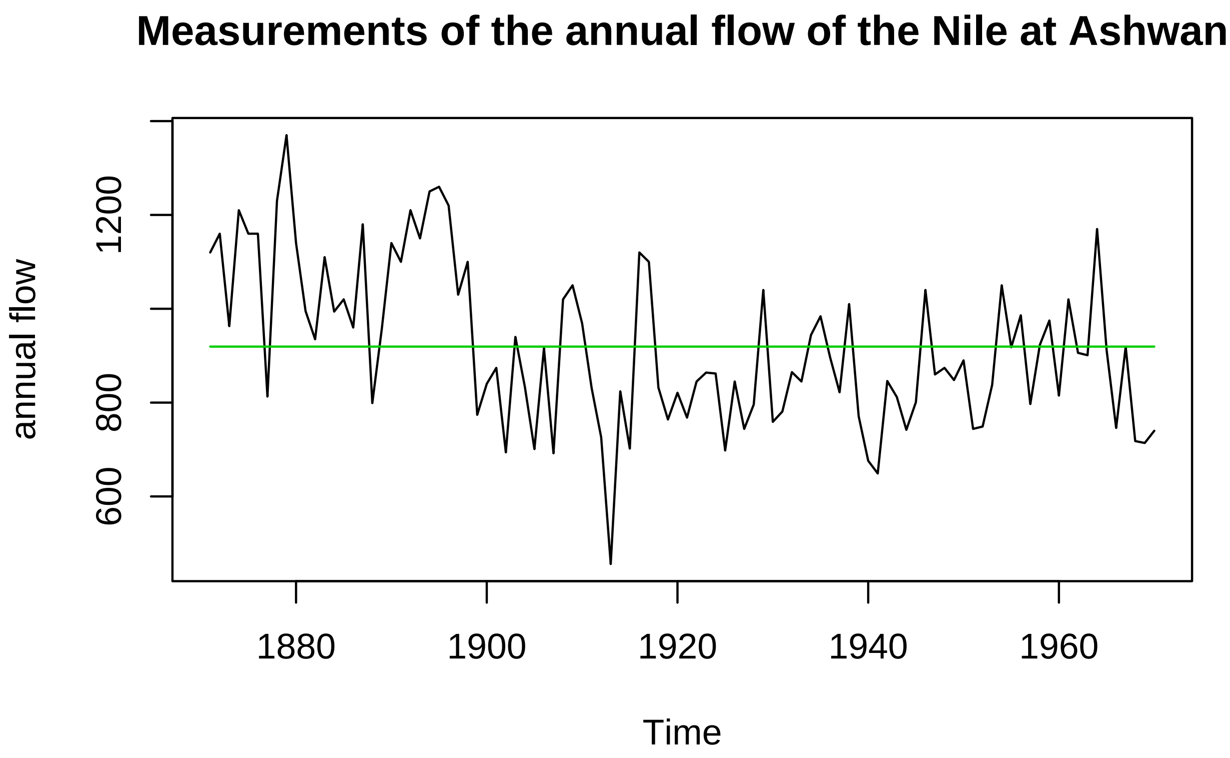

Measurements of the annual flow of the Nile at Ashwan
annual flow
600
800
1200
1880
1900
1920
1940
1960
Time

F tests

```
R> fs <- Fstats(Nile ~ 1)
R> plot(fs)
```

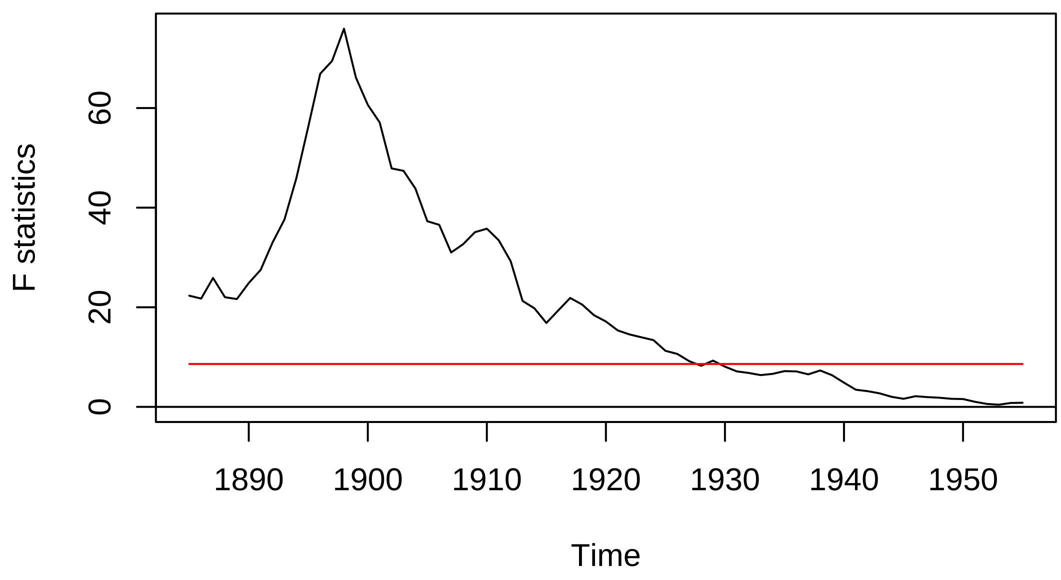

F tests

```
R> fs <- Fstats(Nile ~ 1)
R> plot(fs)

R> lines(breakpoints(fs))
```

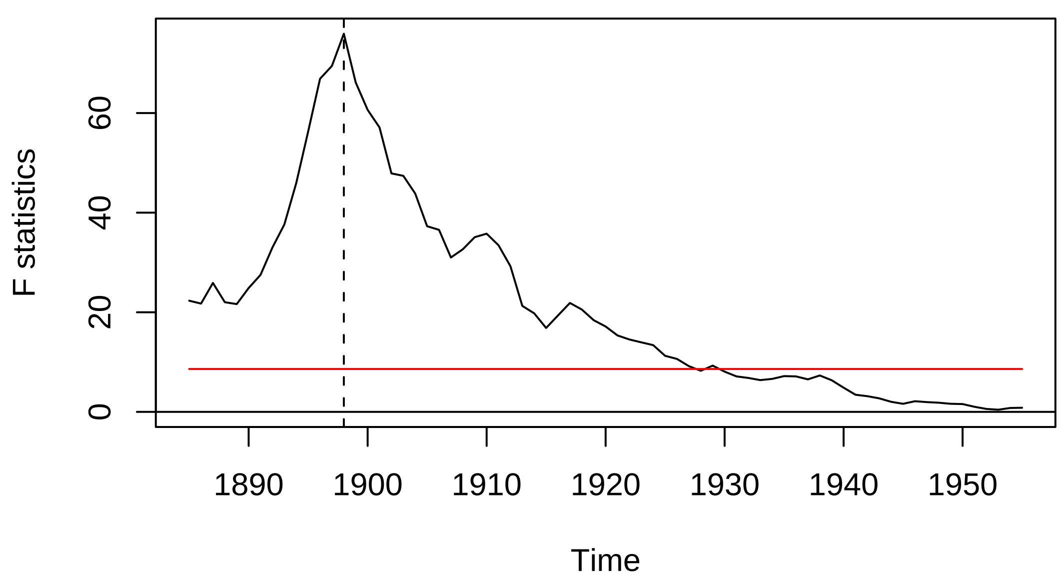

F tests

Additionally to these graphical methods significance tests can be carried out using the function `sctest()`:

```
R> sctest(fs)

        supF test

data:  fs
sup.F = 75.9298, p-value = 2.220e-16
```

F tests

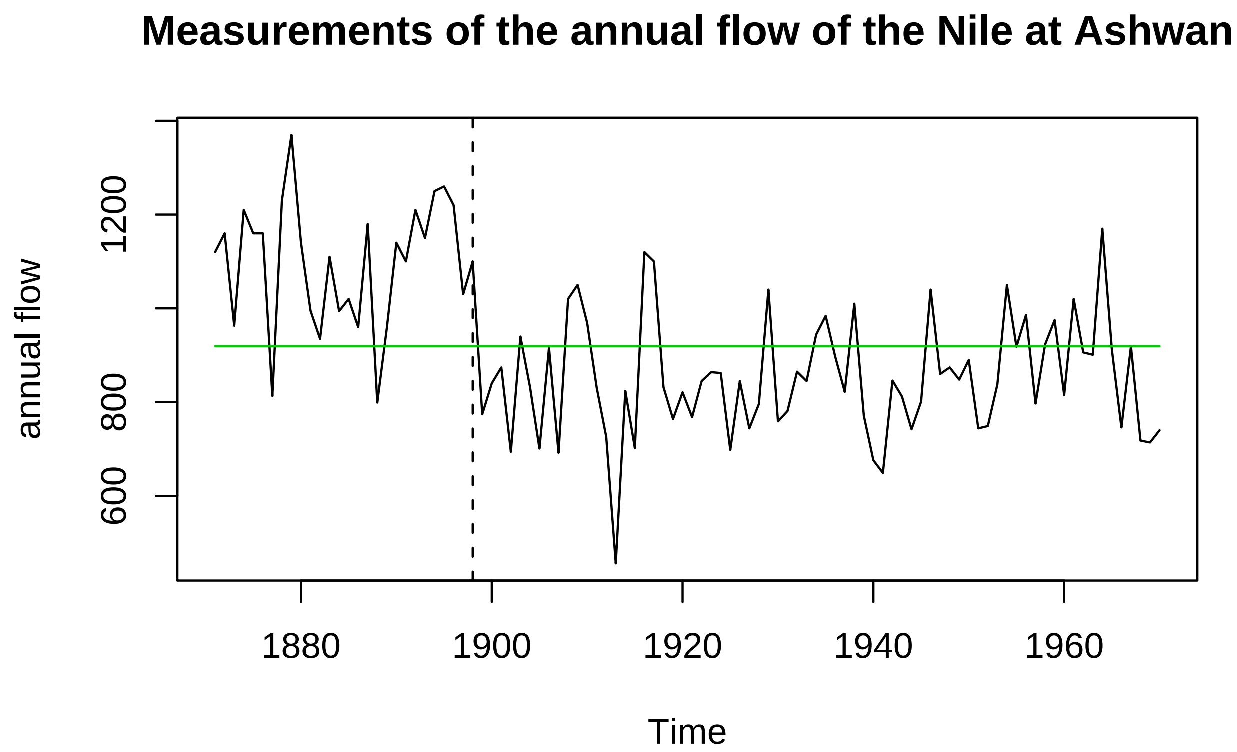
Measurements of the annual flow of the Nile at Ashwan
annual flow
600
800
1200
1880
1900
1920
1940
1960
Time

F tests

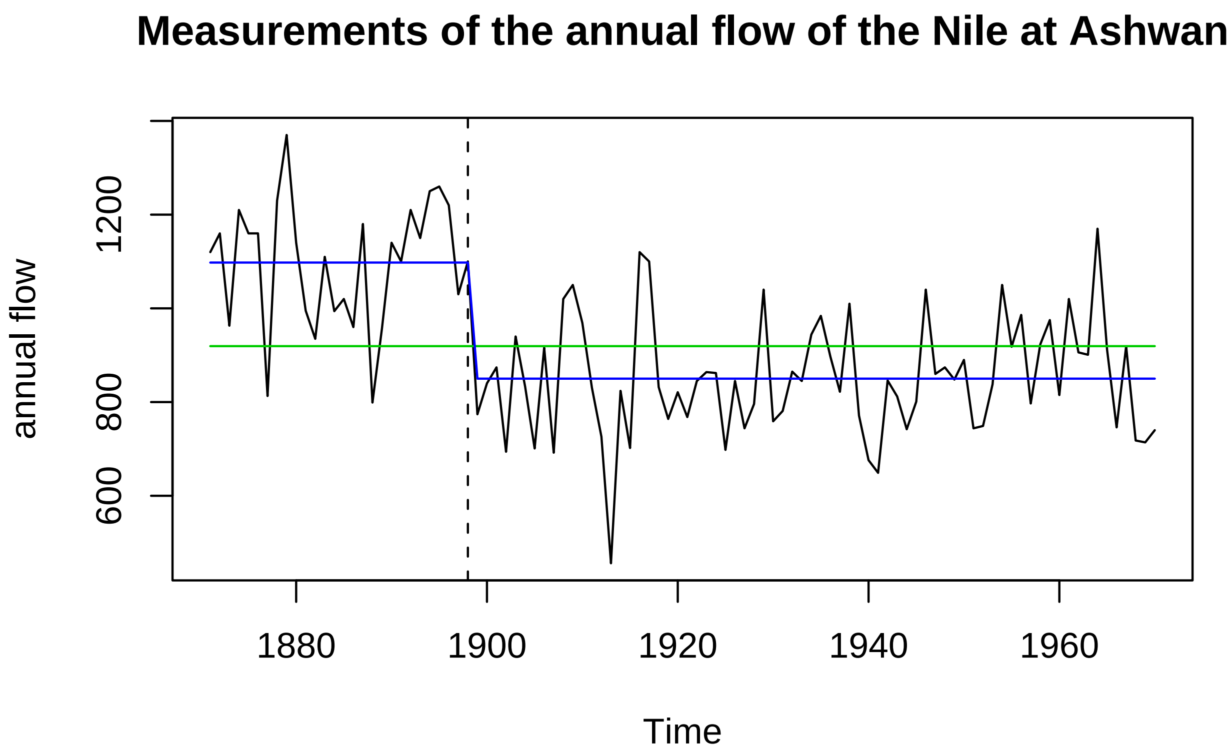

Measurements of the annual flow of the Nile at Ashwan
annual flow
600
800
1200
1880
1900
1920
1940
1960
Time

Generalized fluctuation tests

Generalized fluctuation tests are especially useful, if one has no particular alternative (H_1) in mind. This framework ...

"... includes formal significance tests but its philosophy is basically that of data analysis as expounded by Tukey. Essentially, the techniques are designed to bring out departures from constancy in a graphic way instead of parametrizing particular types of departure in advance and then developing formal significance tests intended to have high power against these particular alternatives." (Brown, Durbin, Evans, 1975)

Generalized fluctuation tests

- empirical fluctuation processes reflect fluctuation in
  - residuals (common OLS residuals or recursive residuals = 1-step prediction error)
  - coefficient estimates
- theoretical limiting process is known
- choose boundaries which are crossed by the limiting process only with a known probability α.
- if the empirical fluctuation process crosses the theoretical boundaries the fluctuation is improbably large $\Rightarrow$ reject the null hypothesis.

Generalized fluctuation tests

Processes based on OLS residuals:

$$\hat{u}_i \quad = \quad y_i - x_i^\top \hat{\beta}^{(n)}$$

OLS-based CUSUM process:

$$W_n^0(t) \quad = \quad \frac{1}{\hat{\sigma}\sqrt{n}} \sum_{i=1}^{\lfloor nt \rfloor} \hat{u}_i \qquad (0 \le t \le 1).$$

Generalized fluctuation tests

Processes based on OLS residuals:

$$\hat{u}_i \quad = \quad y_i - x_i^\top \hat{\beta}^{(n)}$$

OLS-based CUSUM process:

$$W_n^0(t) \quad = \quad \frac{1}{\hat{\sigma}\sqrt{n}} \sum_{i=1}^{\lfloor nt \rfloor} \hat{u}_i \qquad (0 \le t \le 1).$$

OLS-based MOSUM process:

$$M_n^0(t|h) \quad = \quad \frac{1}{\hat{\sigma}\sqrt{n}} \left(\sum_{i=\lfloor nt \rfloor + 1}^{\lfloor nt \rfloor + \lfloor nh \rfloor} \hat{u}_i \right) \qquad (0 \le t \le 1 - h).$$

Generalized fluctuation tests

```
R> ols <- efp(Nile ~ 1, type = "OLS-CUSUM")
R> plot(ols)
```

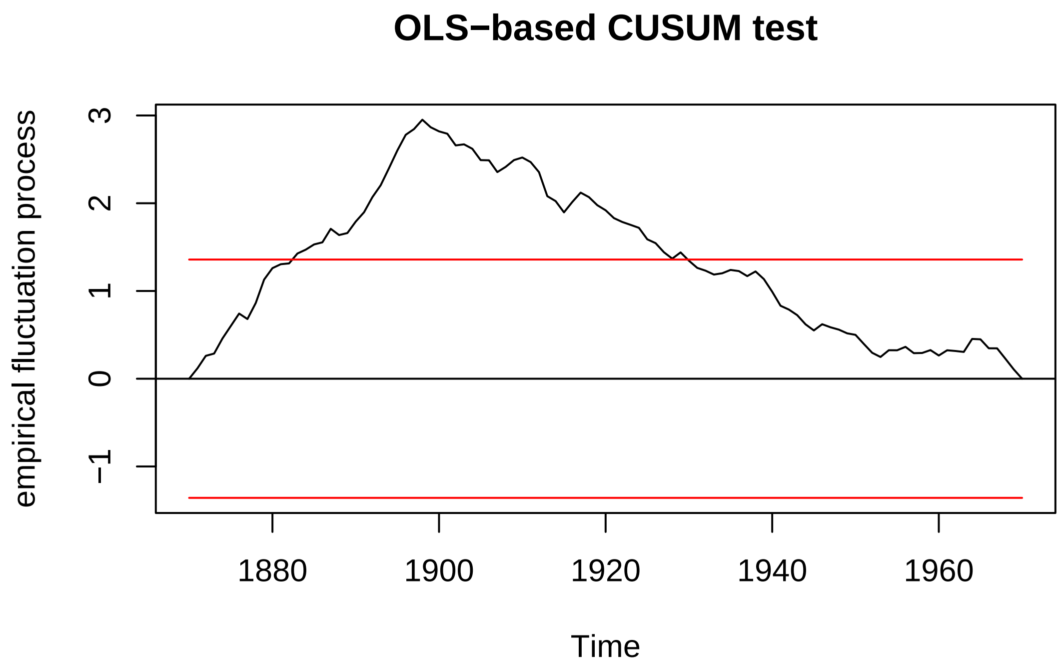

Generalized fluctuation tests

```
R> ols <- efp(Nile ~ 1, type = "OLS-CUSUM")
R> plot(ols)

R> lines(breakpoints(fs))
```

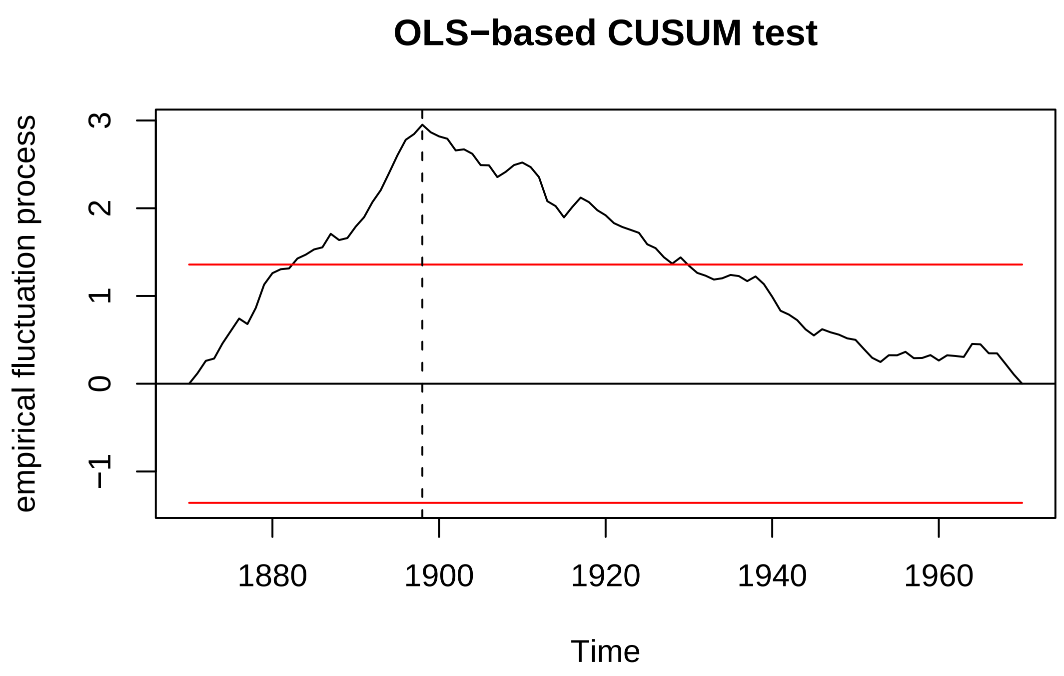

Generalized fluctuation tests

As for the F statistics significance tests can be carried out using the function `sctest()`:

```
R> sctest(ols)


        OLS-based CUSUM test


data:  ols
S0 = 2.9518, p-value = 5.409e-08
```

The GLM

Consider the generalized linear model (GLM) with n independent observations from an exponential family

$$y_i \sim F(\mu_i, \phi) \quad \text{with } \mathsf{E}[y_i] = \mu_i \text{ and dispersion } \phi.$$

The following regression relationship is assumed:

$$g(\mu_i) \quad = \quad x_i^\top \beta_i \qquad (i = 1, \dots, n),$$

where $g(\cdot)$ is a known link function and β_i is a vector of regression coefficients.

The null hypothesis remains

$$H_0: \quad \beta_i \quad = \quad \beta_0.$$

Quantile residuals

Dunn & Smyth (1996) suggest quantile residuals for GLMs:

(i) F **continuous**: The quantile residuals are defined as:

$$r_i = \Phi^{-1}\{F(y_i|\widehat{\mu}_i, \widehat{\phi})\},$$

and are asymptotically standard normal.

Quantile residuals

Dunn & Smyth (1996) suggest quantile residuals for GLMs:

(i) F **continuous**: The quantile residuals are defined as:

$$r_i = \Phi^{-1}\{F(y_i|\widehat{\mu}_i, \widehat{\phi})\},$$

and are asymptotically standard normal.

(ii) F **discrete**: To obtain asymptotic normality the quantile residuals have to be randomized:

$$r_i = \Phi^{-1}(u_i),$$

where u_i is a uniform random variable on $(q_i, p_i]$ with

$$q_i = \lim_{y \uparrow y_i} F(y|\widehat{\mu}_i, \widehat{\phi}),$$
$$p_i = F(y_i|\widehat{\mu}_i, \widehat{\phi})].$$

Illegitimate births

Veichtlbauer, Hanser, Zeileis, Leisch (2002) analyze the annual number of legitimate and illegitimate births in Großarl between 1700 and 1900.
Question: Did moral regulations have any effect on the proportion of illegitimate births?

Two models are compared: a binomial model explaining the illegitimate births just by political influences or additionally by moral regulations (and the number of marriages in the previous year).

Illegitimate births

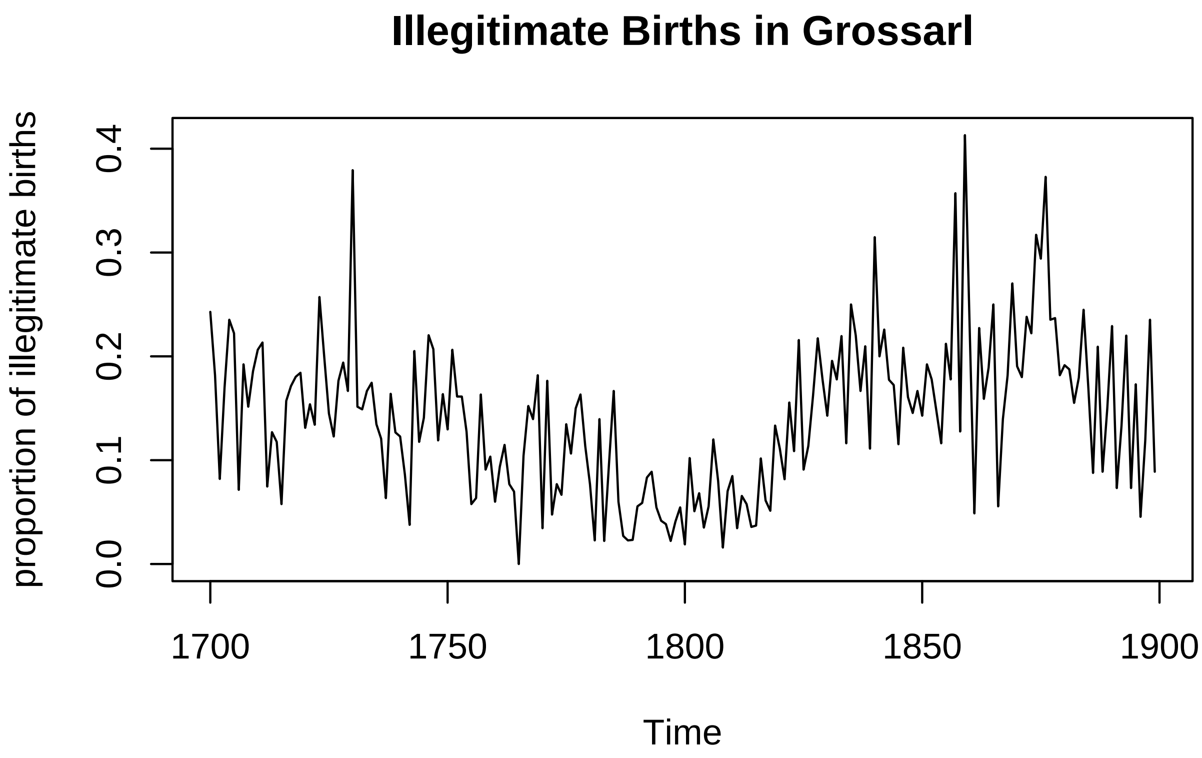

Illegitimate births

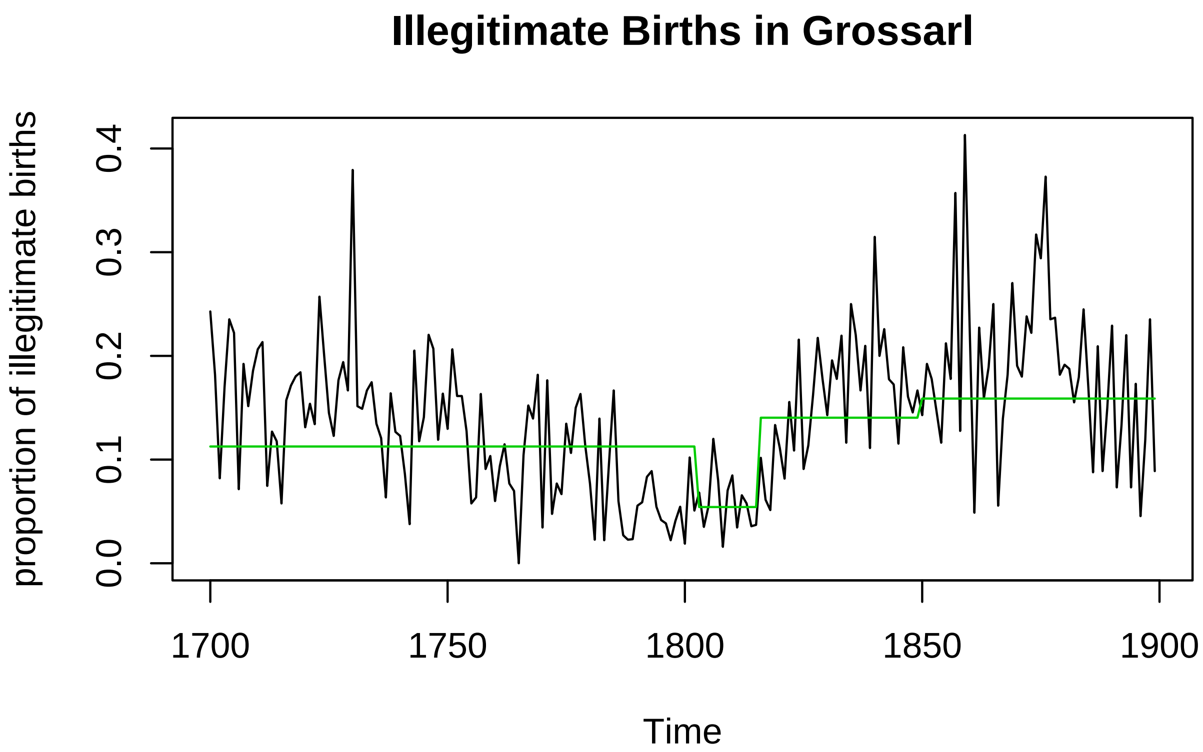

Illegitimate births

```
R> rqr <- gefp(baptisms ~ politics, type = "RQR-CUSUM", family = binomial)
R> plot(rqr)
```

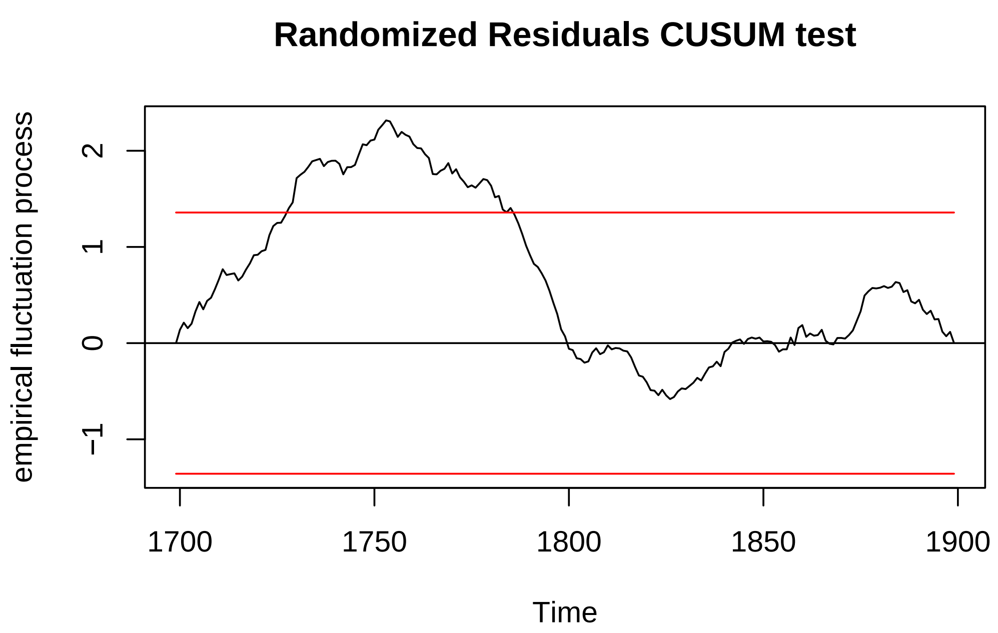

Illegitimate births

```
R> rqr <- gefp(baptisms ~ politics, type = "RQR-CUSUM", family = binomial)
R> plot(rqr)

R> lines(gefp(baptisms ~ politics + morals + lag.marriages,
              type = "RQR-CUSUM", family = binomial))
```

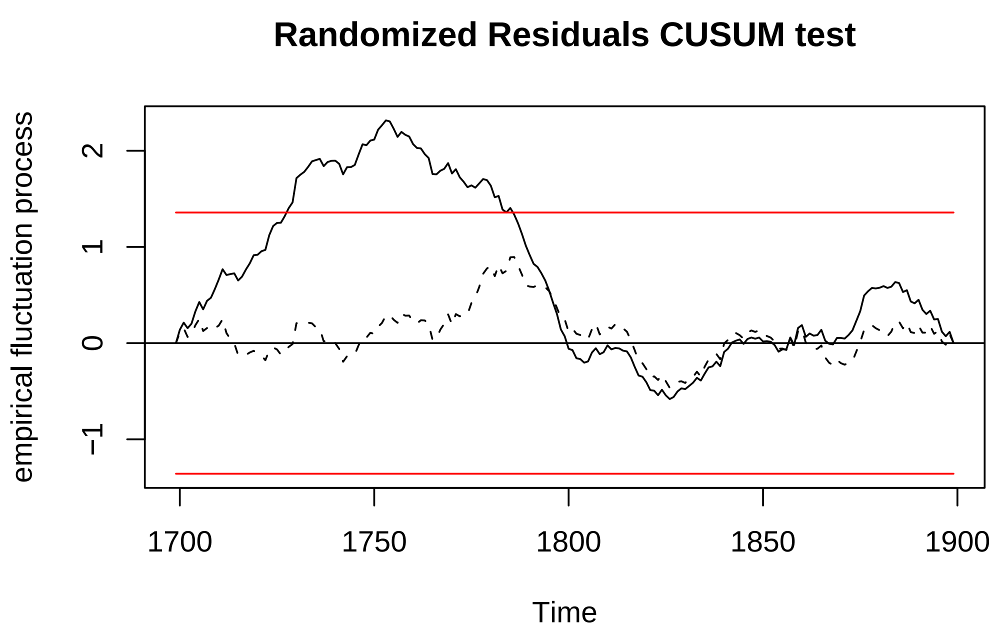

Illegitimate births

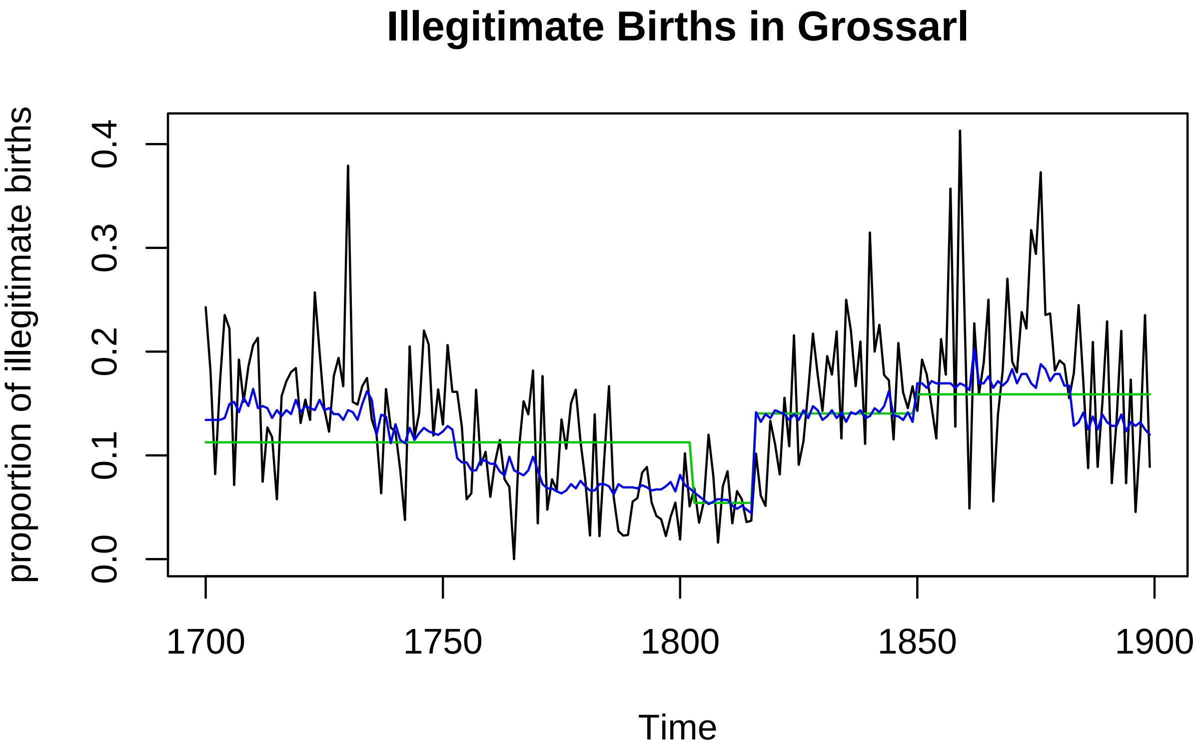

Score-based processes

Instead of capturing the fluctuation in residuals or estimates, fluctuation processes can also be based on Maximum Likelihood (ML) scores (Hjort & Koning, 2002):

$$W_n(t) \quad = \quad \frac{1}{\sqrt{n}} Q(\hat{\beta})^{-1/2} \sum_{i=1}^{\lfloor nt \rfloor} \log L'(y_i | \hat{\beta}) \qquad (0 \leq t \leq 1),$$

where

$$Q(\beta) \quad = \quad -\frac{1}{n} \sum_{i=1}^{n} \log L''(y_i | \beta).$$

Then W_n converges in distribution to a vector of k independent Brownian bridges.

Score-based processes

For a Poisson regression model with the canonical log link this leads to

$$W_n(t) \quad = \quad \frac{1}{\sqrt{n}} \, \widehat{Q}^{-1/2} \sum_{i=1}^{\lfloor nt \rfloor} \{y_i - \mathsf{exp}(x_i^\top \widehat{\beta})\} x_i,$$

where

$$\widehat{Q} \quad = \quad \frac{1}{n} \sum_{i=1}^{n} \mathsf{exp}(x_i^\top \widehat{\beta}) x_i x_i^\top.$$

If regressing only on a constant (i.e., $x_i \equiv 1$), this is equivalent to a fluctuation process based on Pearson residuals.

Boston Homicides

Cooper, Piehl, Braga, Kennedy (2001) investigate the monthly number of youth homicides in Boston: *Does the "Boston Gun Project" have any effect on the number of youth homicides?*

The project was launched in the mid-1990s due to concern about the growing numbers of youth involved in homicides.

Boston Homicides

Monthly Youth Homicide Count in Boston

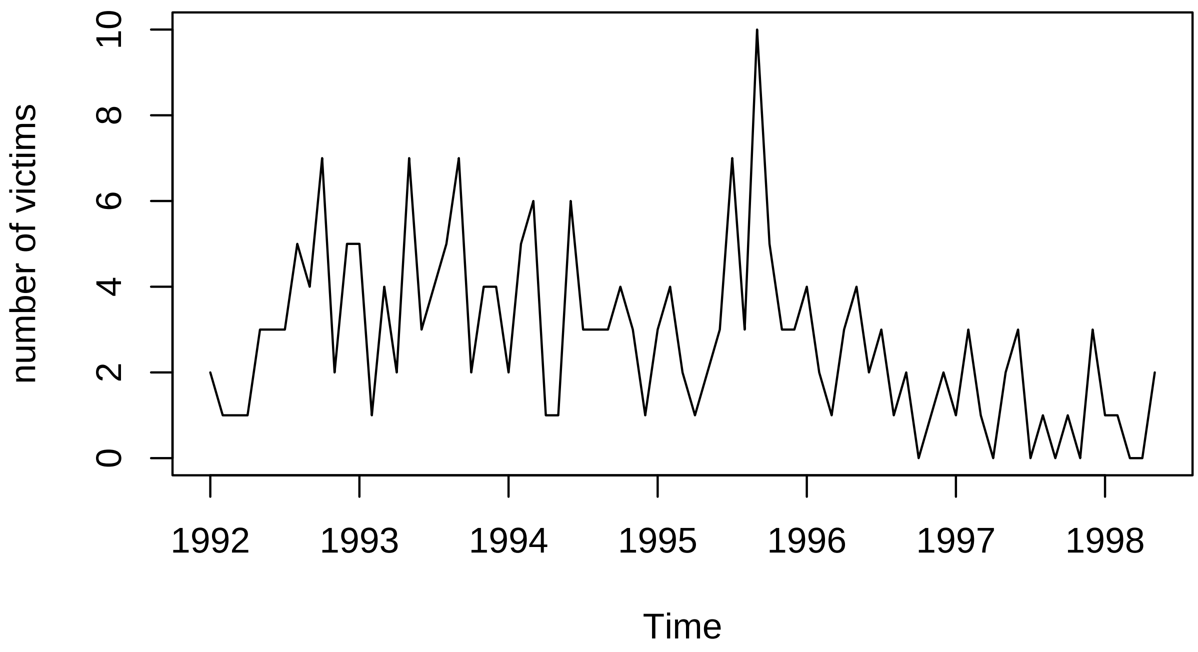

Boston Homicides

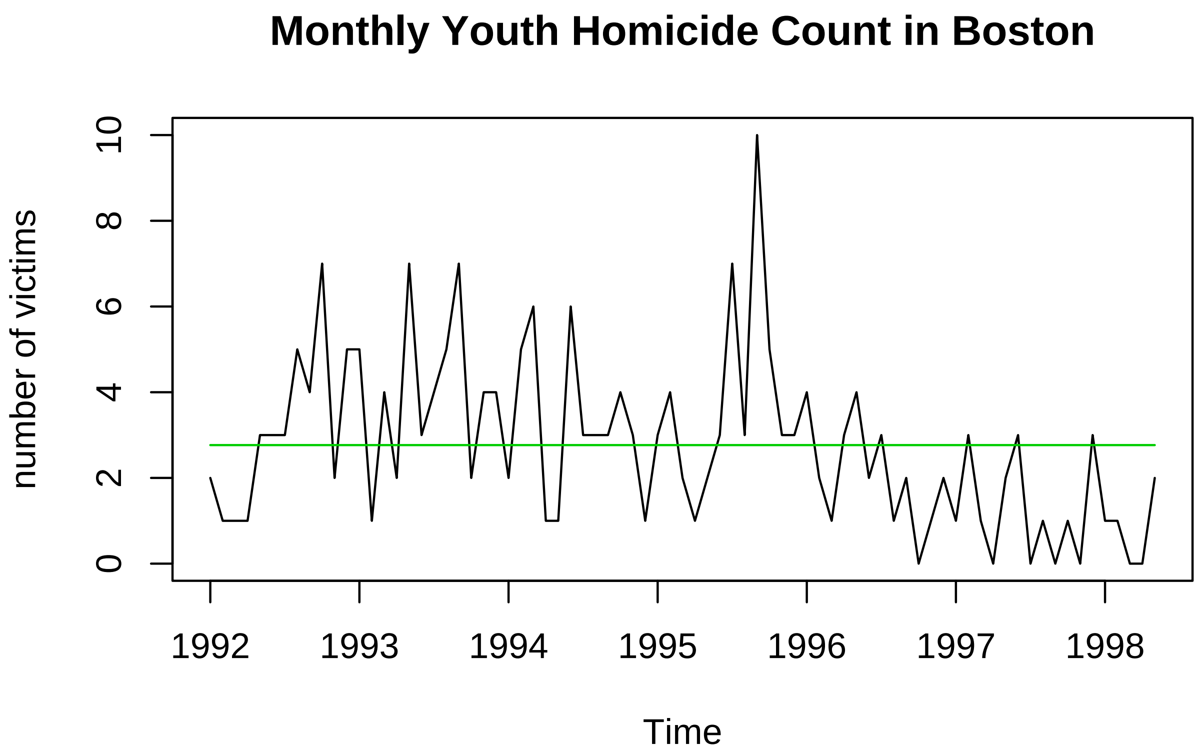

Boston Homicides

```
R> sco <- gefp(BostonHomicide ~ 1, type = "Score-CUSUM", family = poisson)
R> plot(sco)
```

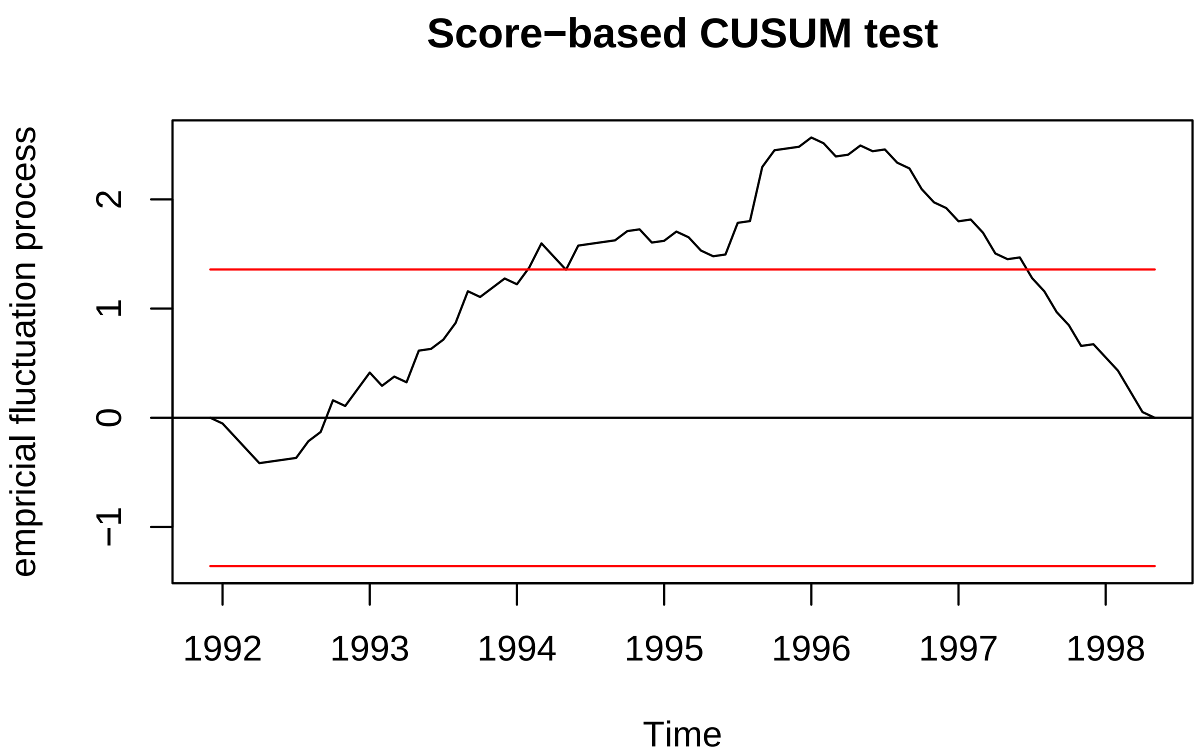

Boston Homicides

```
R> sco <- gefp(BostonHomicide ~ 1, type = "Score-CUSUM", family = poisson)

R> plot(sco)
```

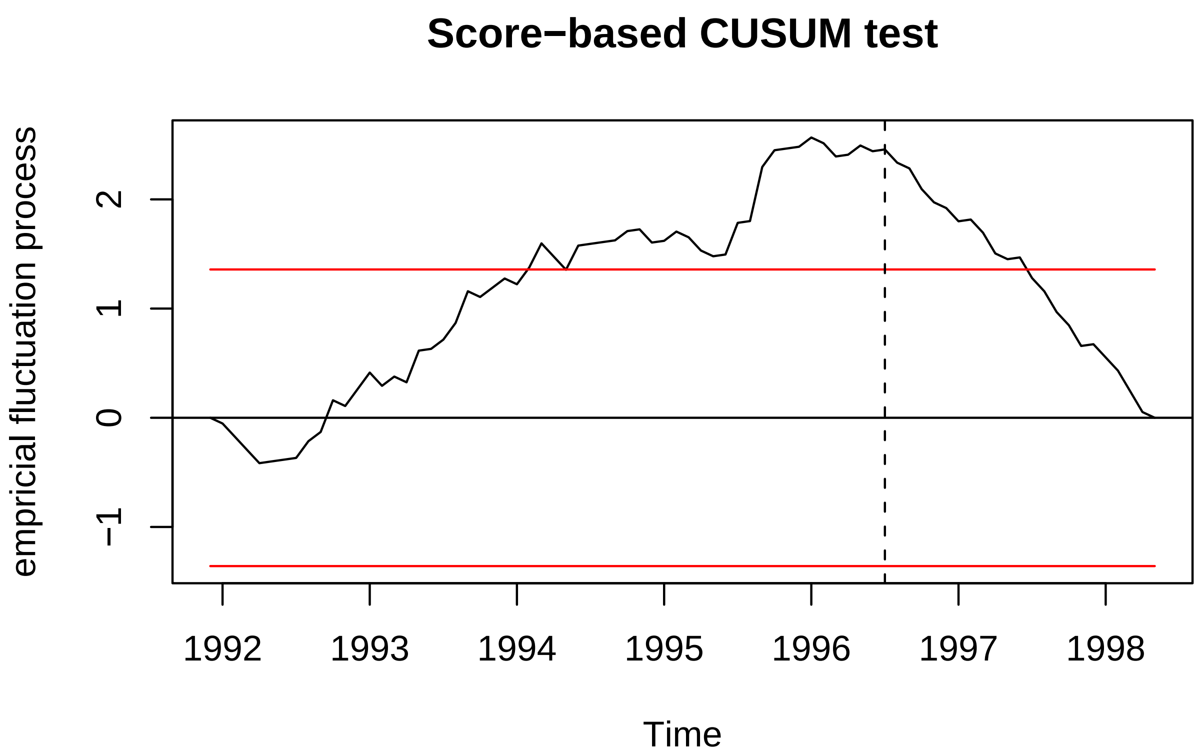

Boston Homicides

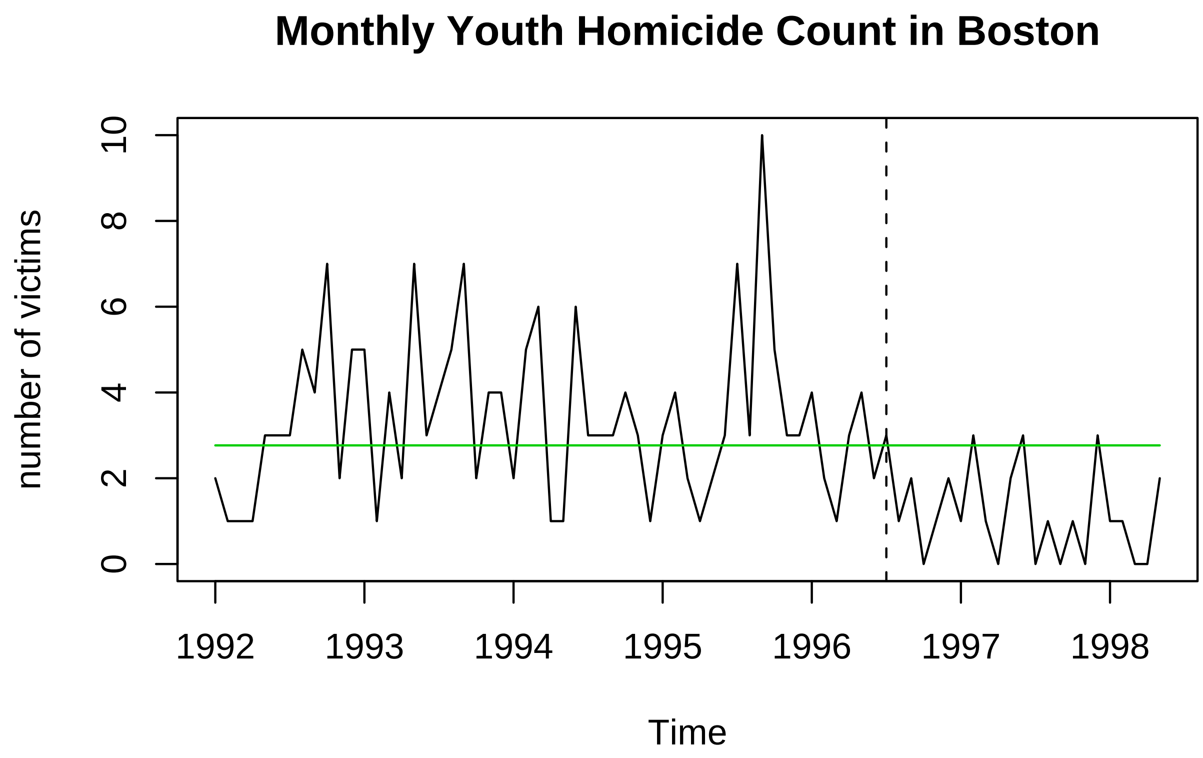

Boston Homicides