Universität Innsbruck

Testing, Monitoring, and Dating Structural Changes in FX Regimes

Achim Zeileis

http://statmath.wu.ac.at/~zeileis/

Overview

- Motivation
  - What is the new Chinese exchange rate regime?
  - Exchange rate regimes
  - Exchange rate regression
- Structural change tools
  - Model frame
  - Testing
  - Monitoring
  - Dating
- Software: *strucchange*, *fxregime*
- Application: Indian exchange rate regimes
- Summary

Motivation

Initial impulse: Ajay Shah, long time *R-help* and *R-SIG-Finance* contributor, contacts Achim Zeileis, *strucchange* package maintainer.

```
Date: Thu, 28 Jul 2005 21:57:10 +0530
From: Ajay Narottam Shah <ajayshah@mayin.org>
To: Achim Zeileis <Achim.Zeileis@wu-wien.ac.at>
Subject: Wonder if this fits (structural breaks work in
a currency regime context)

...

The issues are like this. Many central banks SAY that a
currency regime is X. But they routinely lie. Economists
would like to know the true currency regime. And, we
would like to know the date when something changed.

...
```

Motivation

Of particular interest: China gave up on a fixed exchange rate to the US dollar (USD) on 2005-07-21. The People's Bank of China announced that the Chinese yuan (CNY) would no longer be pegged to the USD but to a basket of currencies with greater flexibility.

Collaboration: Ajay Shah, Ila Patnaik, and Achim Zeileis start to investigate the question *What is the new Chinese exchange rate regime?*

First step: Collect foreign exchange (FX) rates for various currencies for three months up to 2005-10-31.

Motivation

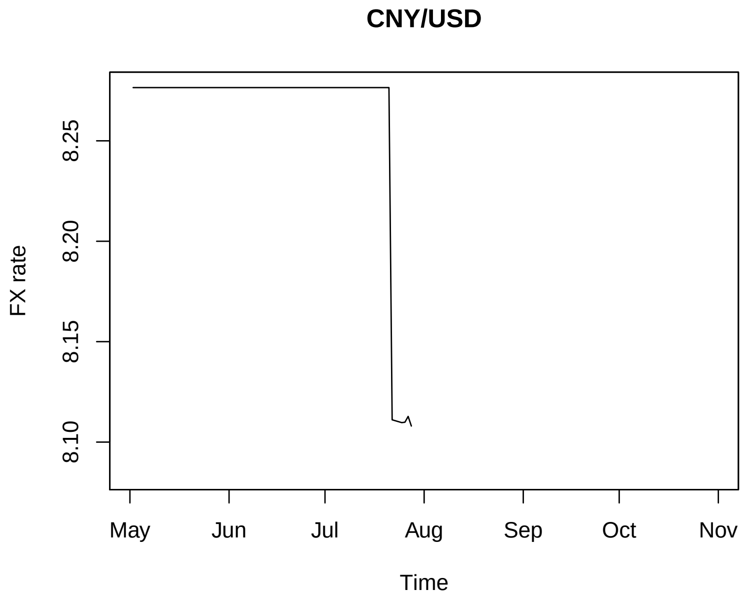

Motivation

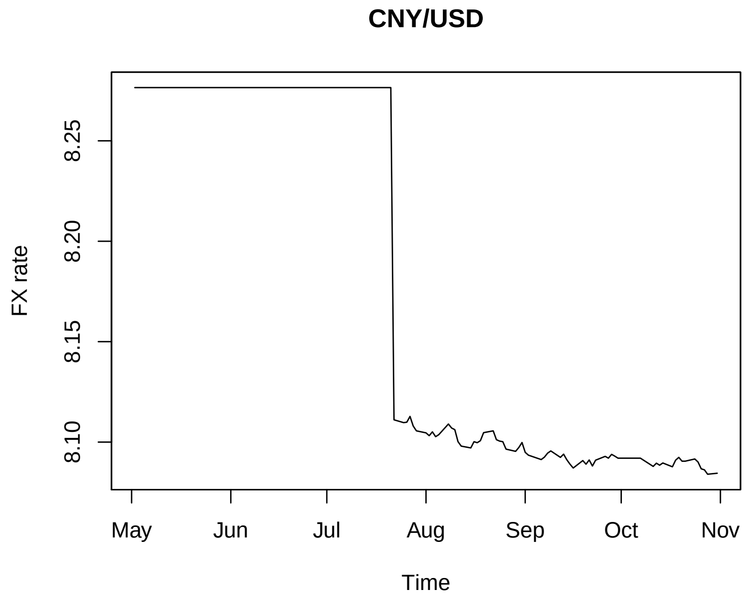

Exchange rate regimes

The FX regime of a country determines how it manages its currency wrt foreign currencies. Broadly, it can be

- *floating:* currency is allowed to fluctuate based on market forces,
- *pegged:* currency has limited flexibility when compared with a basket of currencies or a single currency,
- *fixed:* direct convertibility to another currency.

Problem: The *de facto* and *de jure* FX regime in operation in a country often differ. *(≈ politically correct version of Ajay's original e-mail)*

⇒ Data-driven classification of FX regimes

Exchange rate regression

The workhorse for de facto FX regime classification is a linear regression model based on log-returns of cross-currency exchange rates (with respect to some floating reference currency). In the literature, this is also known as *Frankel-Wei regression*.

For modeling the log-returns of CNY a basket of regressors USD, JPY, EUR, and GBP (all log-returns wrt CHF) is employed.

Fitting the model for the first three months (up to 2005-10-31, $n = 68$) shows that a plain USD peg is still in operation.

Exchange rate regression

Ordinary least squares (OLS) estimation gives:

$$\text{CNY}_i = \underset{(0.004)}{0.005} + \underset{(0.009)}{0.9997}\,\text{USD}_i + \underset{(0.011)}{0.005}\,\text{JPY}_i$$

$$- \underset{(0.027)}{0.014}\,\text{EUR}_i - \underset{(0.015)}{0.008}\,\text{GBP}_i + \widehat{\varepsilon}_i$$

Only the USD coefficient is significantly different from 0 (but not from 1).

The error standard deviation is tiny with $\widehat{\sigma} = 0.028$ leading to $R^2 = 0.998$.

Exchange rate regression

Questions:

1. Is this model for the period 2005-07-26 to 2005-10-31 stable or is there evidence that China kept changing its FX regime after 2005-07-26? (***testing***)
2. Depending on the answer to the first question:
   - Does the CNY stay pegged to the USD in the future (starting from November 2005? (***monitoring***)
   - When and how did the Chinese FX regime change? (***dating***)

Exchange rate regression

In practice: Rolling regressions are often used to answer these questions by tracking the evolution of the FX regime in operation.

More formally: Structural change techniques can be adapted to the FX regression to estimate and test the stability of FX regimes.

Problem: Unlike many other linear regression models, the stability of the error variance (fluctuation band) is of interest as well.

Solution: Employ an (approximately) normal regression estimated by ML where the variance is a full model parameter.

Model frame

Generic idea: Consider a regression model for n ordered observations $y_i \mid x_i$ with k-dimensional parameter θ. Ordering is typically with respect to time in time-series regressions, but could also be with respect to income, age, etc. in cross-section regressions.

To fit the model to observations $i = 1, \dots, n$ an objective function $\Psi(y, x, \theta)$ is used such that

$$\widehat{\theta} = \underset{\theta}{\text{argmin}} \sum_{i=1}^{n} \Psi(y_i, x_i, \theta).$$

This can also be defined implicitly based on the corresponding score function (or estimating function) $\psi(y, x, \theta) = \partial\Psi(y, x, \theta)/\partial\theta$:

$$\sum_{i=1}^{n} \psi(y_i, x_i, \widehat{\theta}) = 0.$$

Model frame

This class of M-estimators includes OLS and maximum likelihood (ML) estimation as well as IV, Quasi-ML, robust M-estimation etc.

Under parameter stability and some mild regularity conditions, a central limit theorem holds

$$\sqrt{n}(\widehat{\theta} - \theta_0) \xrightarrow{d} \mathcal{N}(0, V(\theta_0)),$$

where the covariance matrix is

$$V(\theta_0) = \{A(\theta_0)\}^{-1} B(\theta_0) \{A(\theta_0)\}^{-1}$$

and A and B are the expectation of the derivative of ψ and its variance respectively.

Model frame

For the standard linear regression model

$$y_i \; = \; x_i^\top \beta \; + \; \varepsilon_i$$

with coefficients β and error variance σ^2 one can either treat σ^2 as a nuisance parameter $\theta = \beta$ or include it as $\theta = (\beta, \sigma^2)$.

In the former case, the estimating functions are $\psi = \psi_\beta$

$$\psi_\beta(y, x, \beta) \; = \; (y - x^\top \beta) \, x$$

and in the latter case, they have an additional component

$$\psi_{\sigma^2}(y, x, \beta, \sigma^2) \; = \; (y - x^\top \beta)^2 - \sigma^2.$$

and $\psi = (\psi_\beta, \psi_{\sigma^2})$. This is used for FX regressions.

Model frame

Testing: Given that a model with parameter $\widehat{\theta}$ has been estimated for these n observations, the question is whether this is appropriate or: *Are the parameters stable or did they change through the sample period $i = 1, \dots, n$?*

Monitoring: Given that a stable model could be established for these n observations, the question is whether it remains stable in the future or: *Are incoming observations for $i > n$ still consistent with the established model or do the parameters change?*

Dating: Given that there is evidence for a structural change in $i = 1, \dots, n$, it might be possible that stable regression relationships can be found on subsets of the data. *How many segments are in the data? Where are the breakpoints?*

Testing

To assess the stability of the fitted model with $\widehat{\theta}$, we want to test the null hypothesis

$$H_0: \; \theta_i = \theta_0 \qquad (i = 1, \dots, n)$$

against the alternative that θ_i varies over "time" i.

Various patterns of deviation from H_0 are conceivable: single/multiple break(s), random walks, etc.

To test this null hypothesis, the basic idea is to assess wether the empirical estimating functions $\hat{\psi}_i = \psi(y_i, x_i, \widehat{\theta})$ deviate systematically from their theoretical zero mean.

Testing

To capture systematic deviations the empirical fluctuation process of scaled cumulative sums of empirical estimating functions is computed:

$$efp(t) \;=\; \widehat{B}^{-1/2}\, n^{-1/2} \sum_{i=1}^{\lfloor nt \rfloor} \hat{\psi}_i \qquad (0 \le t \le 1).$$

Under H_0 the following functional central limit theorem (FCLT) holds:

$$efp(\cdot) \;\xrightarrow{d}\; W^0(\cdot),$$

where W^0 denotes a standard k-dimensional Brownian bridge.

Testing

Testing procedure:

- empirical fluctuation processes captures fluctuation in estimating functions
- theoretical limiting process is known
- choose boundaries which are crossed by the limiting process (or some functional of it) 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.

Testing

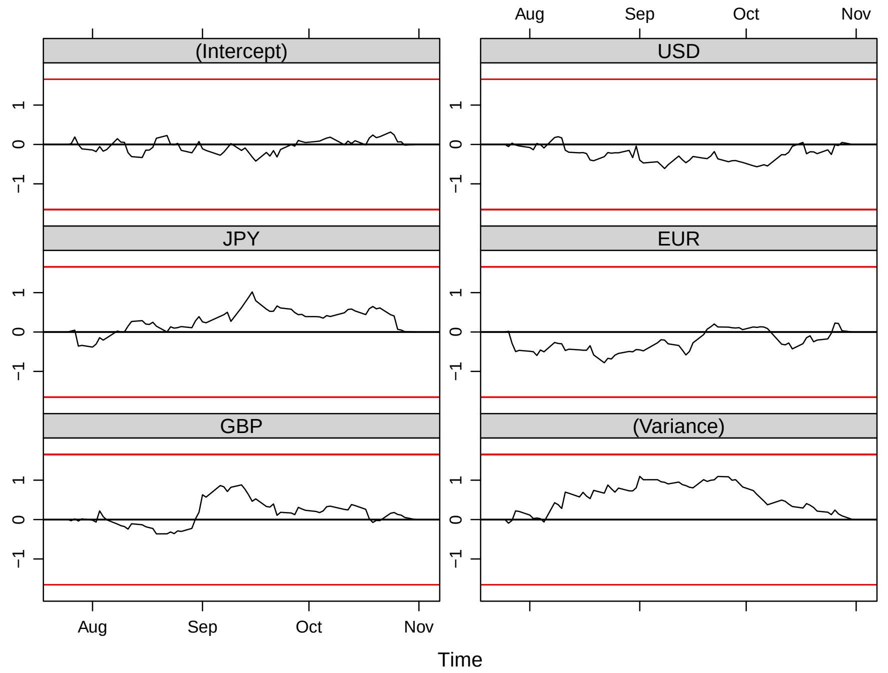

Testing

More formally: These boundaries correspond to critical values for a double maximum test statistic

$$\max_{j=1,\dots,k}\ \max_{i=1,\dots,n} |efp_j(i/n)|$$

which is 1.097 for the Chinese FX regression ($p = 0.697$).

Alternatively: Employ other test statistics for aggregation.

Special cases: This class contains various well-known tests from the statistics and econometrics literature, e.g., Andrews' sup*LM* test, Nyblom-Hansen test, OLS-based CUSUM/MOSUM tests.

Testing

In empirical samples, $efp(\cdot)$ is a $k \times n$ array. For significance testing, aggregate it to a scalar test statistic by a functional $\lambda(\cdot)$

$$\lambda \left(efp_j \left(\frac{i}{n} \right) \right),$$

where $j = 1, \dots, k$ and $i = 1, \dots n$.

Typically, $\lambda(\cdot)$ can be split up into

- $\lambda_{\text{comp}}(\cdot)$ aggregating over components j (e.g., absolute maximum, Euclidian norm),
- $\lambda_{\text{time}}(\cdot)$ aggregating over time i (e.g., max, mean, range).

The limiting distribution is given by $\lambda(W^0)$ and can easily be simulated (or some closed form results are also available).

Testing

Nyblom-Hansen test: The test was designed for a random-walk alternative and employs a Cramér-von Mises functional.

$$\frac{1}{n} \sum_{i=1}^{n} \left|\left| efp \left(\frac{i}{n} \right) \right|\right|_2^2 .$$

It aggregates $efp(\cdot)$ over the components first, using the squared Euclidian norm, and then over time, using the mean.

For the Chinese FX regression this is 1.012 ($p = 0.364$).

Testing

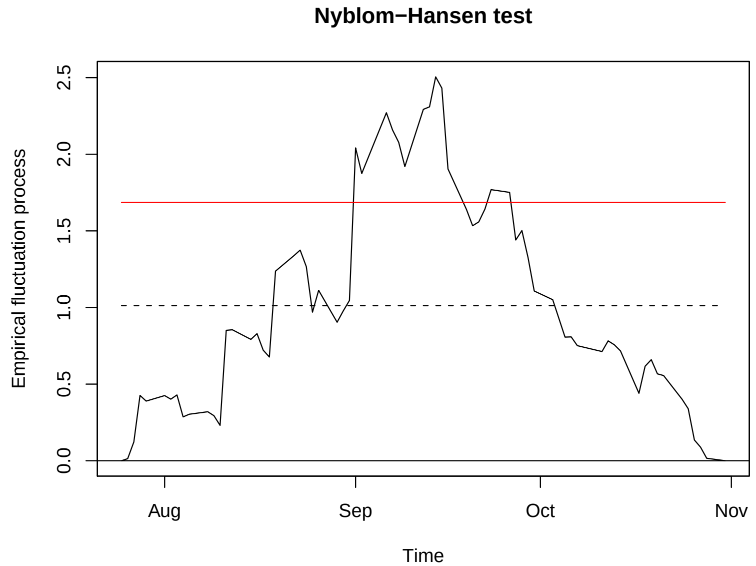

Testing

Andrews' sup*LM* test: This test is designed for a single shift alternative (with unknown timing) and employs the supremum of *LM* statistics for this alternative.

$$\sup_{t \in \Pi} LM(t) \quad = \quad \sup_{t \in \Pi} \frac{||efp(t)||_2^2}{t\,(1-t)}.$$

It aggregates $efp(\cdot)$ over the components first, using a weighted squared Euclidian norm, and then over time, using the maximum (over a compact interval $\Pi \subset [0, 1]$).

For the Chinese FX regression this is 10.055 ($p = 0.766$), using $\Pi = [0.1, 0.9]$.

Testing

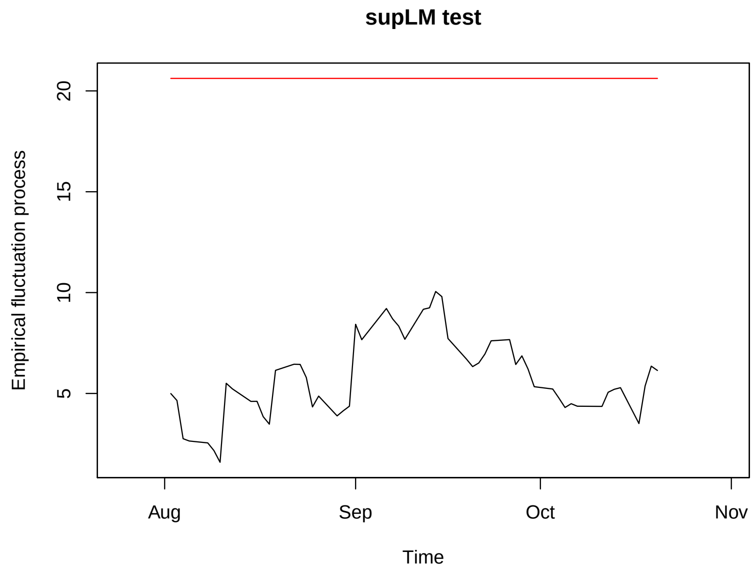

Monitoring

Idea: Fluctuation tests can be applied sequentially to monitor regression models.

More formally: Sequentially test the null hypothesis

$$H_0 : \ \theta_i = \theta_0 \qquad (i > n)$$

against the alternative that θ_i changes at some time in the future $i > n$ (corresponding to $t > 1$).

Basic assumption: The model parameters are stable $\theta_i = \theta_0$ in the history period $i = 1, \dots, n$ ($0 \le t \le 1$).

Monitoring

Test statistics: Update $efp(t)$, and re-compute $\lambda(efp(t))$ in the monitoring period $1 \le t \le T$.

Critical values: For sequential testing not only a single critical value is needed, but a full boundary function $b(t)$ that satisfies

$$1 - \alpha \ = \ \mathsf{P}(\lambda(W^0(t)) \le b(t) \ | \ t \in [1, T])$$

Various boundary (or weighting) functions are conceivable that can direct power to early or late changes or try to spread the power evenly.

In 2005: Ajay Shah, Ila Patnaik, and Achim Zeileis establish a webpage and start monitoring the CNY regime. A double maximum functional with boundary $b(t) = c \cdot t$ is employed (where c controls the significance level, using $T = 4$ and $\alpha = 0.05$).

Monitoring

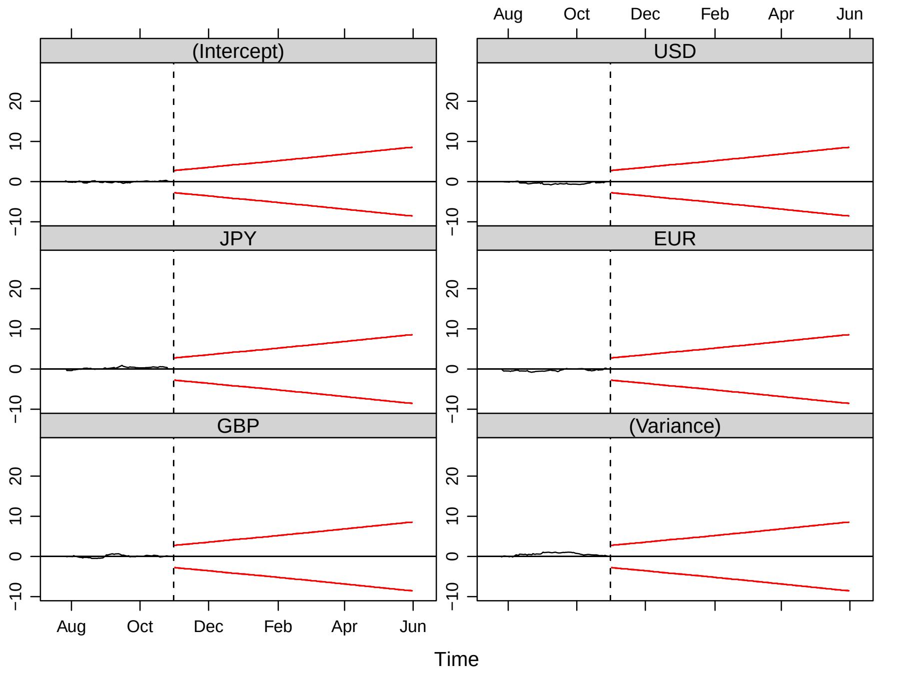

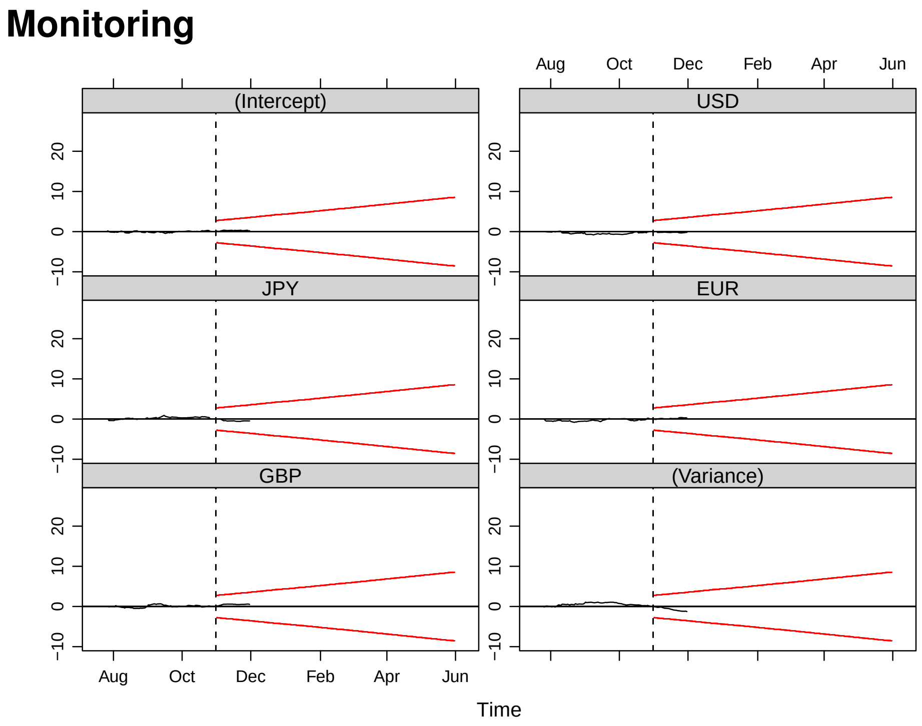
Monitoring
(Intercept)
USD
JPY
EUR
GBP
(Variance)
20 10 0 −10
Aug Oct Dec Feb Apr Jun
Time

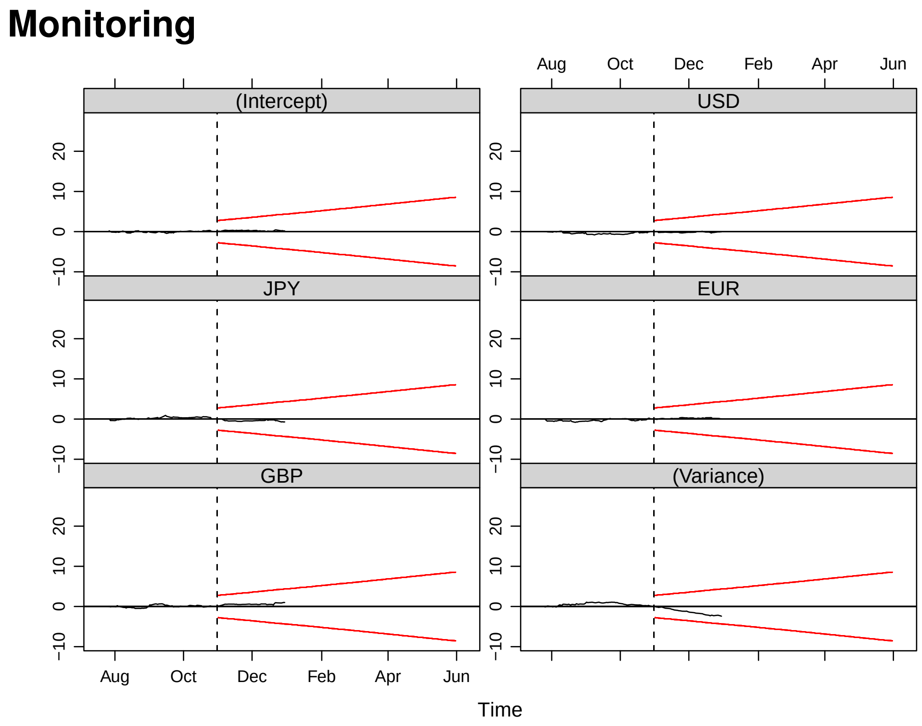
Monitoring
(Intercept)
USD
JPY
EUR
GBP
(Variance)
20 10 0 −10
Aug Oct Dec Feb Apr Jun
Time

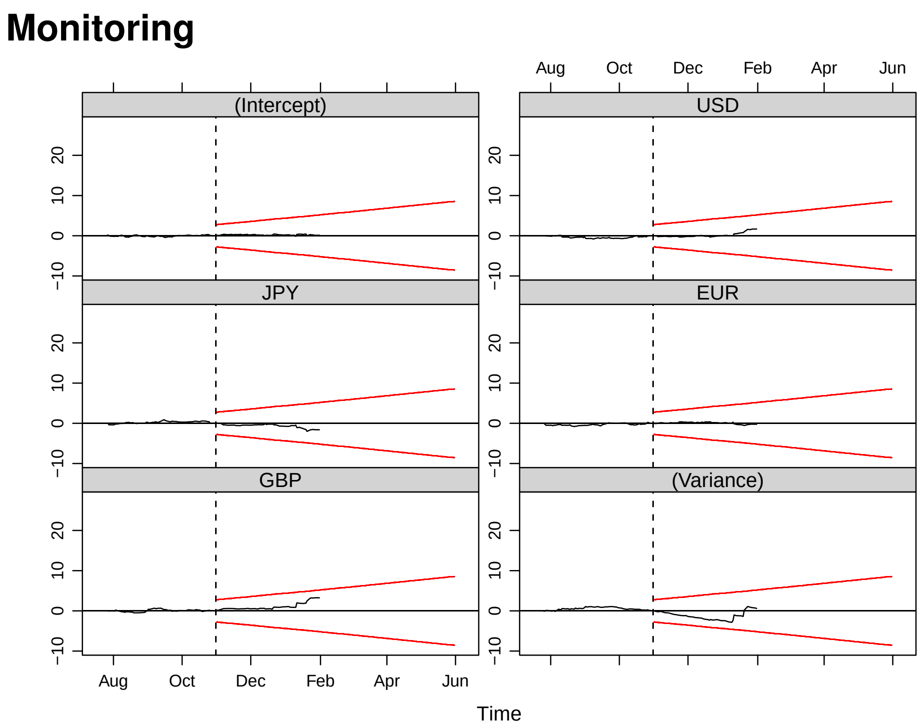
Monitoring
(Intercept)
USD
JPY
EUR
GBP
(Variance)
20 10 0 −10
Aug Oct Dec Feb Apr Jun
Time

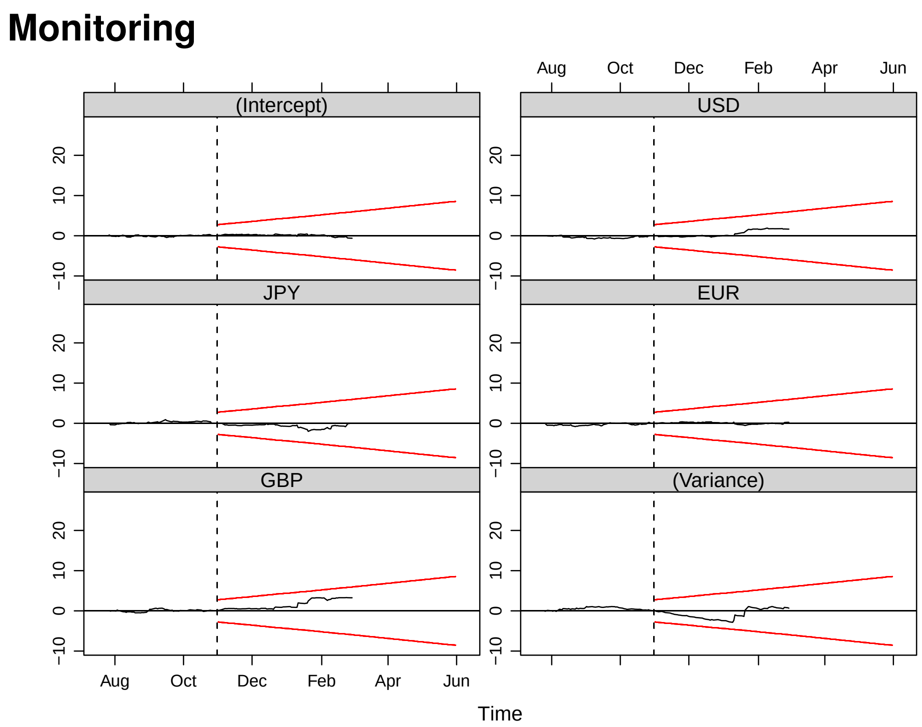
Monitoring
(Intercept)
USD
JPY
EUR
GBP
(Variance)
20 10 0 −10
Aug Oct Dec Feb Apr Jun
Time

Monitoring

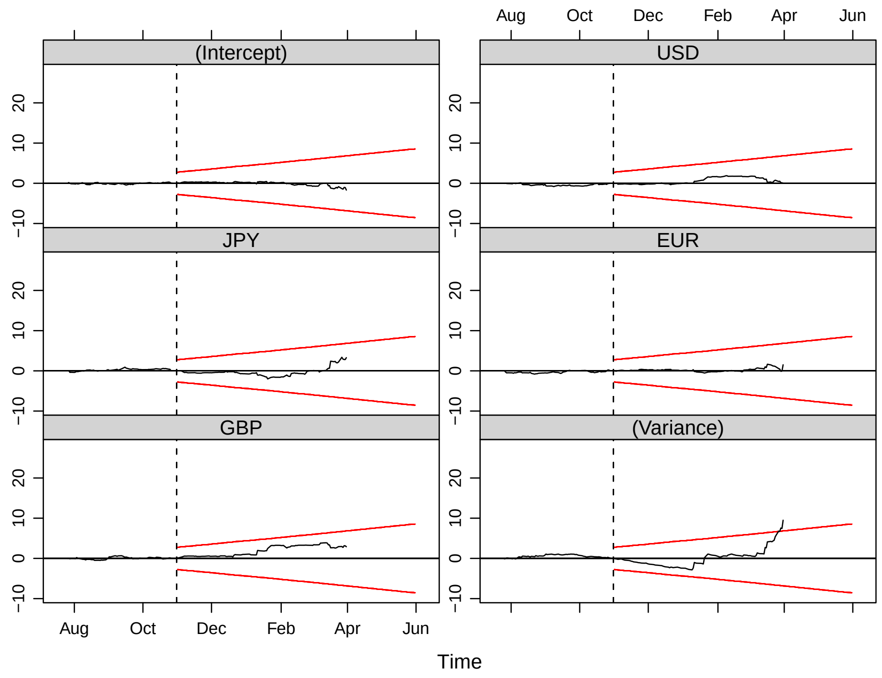

Monitoring

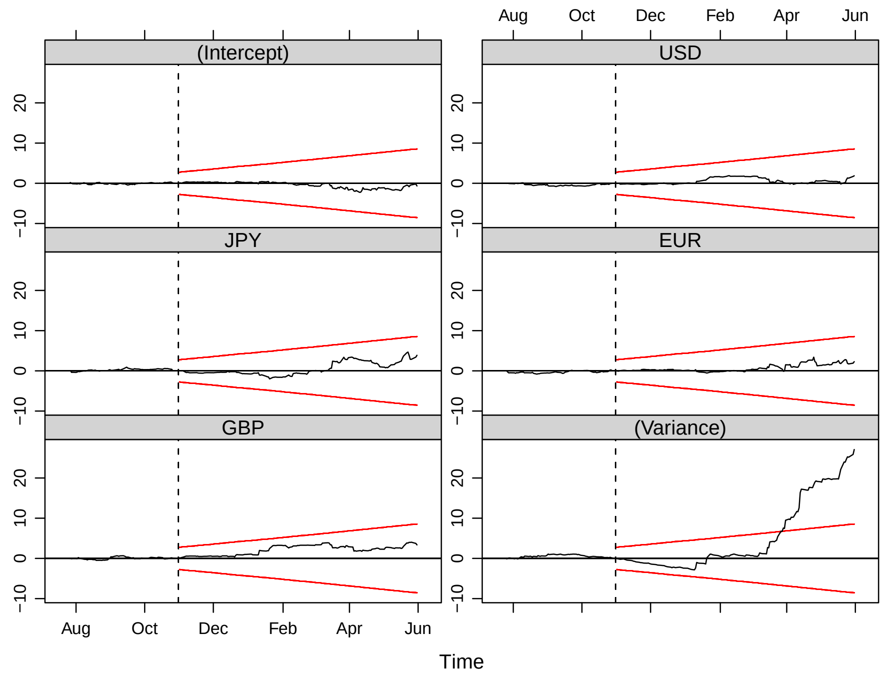

Monitoring

This signals a clear increase in the error variance.

The change is picked up by the monitoring procedure on 2006-03-27.

The other regression coefficients did not change significantly, signalling that they are not part of the basket peg.

Using data from an extended period up to 2009-07-31, we fit a segmented model to determine where and how the model parameters changed.

Dating

Segmented regression model: A stable model with parameter vector $\theta^{(j)}$ holds for the observations in $i = i_{j-1} + 1, \dots, i_j$. The segment index is $j = 1, \dots, m + 1$.

The set of m breakpoints $\mathcal{I}_{m,n} = \{i_1, \dots, i_m\}$ is called m-partition. Convention: $i_0 = 0$ and $i_{m+1} = n$.

The value of the segmented objective function Ψ is

$$\begin{aligned} PSI(i_1, \dots, i_m) &= \sum_{j=1}^{m+1} psi(i_{j-1} + 1, i_j), \\ psi(i_{j-1} + 1, i_j) &= \sum_{i=i_{j-1}+1}^{i_j} \Psi(y_i, x_i, \widehat{\theta}^{(j)}). \end{aligned}$$

Dating

Thus, $psi(i_{j-1}+1, i_j)$ is the minimal value of the objective function for the model fitted on the jth segment.

Dating tries to find

$$(\hat{\imath}_1, \dots, \hat{\imath}_m) = \underset{(i_1,\dots,i_m)}{\text{argmin}}\ PSI(i_1, \dots, i_m)$$

over all partitions $(i_1, \dots, i_m)$ with $i_j - i_{j-1} + 1 \geq \lfloor nh \rfloor \geq k$.

Bellman principle of optimality:

$$PSI(\mathcal{I}_{m,n}) = \min_{mn_h \leq i \leq n - n_h} [PSI(\mathcal{I}_{m-1,i}) + psi(i+1, n)]$$

Dating

It is well-known that this problem can be solved by a dynamic programming algorithm of order $O(n^2)$ that essentially relies on a triangular matrix of $psi(i,j)$ for all $1 \leq i < j \leq n$.

In linear regressions this approach has been popularized by Bai & Perron and it is common practice to use the residual sum of squares as objective function:

$$\Psi_{\text{RSS}}(y_i, x_i, \beta) = (y_i - x_i^\top \beta)^2.$$

To capture changes in the variances as well the (negative) log-likelihood from a normal model can be employed:

$$\Psi_{\text{NLL}}(y_i, x_i, \beta, \sigma) = -\log\left(\sigma^{-1}\phi\left(\frac{y_i - x_i^\top \beta}{\sigma}\right)\right).$$

Dating

Thus, for a given number of breaks m, the optimal breaks $\hat{\imath}_1, \dots, \hat{\imath}_m$ be found.

To determine the number of breaks, some model selection has to be done, e.g., via information criteria or sequential tests. Here, we use the LWZ criterion (modified BIC):

$$\begin{aligned} IC(m) &= 2 \cdot NLL(\mathcal{I}_{m,n}) + \text{pen} \cdot ((m+1)k + m), \\ \text{pen}_{\text{BIC}} &= \log(n), \\ \text{pen}_{\text{LWZ}} &= 0.299 \cdot \log(n)^{2.1}. \end{aligned}$$

Dating

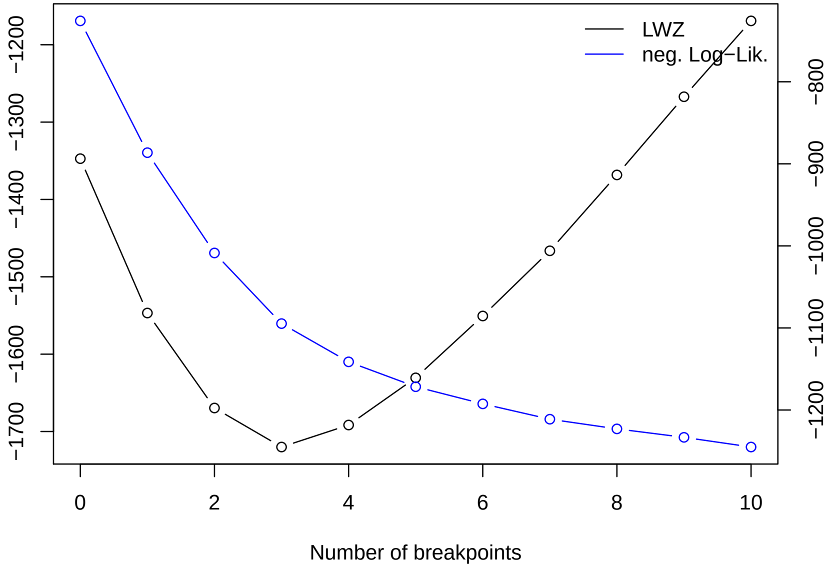

Dating

The estimated breakpoints and parameters are:

start/end	β_0	β_{USD}	β_{JPY}	β_{EUR}	β_{GBP}	σ	R^2
2005-07-26	**−0.005**	**0.999**	0.005	−0.015	0.007	0.028	0.998
2006-03-14	(0.002)	(0.005)	(0.005)	(0.017)	(0.008)		
2006-03-15	**−0.025**	**0.969**	−0.009	0.026	−0.013	0.106	0.965
2008-08-22	(0.004)	(0.012)	(0.010)	(0.023)	(0.012)		
2008-08-25	−0.015	**1.031**	−0.026	0.049	0.007	0.263	0.956
2008-12-31	(0.030)	(0.044)	(0.030)	(0.059)	(0.035)		
2009-01-02	0.001	**0.981**	0.008	−0.008	0.009	0.044	0.998
2009-07-31	(0.004)	(0.005)	(0.004)	(0.009)	(0.004)		

corresponding to

1. tight USD peg with slight appreciation,
2. slightly relaxed USD peg with some more appreciation,
3. slightly relaxed USD peg without appreciation,
4. tight USD peg without appreciation.

Dating

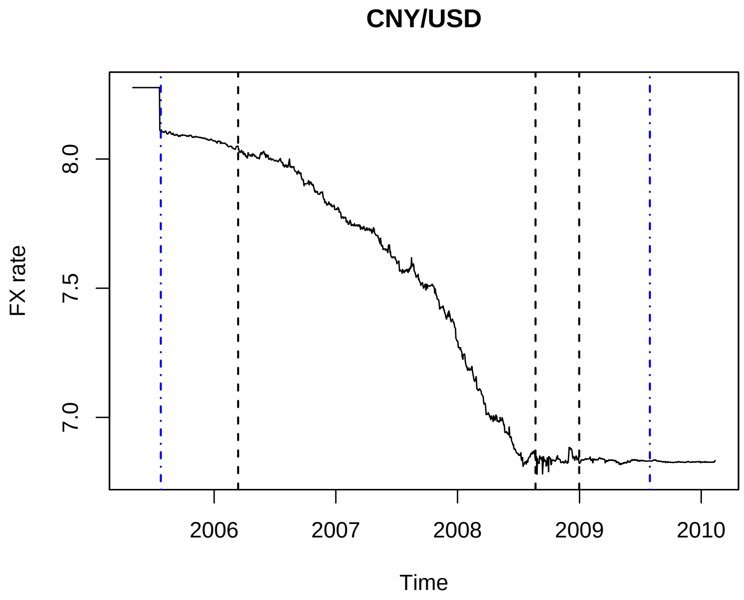

Dating

Epilogue: What happened since summer 2009?

Estimation based on 2009-08-04 through 2010-01-29 ($n = 122$) gives:

$$\begin{aligned} \text{CNY}_i &= \underset{(0.002)}{0.001} + \underset{(0.003)}{0.9953}\,\text{USD}_i + \underset{(0.003)}{0.002}\,\text{JPY}_i \\ &\quad + \underset{(0.011)}{0.007}\,\text{EUR}_i + \underset{(0.003)}{0.004}\,\text{GBP}_i + \hat{\varepsilon}_i \end{aligned}$$

Only the USD coefficient is significantly different from 0 (but not from 1).

The error standard deviation became even smaller with $\hat{\sigma} = 0.018$ leading to $R^2 = 0.999$.

Software

All methods are implemented in the R system for statistical computing and graphics and are freely available in the contributed packages *strucchange* and *fxregime* from the Comprehensive R Archive Network:

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

Software: *strucchange*

Classical structural change tools for OLS regression:

- Testing: `efp()`, `Fstats()`, `sctest()`.
- Monitoring: `mefp()`, `monitor()`.
- Dating: `breakpoints()`.
- Vignette: `"strucchange-intro"`.

Object-oriented structural change tools:

- Testing: `gefp()`, `efpFunctional()` (including special cases: `maxBB`, `meanL2BB`, `supLM`, `...`).
- Monitoring: Object-oriented implementation still to do.
- Dating: Some currently unexported support in `gbreakpoints()` in *fxregime*.
- Vignette: None, but CSDA paper.

Software: *fxregime*

Structural change tools for exchange rate regression based on normal (quasi-)ML:

- Data: `FXRatesCHF` (“`zoo`” series with US Federal Reserve exchange rates in CHF for various currencies).
- Preprocessing: `fxreturns()`.
- Model fitting: `fxlm()`.
- Testing: `gefp()` from *strucchange*.
- Monitoring: `fxmonitor()`.
- Dating: `fxregimes()` based on currently unexported `gbreakpoints()`; `refit()` method for fitting segmented regression.
- Vignettes: `"CNY"`, `"INR"`.

Application: Indian FX regimes

India also has an expanding economy with a currency receiving increased interest over the last years. We track the evolution of the INR FX regime since trading in the INR began.

```
R> head(FXRatesCHF[, c(1:6, 13)], 3)

             USD  JPY   DUR EUR   DEM    GBP INR
1971-01-04 0.232 82.8 0.429  NA 0.844 0.0967  NA
1971-01-05 0.232 83.0 0.429  NA 0.845 0.0968  NA
1971-01-06 0.232 83.0 0.429  NA 0.845 0.0968  NA

R> inr <- fxreturns("INR", data = FXRatesCHF,
+    other = c("USD", "JPY", "DUR", "GBP"), frequency = "weekly",
+    start = as.Date("1993-04-01"), end = as.Date("2008-01-04"))
R> head(inr, 3)

               INR     USD     JPY   DUR      GBP
1993-04-09  0.9773  0.9773  0.0977 0.567 -0.02236
1993-04-16 -0.0339 -0.0339 -0.5387 0.625  0.14295
1993-04-23  3.2339  3.2339  1.4331 1.264  0.00876
```

Application: Indian FX regimes

Using weekly returns from 1993-04-09 through to 2008-01-04 (yielding $n = 770$ observations), we fit a single FX regression using the same basket as above.

```
R> inr_lm <- fxlm(INR ~ USD + JPY + DUR + GBP, data = inr)
R> coef(inr_lm)

(Intercept)         USD         JPY         DUR         GBP
     0.0280      0.9185      0.0405      0.1046      0.0484
 (Variance)
     0.3375
```

Application: Indian FX regimes

As we would expect multiple changes, we assess its stability with the Nyblom-Hansen test. Alternatively, a MOSUM test could be used. The double maximum test has less power.

```
R> inr_efp <- gefp(inr_lm, fit = NULL)
R> sctest(inr_efp, functional = meanL2BB)

M-fluctuation test

data:  inr_efp
f(efp) = 3.11, p-value = 0.005

R> sctest(inr_efp, functional = maxBB)

M-fluctuation test

data:  inr_efp
f(efp) = 1.72, p-value = 0.03099
```

Application: Indian FX regimes

```
R> plot(inr_efp, functional = meanL2BB)
```

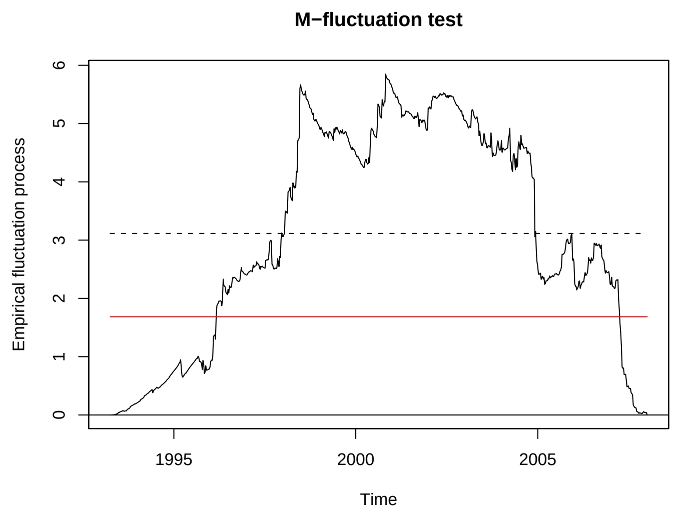

Application: Indian FX regimes

```
R> plot(inr_efp, functional = maxBB, aggregate = FALSE,
+    ylim = c(-2, 2))
```

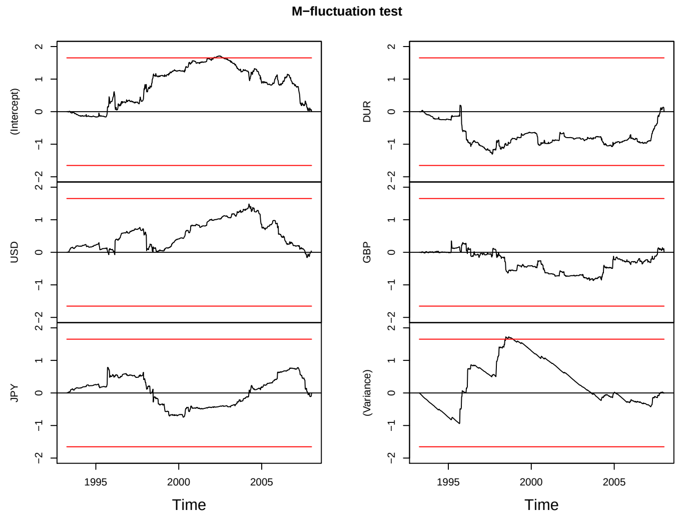

Application: Indian FX regimes

Dating is computationally more demanding. The dynamic programming algorithm can be parallelized, though. This is easily available (thanks to Anmol Sethy) by means of optional *foreach* support in *fxregime*.

```
R> library("foreach")
R> library("doMC")
R> registerDoMC(2)
R> inr_reg <- fxregimes(INR ~ USD + JPY + DUR + GBP, data = inr,
+    h = 20, breaks = 10, hpc = "foreach")
```

Application: Indian FX regimes

```
R> plot(inr_reg)
```

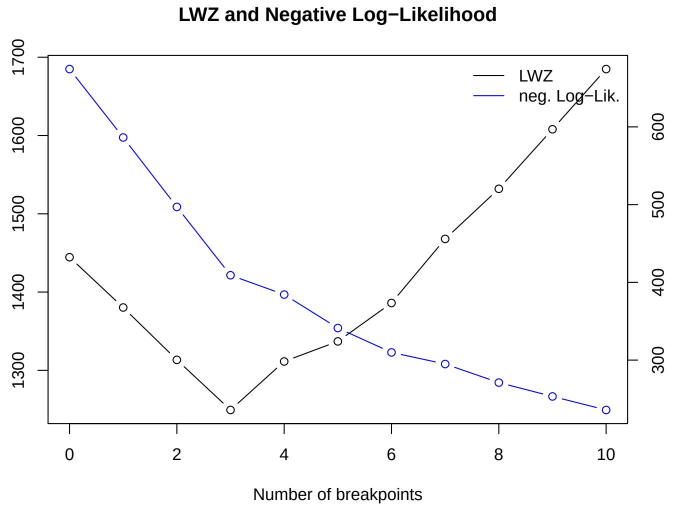

Application: Indian FX regimes

Various methods for extracting information can be applied directly. Otherwise, refitting of FX regressions gives access to all quantities that might be of interest.

```
R> coef(inr_reg)[, 1:5]
                        (Intercept)   USD     JPY    DUR      GBP
1993-04-09--1995-03-03    -0.00574 0.972 0.02347  0.0113  0.02037
1995-03-10--1998-08-21     0.16113 0.943 0.06692 -0.0261  0.04236
1998-08-28--2004-03-19     0.01861 0.993 0.00976  0.0983 -0.00322
2004-03-26--2008-01-04    -0.05761 0.746 0.12561  0.4354  0.12137
R> inr_rf <- refit(inr_reg)
R> sapply(inr_rf, function(x) summary(x)$r.squared)
1993-04-09--1995-03-03 1995-03-10--1998-08-21 1998-08-28--2004-03-19
                 0.989                  0.729                  0.969
2004-03-26--2008-01-04
                 0.800
```

Application: Indian FX regimes

Somewhat more compactly:

start/end	β_0	β_{USD}	β_{JPY}	β_{DUR}	β_{GBP}	σ	R^2
1993-04-09	−0.006	**0.972**	0.023	0.011	0.020	0.157	0.989
1995-03-03	(0.017)	(0.018)	(0.014)	(0.032)	(0.024)		
1995-03-10	**0.161**	**0.943**	0.067	−0.026	0.042	0.924	0.729
1998-08-21	(0.071)	(0.074)	(0.048)	(0.155)	(0.080)		
1998-08-28	0.019	**0.993**	0.010	**0.098**	−0.003	0.275	0.969
2004-03-19	(0.016)	(0.016)	(0.010)	(0.034)	(0.021)		
2004-03-26	−0.058	**0.746**	**0.126**	**0.435**	**0.121**	0.579	0.800
2008-01-04	(0.042)	(0.045)	(0.042)	(0.116)	(0.056)		

corresponding to

1. tight USD peg,
2. flexible USD peg,
3. tight USD peg,
4. flexible basket peg.

Next steps

Current activities: Application to wider range of currencies.

Of particular interest: Classification of exchange rate regimes and monitoring.

Open problems:

- Fully automatic selection of breakpoints.
- Sequential usage of BIC/LWZ, i.e., with growing sample size n.
- Differences between subsequent regimes that are statistically significant but not practically relevant (or vice versa).

First steps: Anmol Sethy started to build infrastructure for larger FX rates database from mixed sources.

First results: World map of R^2 from FX regressions (basket: USD, EUR, JPY, GBP), November 2009, based on segmented weekly data.

Next steps

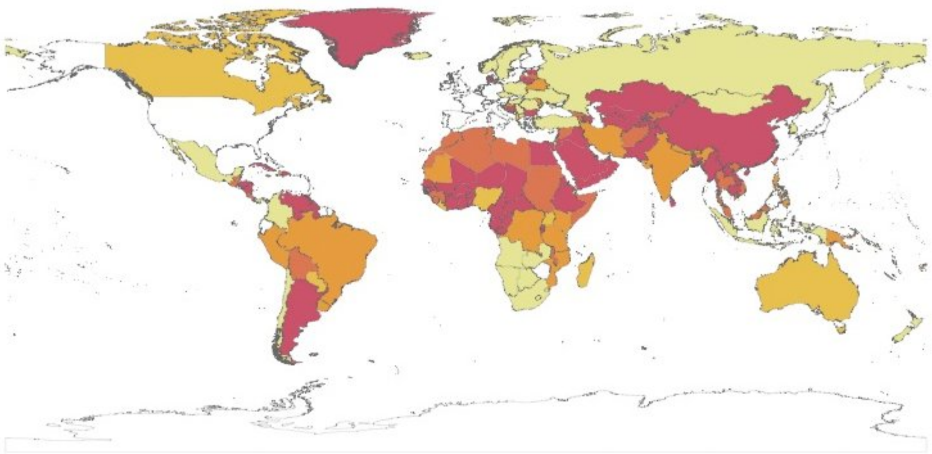

References

Zeileis A, Shah A, Patnaik I (2010). "Testing, Monitoring, and Dating Structural Changes in Exchange Rate Regimes." *Computational Statistics & Data Analysis*, **54**(6), 1696–1706. `doi:10.1016/j.csda.2009.12.005`.

Zeileis A (2005). "A Unified Approach to Structural Change Tests Based on ML Scores, *F* Statistics, and OLS Residuals." *Econometric Reviews*, **24**(4), 445–466. `doi:10.1080/07474930500406053`

Zeileis A, Leisch F, Kleiber C, Hornik K (2005). "Monitoring Structural Change in Dynamic Econometric Models." *Journal of Applied Econometrics*, **20**(1), 99–121. `doi:10.1002/jae.776`

Zeileis A, Kleiber C, Krämer W, Hornik K, (2003). "Testing and Dating of Structural Changes in Practice." *Computational Statistics & Data Analysis*, **44**(1–2), 109–123. `doi:10.1016/S0167-9473(03)00030-6`

Zeileis A, Leisch F, Hornik K, Kleiber C (2002). "*strucchange*: An R Package for Testing for Structural Change in Linear Regression Models." *Journal of Statistical Software*, **7**(2), 1–38. URL `http://www.jstatsoft.org/v07/i02/`.

Summary

- Exchange rate regime analysis can be complemented by structural change tools.
- Both coefficients (currency weights) and error variance (fluctuation band) can be assessed using an (approximately) normal regression model.
- Estimation, testing, monitoring, and dating are all based on the same model, i.e., the same objective function.
- Traditional significance tests can be complemented by graphical methods conveying timing and component affected by a structural change.
- Software is freely available, both for the general method and the application to FX regimes.