Regime-switching models for electricity spot prices

An empirical comparison

Rafał Weron

(joint work with **Joanna Janczura**)

Institute of Organization and Management
Wrocław University of Technology

Motivation

- Risk management and derivatives pricing often require a model for spot electricity prices that is:

- Realistic
  - Why would we want an unrealistic model?!
- Parsimonious
  - Faster simulation, smaller calibration errors
- Statistically sound
  - We can calibrate any model to any dataset …
  - … but does it **really** fit the data? Does it **make sense**?

Agenda

- Reduced form models for the spot price
  - Jump-diffusion (JD) models
  - First generation Markov Regime-Switching (MRS) models
  - Second generation MRS models
- Dealing with seasonality
- Empirical study

Reduced form models for the spot price

- Typically the deseasonalized spot electricity price X_t is assumed to follow some kind of a jump-diffusion (JD) process:

$$dX = \mu(X,t)dt + \sigma(X,t)dB + dq(X,t)$$

Generalized drift (possibly mean reverting)

Generalized volatility (possibly heteroskedastic)

Pure jump process with given intensity and severity – e.g. a compound Poisson process: $dq(X,t) = XdN_t$

- Clewlow & Strickland, 2000; Eydeland & Geman, 2000; Kaminski, 1999

Problems with JD models

- After a jump the price is forced back to its normal level
  - by mean reversion (**MRJD**)
  - by mean reversion coupled with downward jumps
    - Deng 1999; Escribano et al., 2002; Geman & Roncoroni, 2006
  - by a combination of mean reversions with different rates
    - Benth et al., 2007

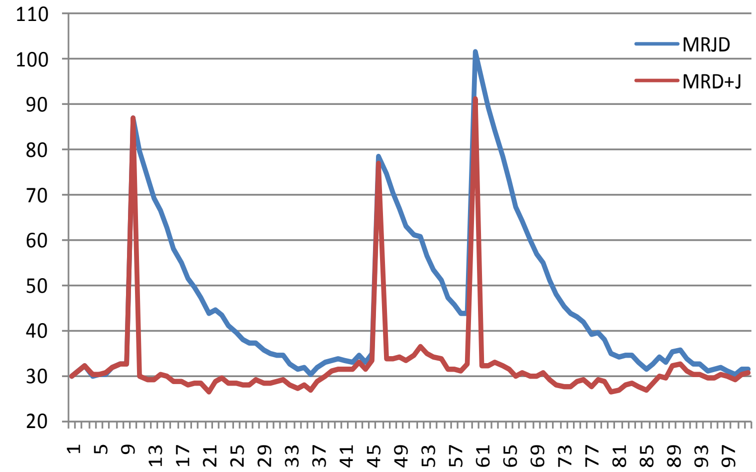

- Alternatively, a positive jump may be always followed by a negative jump of approximately the same size – especially on the daily scale (**MRD+J**)
    - Weron et al., 2004; Weron, 2008

Problems with JD models cont.

- What about periods of consecutive jumps?
  - Grid congestion, outage
- Solution:
  - Allow the process to 'stay' in the 'jump regime' with some probability
  - **Regime-switching models**

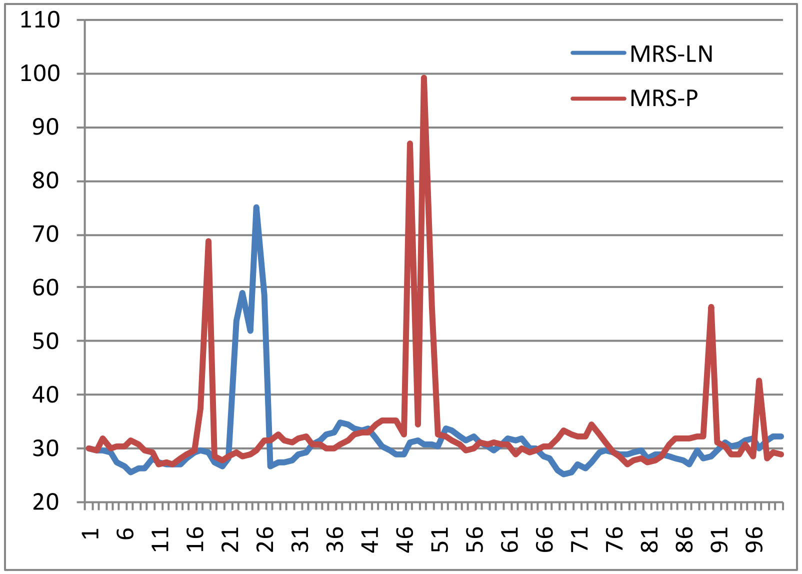

Markov regime switching (MRS) models

- Assume that the switching mechanism can be governed by a latent random variable that follows a Markov chain with two (or more) possible states

A two-state regime model: $X_t = \{1,2\}$

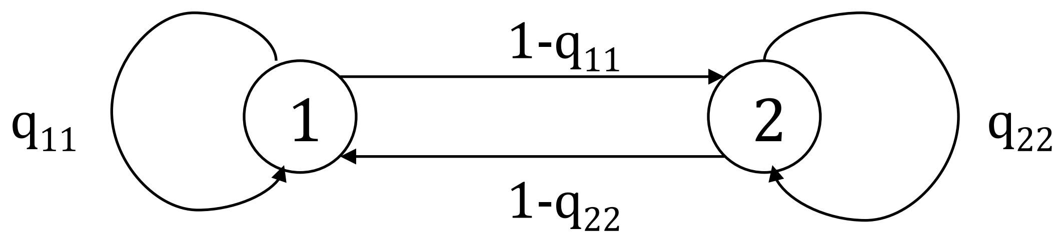

Transition probabilities:

$$\mathbf{Q} = (q_{ij}) = \begin{pmatrix} q_{11} & q_{12} \\ q_{21} & q_{22} \end{pmatrix} = \begin{pmatrix} q_{11} & 1 - q_{11} \\ 1 - q_{22} & q_{22} \end{pmatrix}$$

- The regimes are only latent, not directly observable
- Estimation via EM (expectation–maximization) algorithm

(Dempster et al., 1977; Hamilton, 1990; Kim, 1993)

First generation MRS models for electricity spot log-prices

- Ethier & Mount (1998) proposed a model with 2 regimes governed by AR(1) processes
  - Concluded that there was strong support for the existence of different means and variances in the two regimes
- Huisman & Mahieu (2003) proposed a model with 3 regimes in which
  - The initial jump regime was immediately followed by the reversing regime and then moved back to the base regime
- Huisman & de Jong (2003) proposed a 2-regime model
  - With a stable, mean-reverting AR(1) regime and
  - An independent spike (IS) regime modeled by a normal variable with a higher mean and variance

First generation MRS models for electricity spot log–prices cont.

- Bierbrauer et al. (2004), Weron et al. (2004) modified the IS 2-regime model
  - Used log-normal and Pareto distributed spike regimes
- De Jong (2006) modified the same IS 2-regime model
  - Introduced autoregressive, Poisson driven spike regime dynamics
- Kosater & Mosler (2006) modified both 2-regime models (Ethier & Mount, 1998; Huisman & de Jong, 2003)
  - Used a dummy variable to switch between two sets of parameters of the spike regime for business days and holidays
- Bierbrauer et al. (2007) proposed yet another modification of the IS 2-regime model
  - Used exponentially distributed spikes

Problems with first generation MRS models for log-prices

- Some authors reported that the ‘expected spike sizes’ ($\equiv \mathrm{E}(Y_{t,\text{spike}}) - \mathrm{E}(Y_{t,\text{base}})$) were **negative**
  - See e.g. De Jong (2006), Bierbrauer et al. (2007)
  - … but were not considered as evidence for model misspecification
- Regime classification was not checked but …
  - … the calibration scheme generally assigns all extreme prices to the spike regime
    - The ‘sudden drops’ in the log-price are not interesting for price modeling and derivatives valuation
    - They appear extreme only because of the log transform

Second generation MRS models for electricity (log-)spot prices

- Fundamental extensions to improve spike occurrence:
  - Mount et al. (2006) proposed a 2-regime model with
    - Two AR(1) regimes for log-prices and
    - Transition probabilities dependent on the reserve margin
  - Huisman (2008) extended the IS 2-regime model for log-prices
    - Considered temperature dependent transition probabilities
- Statistical refinements to improve goodness-of-fit:
  - Weron (2008) suggested to fit prices, not log-prices
  - Janczura and Weron (2009) extended the IS 2-regime model
    - Introduced CIR-type dynamics for the base regime and
    - Median-shifted spike regime distributions

Agenda

- Reduced form models for the spot price
- Dealing with seasonality
  - Price spikes
  - Short-term seasonality
  - Long-term seasonality
- Empirical study

Dealing with seasonality

- Spot price P_t (denoted also by S_t) is typically modeled as a sum (or a product) of
  - A 'deterministic' (seasonal) component Λ_t and
  - A purely stochastic component X_t
- A seasonal model for Λ_t is usually calibrated and removed from the data prior to estimating X_t
  - Extreme observations may impact the estimate of Λ_t
- Possible solutions
  - Use preprocessing – detect & replace price spikes with more 'normal' values; add them back before estimating X_t
  - Use methods 'immune' to outliers (quantiles, wavelets)

Short-term seasonality

- Differencing
  - Simplest form: $X_t = P_t - P_{t\text{-}7}$
- Removing the mean or median week

  Mo Tu We Th Fr Sa Su (week #1)

  Mo Tu We Th Fr Sa Su (week #2)

  … …

- Moving average (MA) method
  - Calculate $m_t = (P_{t\text{-}3} + \ldots + P_{t+3})/7$
  - Subtract the mean of deviations $(P_{k+7j} - m_{k+7j})$ from P_t

Long-term seasonality

- Fitting piecewise constant functions (dummy variables) for each month
  - Bhanot (2000), Haldrup & Nielsen (2006), Knittel & Roberts (2005), Lucia & Schwartz (2002)
  - For each day of the week → corresponds to mean/median week or moving average method
  - Can be used for modeling holidays
- Fitting (a sum of) sinusoids with trend
  - Bierbrauer et al. (2007), Borovkova & Permana (2006), Cartea & Figueroa (2005), De Jong (2006), Geman & Roncoroni (2006), Lucia & Schwartz (2002), Pilipovic (1997), Weron (2006)

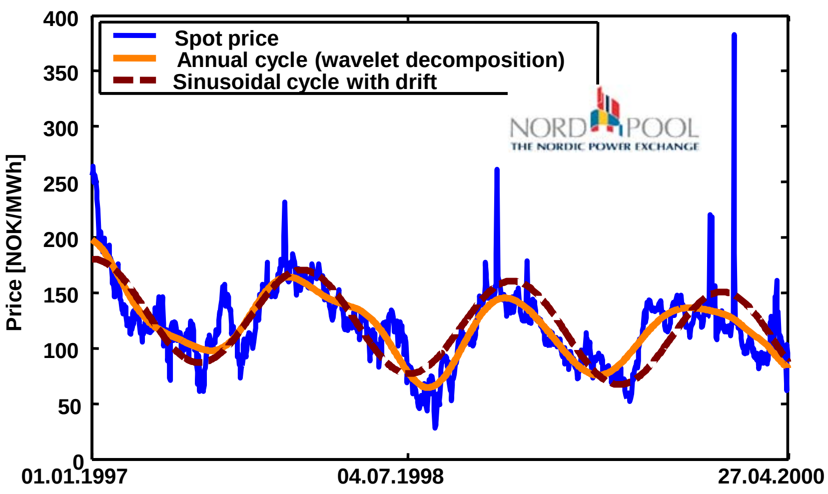

- Wavelet smoothing
  - Weron et al. (2004), Trück et al. (2007), Weron (2008), Janczura & Weron (2009)

Fitting sinusoids

- Sinusoid with linear trend
  - $\Lambda_t = A \cdot \sin(2\pi(t+B)) + C + D \cdot t$
  - t – time in years

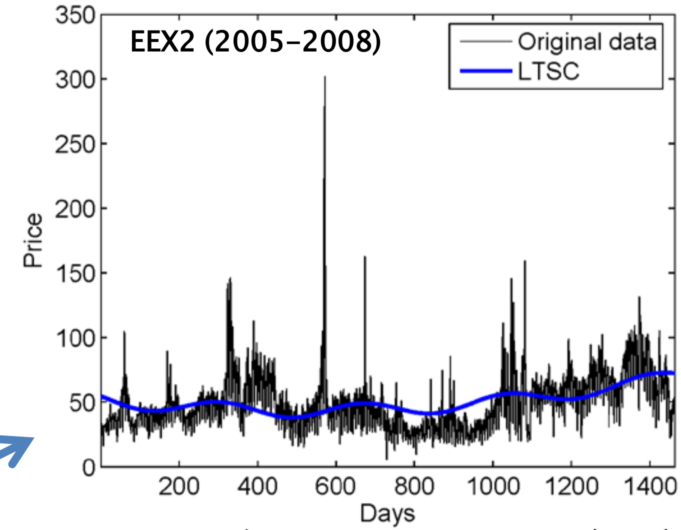

- … with cubic trend
  - $\Lambda_t = A \cdot \sin(2\pi(t+B)) + C + D \cdot t + E \cdot t^2$

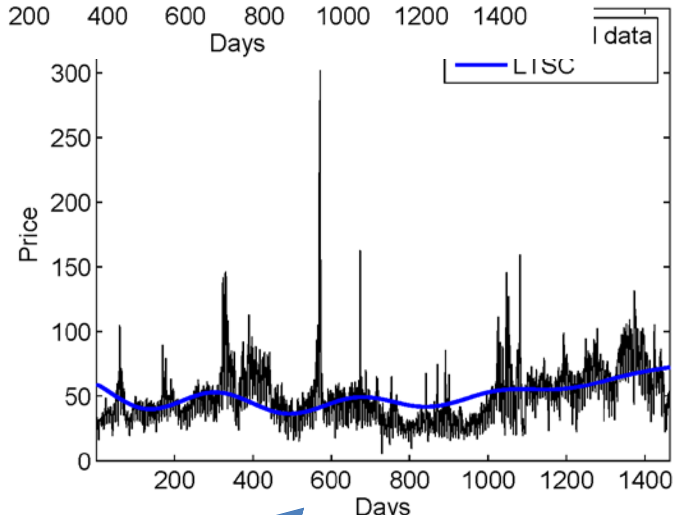

- … with linear trend and linear amplitude
  - $\Lambda_t = (A + F \cdot t) \cdot \sin(2\pi(t+B)) + C + D \cdot t$
- … with cubic trend and linear amplitude
  - $\Lambda_t = (A + F \cdot t) \cdot \sin(2\pi(t+B)) + C + D \cdot t + E \cdot t^2$

Wavelet smoothing 1/2

- Any signal (here: the spot price) can be built up as a sequence of projections onto
  - one father wavelet S_J and
  - a sequence of mother wavelets $\{D_j\}$

$$x(t) = S_J + D_J + D_{J-1} + \dots + D_1$$

  - where 2^J is the maximum scale sustainable by the number of observations

  - **Sines** are localized in frequency (characteristic scale)
  - **Wavelets** are also localized in time (space)

Wavelet smoothing 2/2

- At the coarsest scale the signal can be estimated by S_J
  - By adding a mother wavelet D_j of a lower scale $j = J-1, J-2, \ldots$, we obtain a better estimate of the original signal
    $\rightarrow$ **lowpass filtering**
- For daily data the S_3, S_5 and S_8 approximations roughly correspond to
  - weekly ($2^3 = 8$ days),
  - monthly ($2^5 = 32$ days) and
  - annual ($2^8 = 256$ days) smoothing

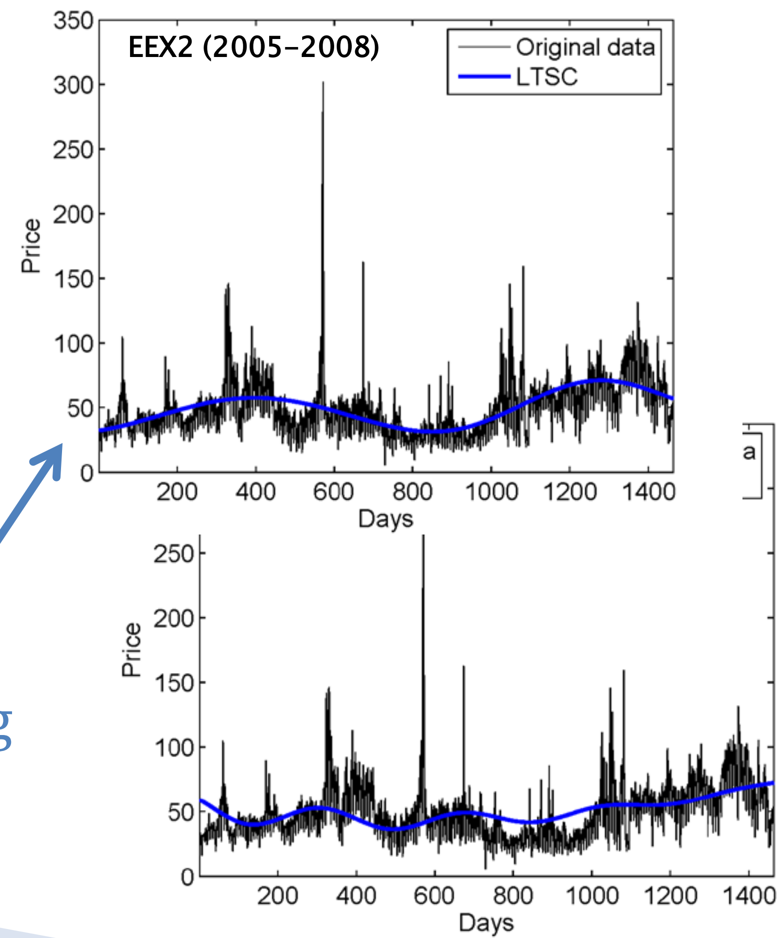

Forecasting long-term seasonality

- Piecewise constant functions
  - What about trends?
- Sinusoids
  - How to predict the trend?
  - Model risk
- Wavelets
  - Extrapolate the smoothed price?
- Forward prices
  - Smooth interpolation of forward prices
  - What about the far end?
  - What about risk premia?

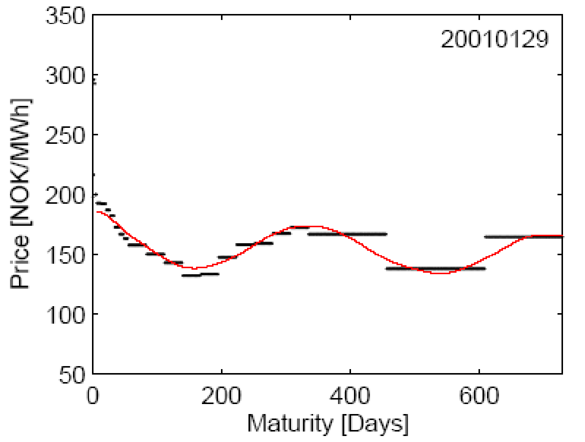

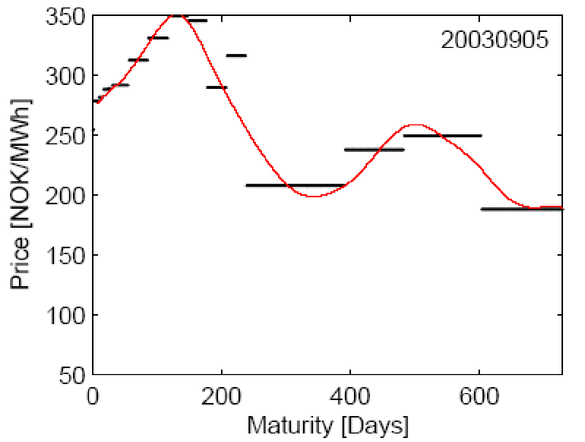

Agenda

- Reduced form models for the spot price
- Dealing with seasonality
- Empirical study
  - First generation MRS models
  - Shifted spike regime distributions
  - CIR base regime dynamics
  - Testing goodness-of-fit
  - 3-regime models

Data

- Mean daily (baseload) spot prices from EEX and PJM
  - EEX1: 1.1.2001-2.1.2005 EEX2: 3.1.2005-3.1.2009

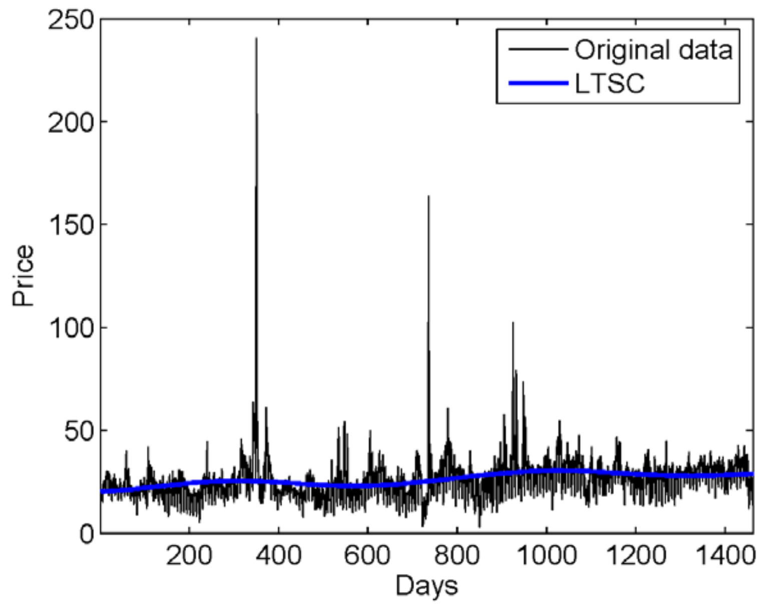

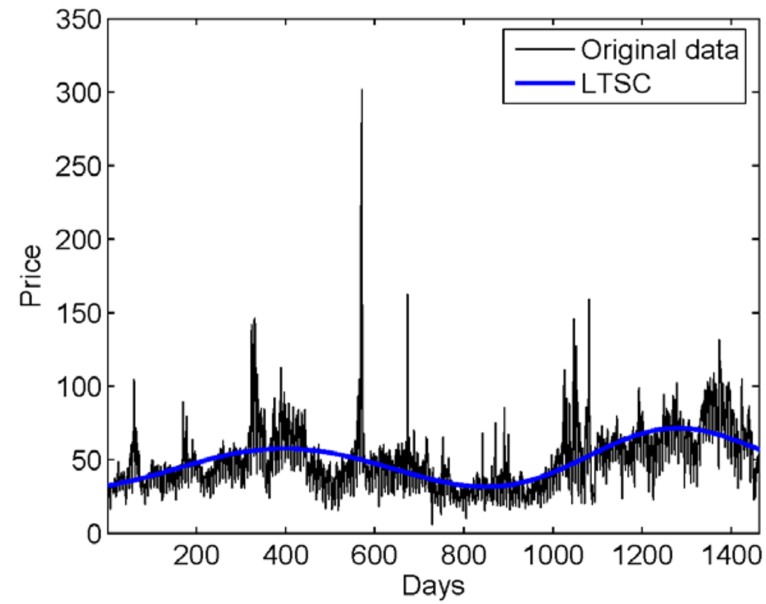

  - PJM: 1.1.2001-2.1.2005

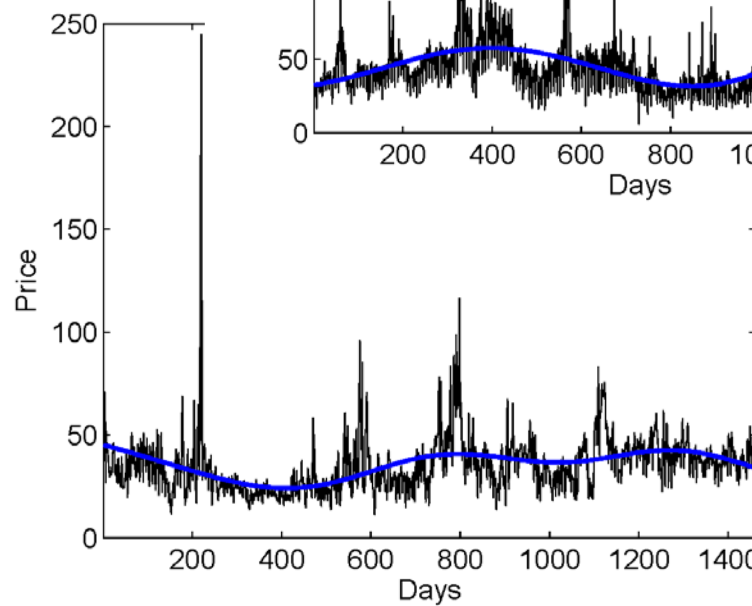

Dealing with seasonality

- Let the seasonal component Λ_t be composed of
  - **A long-term seasonal trend T_t**
    - Changing climate/consumption conditions throughout the year and the long-term non-periodic structural changes
  - And a weekly periodic part s_t

- First, T_t is estimated from daily spot prices P_t
  - Using 'annual' wavelet smoothing (J=8)
    - Fitting a sinusoid with cubic trend within a LS framework
- Next, s_t is removed by applying the MA method
- Finally, the deseasonalized prices, i.e. $P_t - T_t - s_t$, are shifted: min(new process) = min(P_t)

Loss distributions

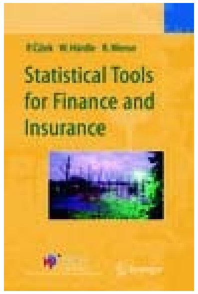

- Log-normal (LN), for $x > 0$

$$\log(X_{t,2}) \sim N(\mu, \sigma^2)$$

  - ML estimation straightforward
  - Apply Gaussian mean and std estimators to log-prices

Pareto (P), for $x \geq \lambda$

$$X_{t,2} \sim F_P(x; \alpha, \lambda) = 1 - \left(\frac{\lambda}{x}\right)^{\alpha}$$

The log-likelihood is increasing with scale parameter λ

Since $x \geq \lambda$, set $\lambda = \min(X_t)$

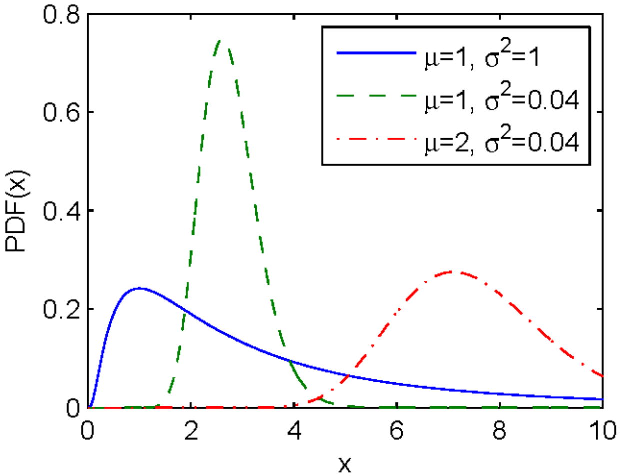

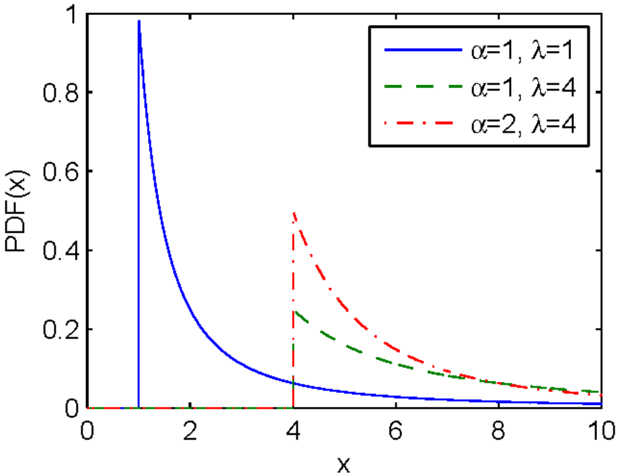

Vas–LN and Vas–P models fitted to (deseasonalized) log–prices

Prob. of staying in regime *i*

Prob. of being in regime *i*

E(X_{spike}) lower than E(X_{base}) ?!

Regime	Model	E(X_{t,i})	Var(X_{t,i})	q_ii	P(R=i)
		EEX (2001-2005)			
Base	Vas	3.5220	0.0235	0.9546	0.8844
Spike	LN	3.4745	0.3882	0.6530	0.1156
		PJM (2001-2005)			
Base	Vas	3.7049	0.0645	0.9679	0.9465
Spike	LN	3.5391	0.1937	0.4322	0.0536
		EEX (2001-2005)			
Base	Vas	3.5255	0.0315	0.9769	0.9663
Spike	P	Inf	Inf	0.3388	0.0337
		PJM (2001-2005)			
Base	Vas	3.6904	0.0618	0.9891	0.9797
Spike	P	3.7061	5.2432	0.4752	0.0204

Vas-LN and Vas-P models fitted to (deseasonalized) log-prices

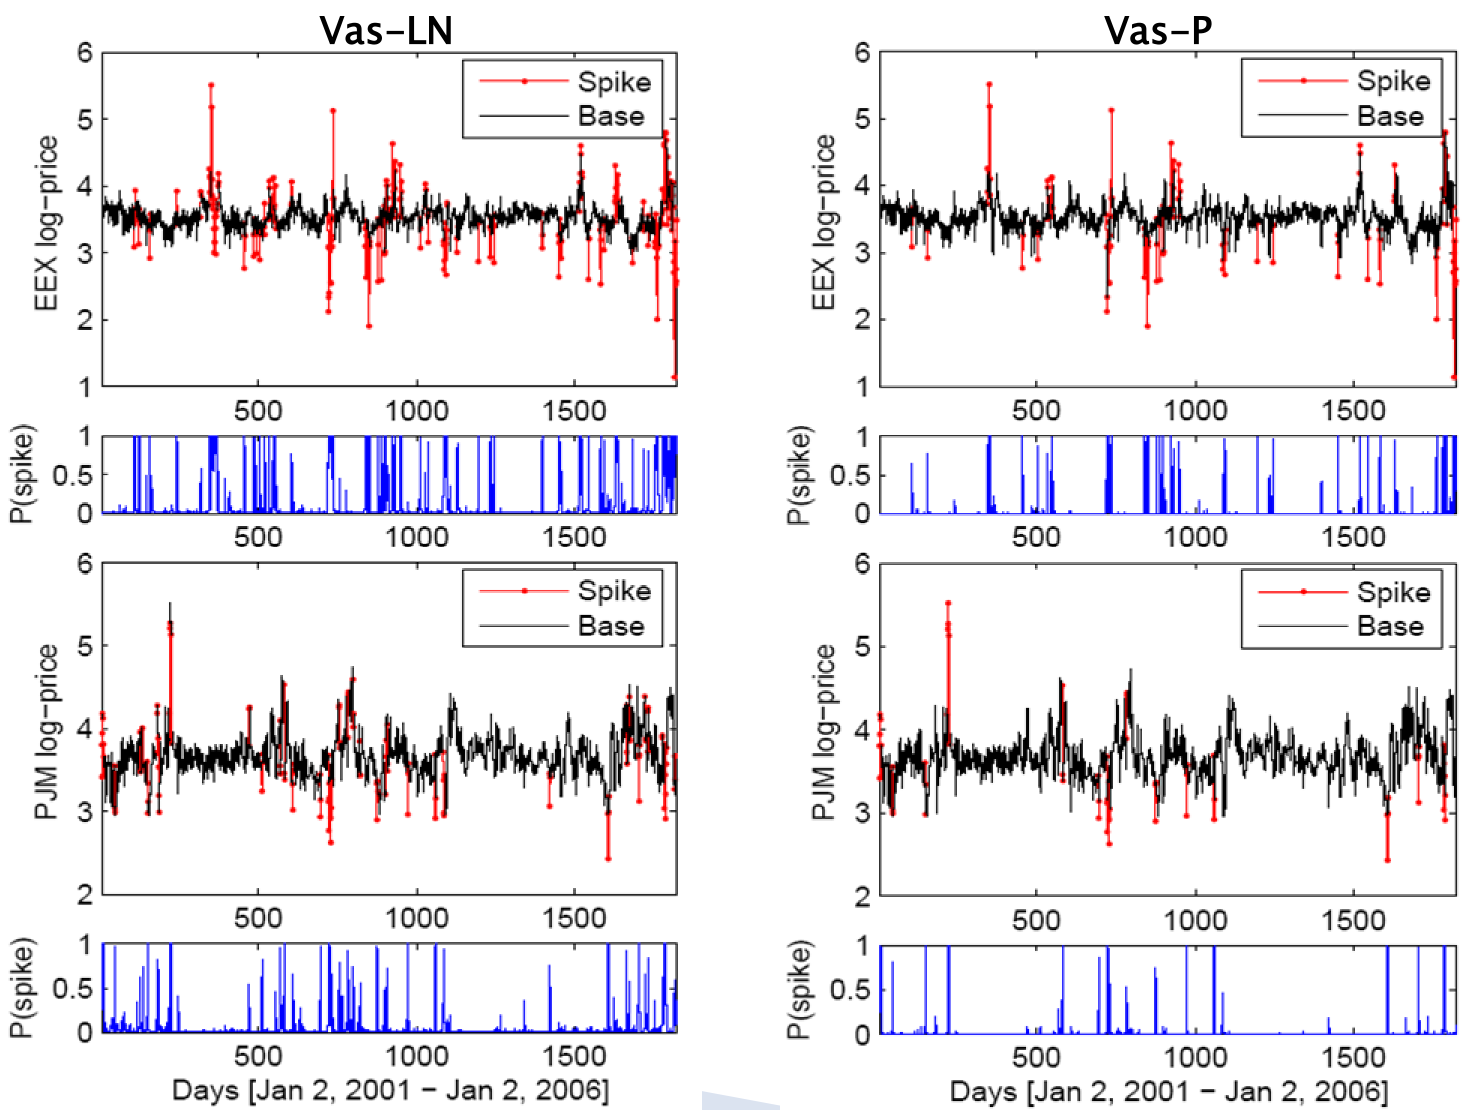

Vas–P model fitted to prices

Regime	E(X_{t,i})	Var(X_{t,i})	q_ii	P(R=i)
	EEX (2001-2005)			
Base	33.7241	30.3112	0.9714	0.9500
Spike	Inf	Inf	0.4563	0.0500
	PJM (2001-2005)			
Base	39.5883	87.6328	0.9802	0.9587
Spike	Inf	Inf	0.5400	0.0413

- Almost all spikes are identified correctly
- The number of ‘sudden drops’ classified as spikes is much lower
- The unconditional probabilities of being in the spike regime P(R = 2) are 2x higher
- **But … some low prices are still classified as spikes**

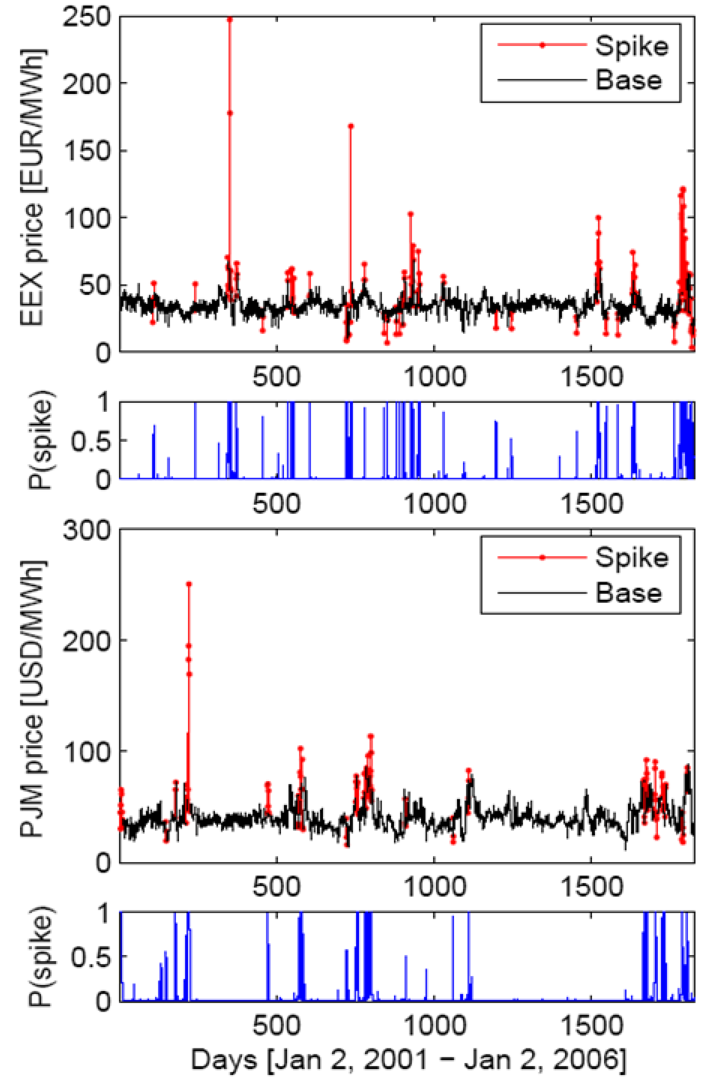

Shifted spike distributions

- Perhaps spike distributions should assign zero probability to prices below a certain quantile
- Let m = median(X_t)
  - Shifted log-normal (SLN), for $x > m$

$$\log\left(X_{t,2} - m\right) \sim N(\mu, \sigma^2)$$

  - Shifted Pareto (SP), for $x > \lambda \geq m$

$$X_{t,2} \sim F_P(x; \alpha, \lambda) = 1 - \left(\frac{\lambda}{x}\right)^{\alpha}$$

EEX1 (2001–2004) log-prices:
Comparison of spike regimes

$E(X_2)=3.33 < E(X_1)=3.42$!!!

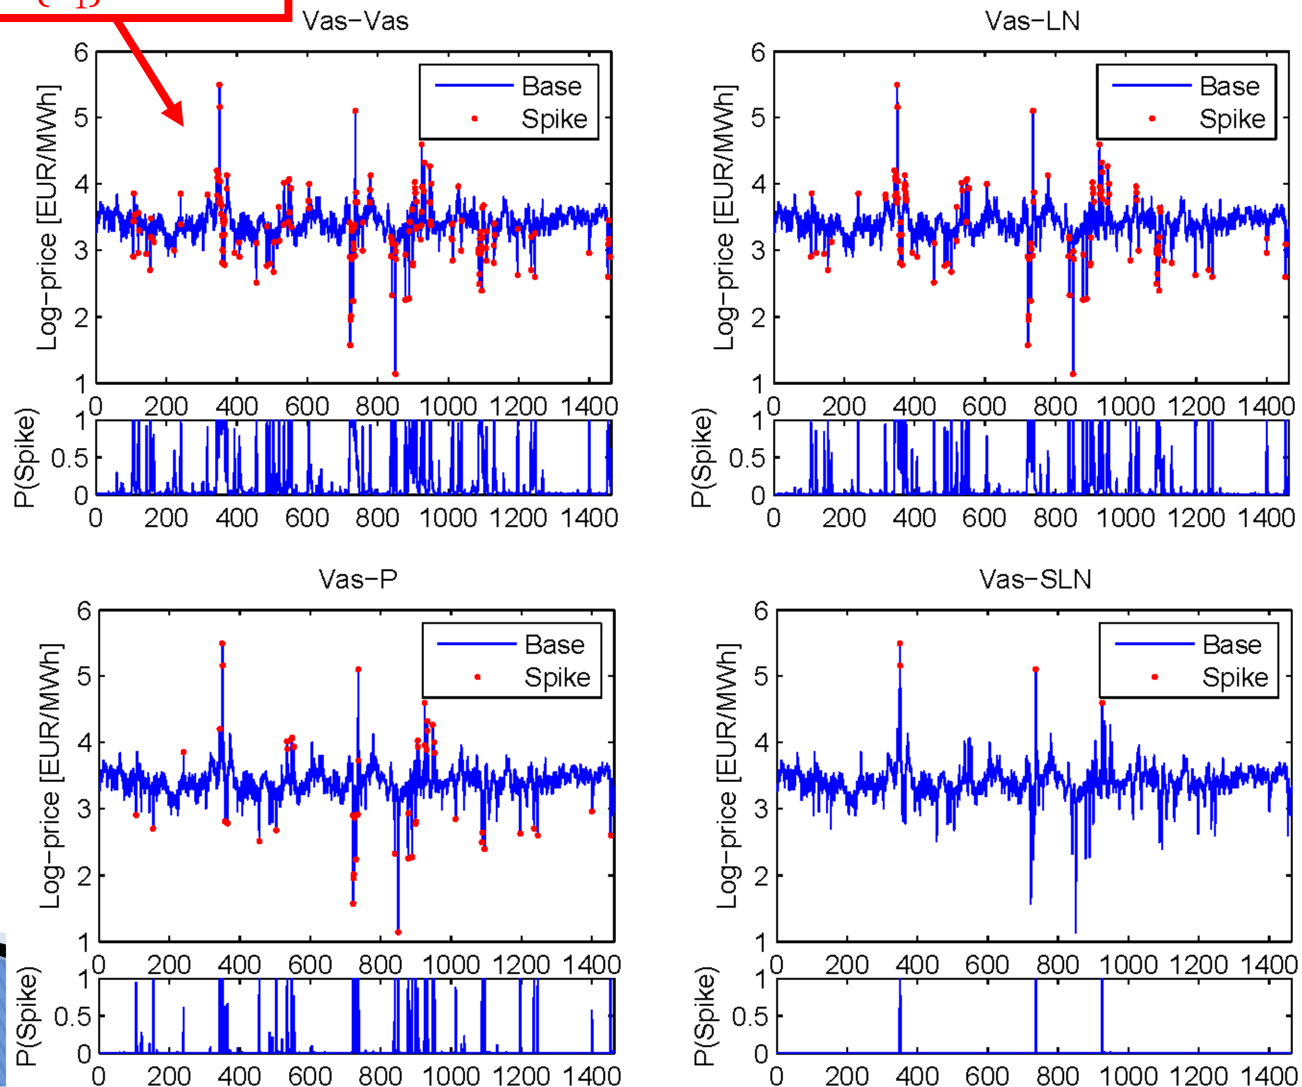

EEX1 (2001–2004) prices:
Comparison of spike regimes

$\mathrm{E}(X_2)=30.3 > \mathrm{E}(X_1)=41.6$

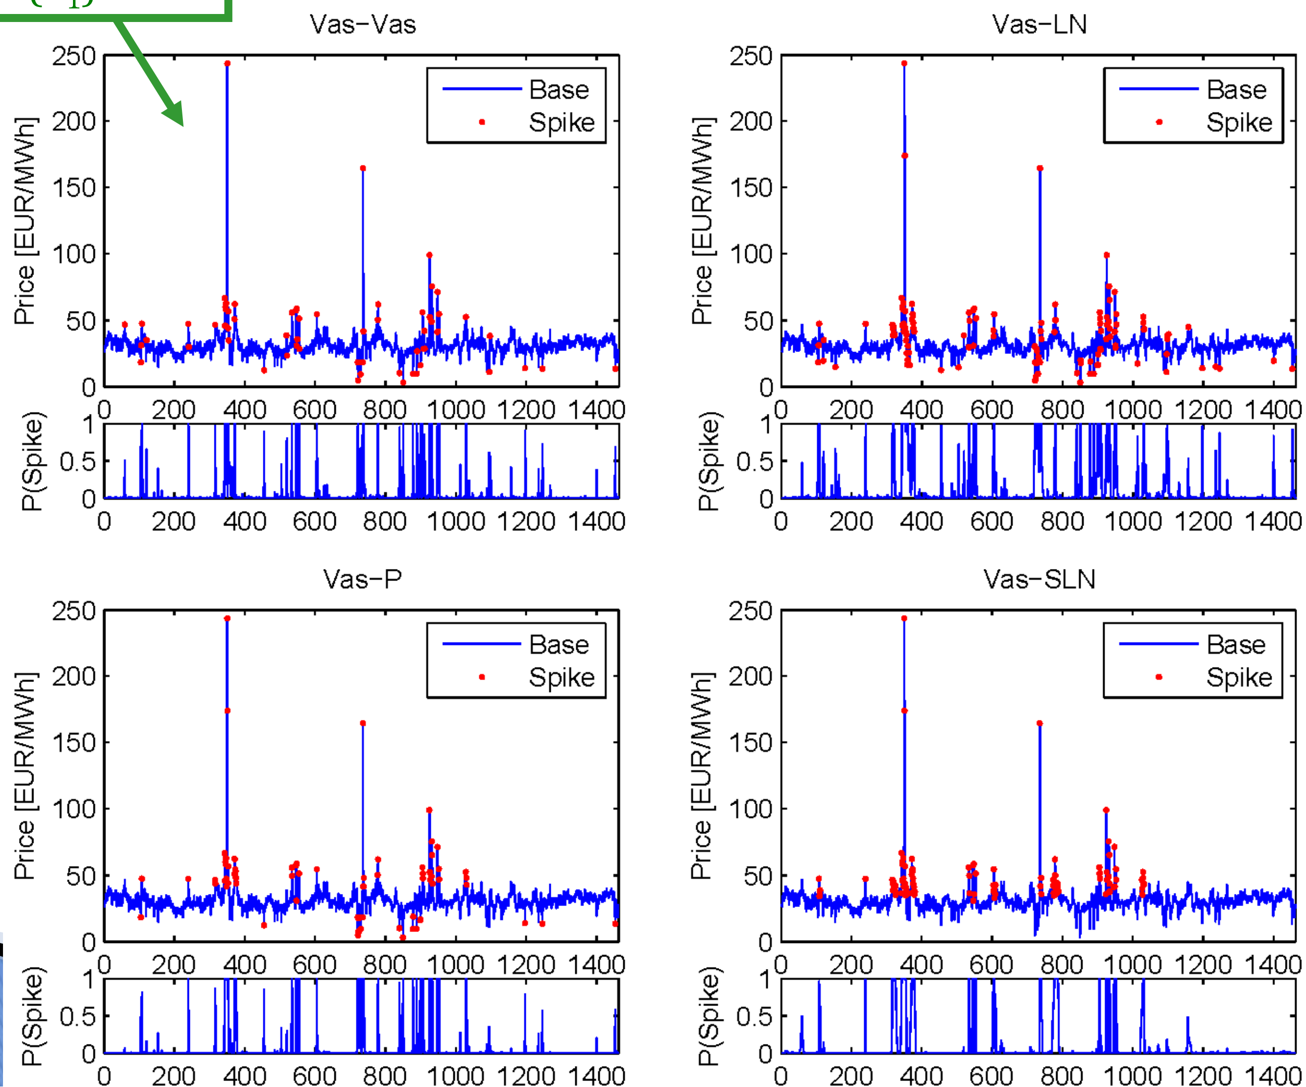

Conclusions

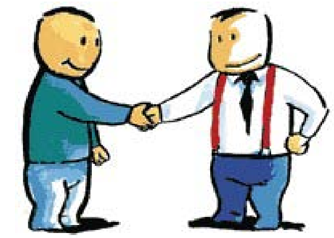

- When fitting MRS models with shifted spike distributions
  - Practically all spikes are identified correctly
  - There are no 'sudden drops' classified as spikes
- This suggests that the shifted distributions are more suitable for the spikes
  - But … there are clusters of 'normal' prices classified as spikes

Vas-SLN

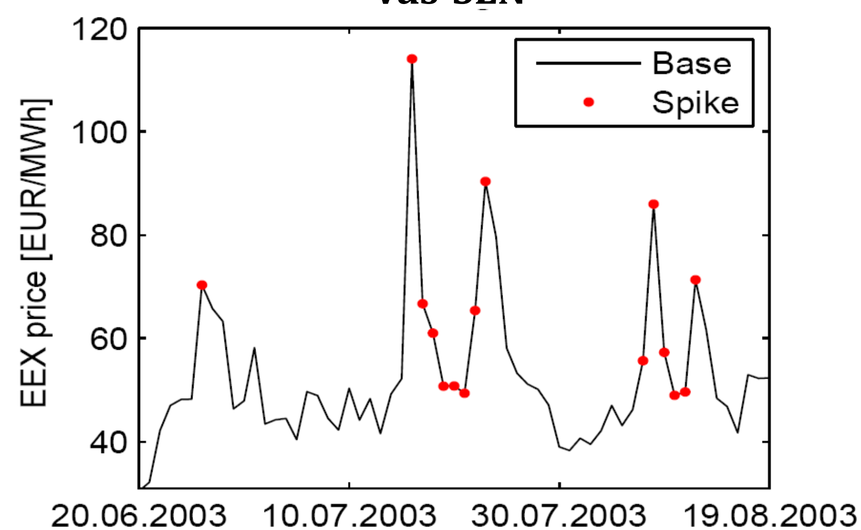

Diffusion processes revisited

- Perhaps different dynamics should be used for the base regime
- Introduce heteroskedasticity ($\gamma \neq 0$)

$$dX_t = (\alpha - \beta X_t)dt + \sigma X_t^{\gamma} dW_t$$

- For $\gamma=0$, Vasicek MRD process (Vasicek, 1977)
- For $\gamma=0.5$, CIR square root process (Cox, Ingersoll, Ross, 1985)
- For $\gamma=1$, BS process (Brennan, Schwartz, 1980)

EEX1 (2001–2004):
Vas vs. CIR base regime

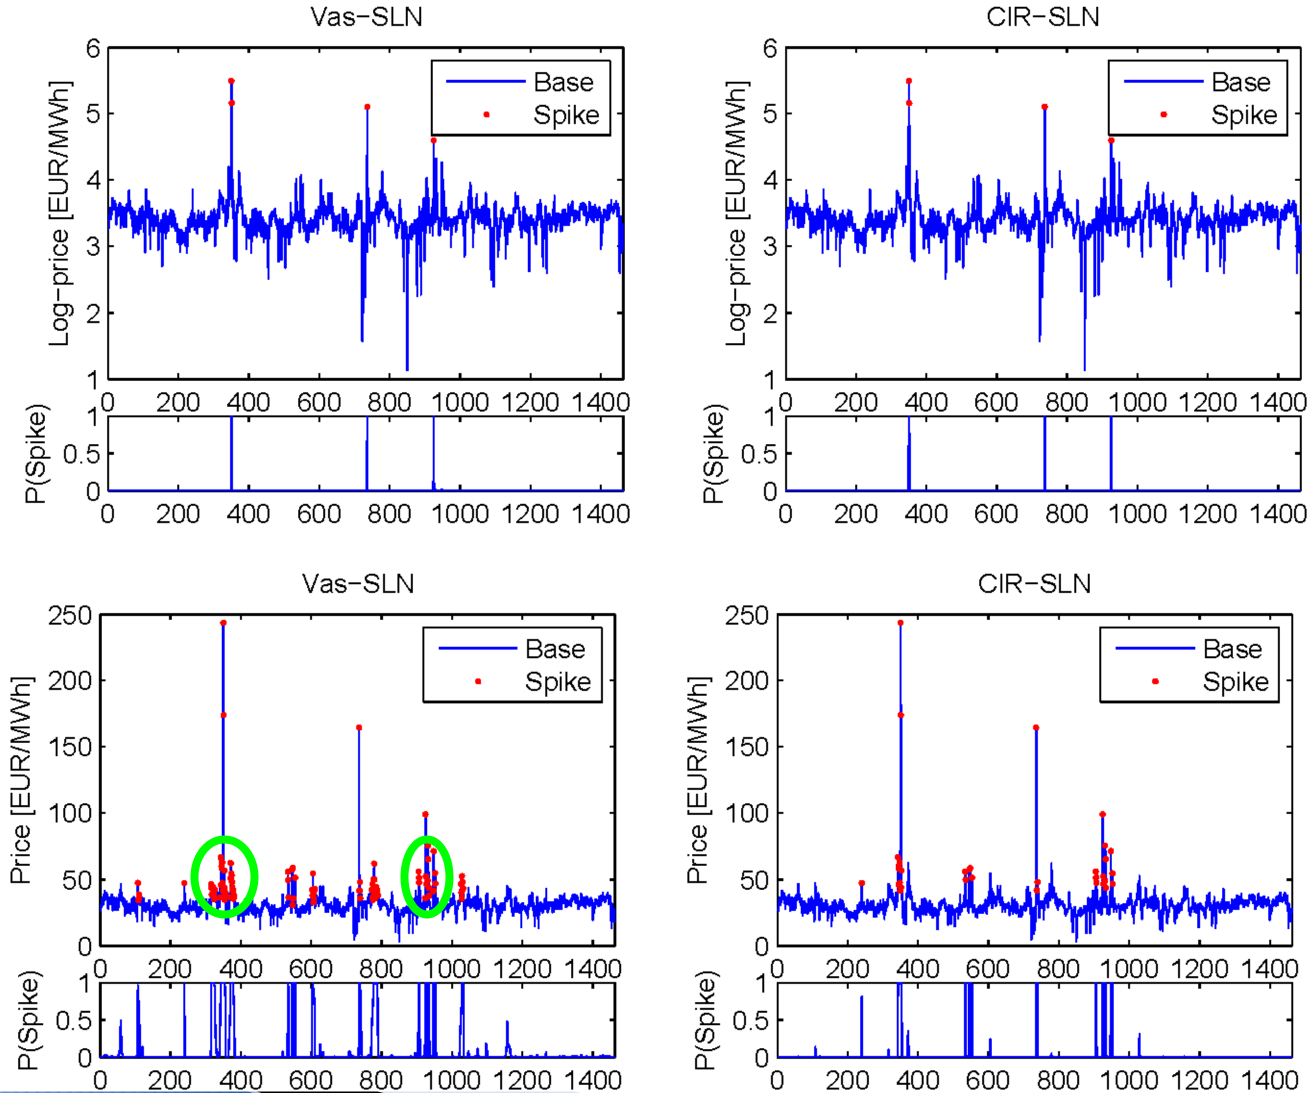

EEX2 (2005–2008):
Vas vs. CIR base regime

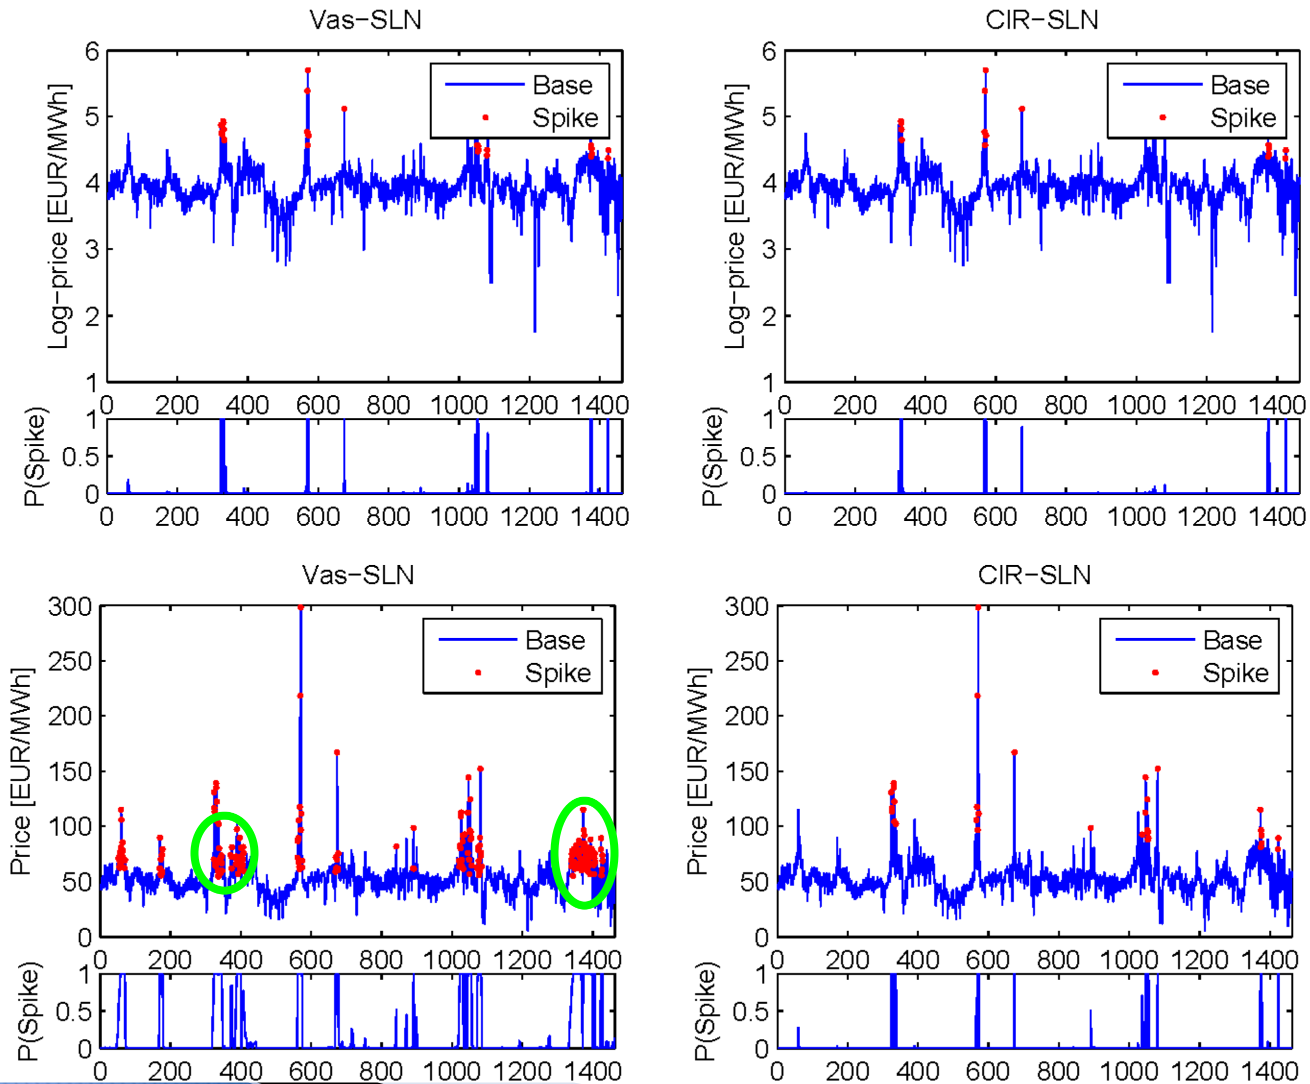

Goodness–of–fit

- Testing the goodness-of-fit for processes is not straightforward
- We can
  - Test the marginal distributions using an EDF-type test, like the Kolmogorov-Smirnov (K-S) test
- … but
  - The K-S test cannot be applied directly …
  - In the considered models neither the prices themselves nor their differences or returns are i.i.d.

Testing procedure

- Data is split into 2 subsets (3 for the 3-regime model)
  - **Spikes**, i.e. prices with probability P(R_t= 2) > 0.5
    - Price **drops**, i.e. prices with probability P(R_t= 3) > 0.5
  - The **base** regime, i.e. all remaining prices
- Discretization of the base regime dynamics:

$$dX_t = (\alpha - \beta X_t)dt + \sigma X_t^{\gamma} dW_t$$

  - Where γ = 0 for Vasicek and γ = 0.5 for CIR models
- … for *dt* = 1 leads to:

$$\varepsilon_t = (X_{t+1} - (1-\beta)X_t - \alpha)/\sigma X_t^{\gamma}$$

  - Where ε_t's are i.i.d. Gaussian random variables

Testing procedure cont.

- Applying the above transformation to base regime data we obtain 2 (or 3) i.i.d. samples:
  - f_S-distributed, e.g. lognormal or Pareto, for the spike regime
    - f_D-distributed, e.g. lognormal, for the drop regime
  - And Gaussian for the base regime

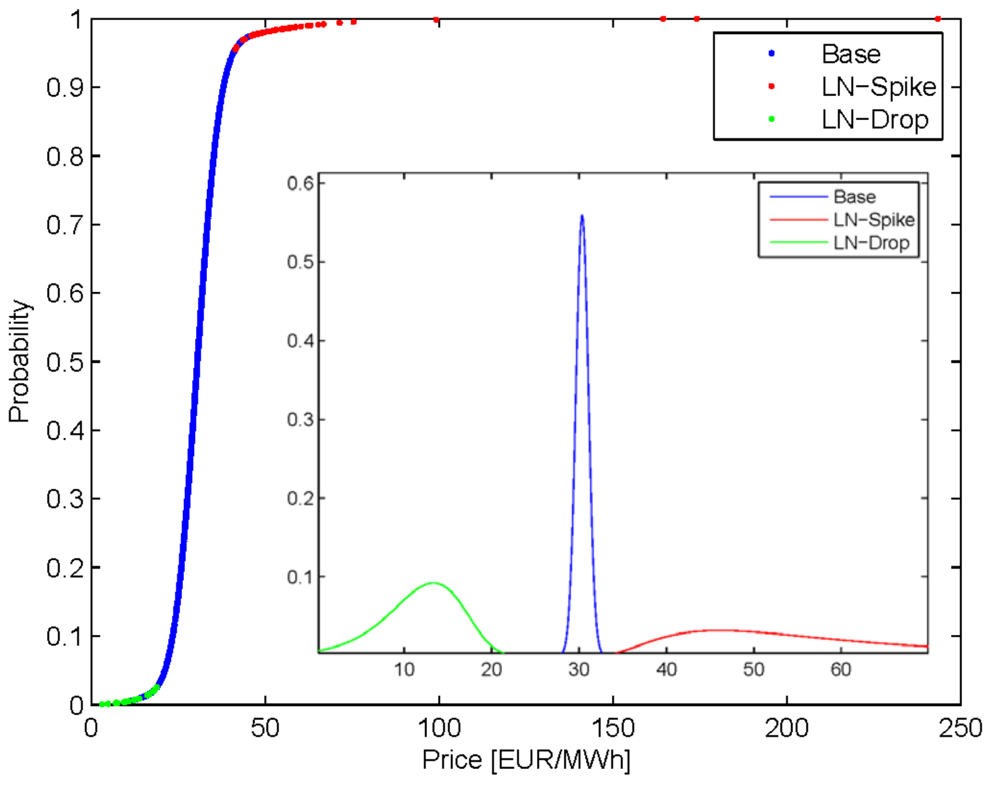

Testing procedure cont.

- Combining these samples yields
  - A sample of independent variables with the distribution being a mixture of 2 (or 3) laws:
    - f_S, (f_D,) and Gaussian
- The probability that a given price X_t comes from
  - The spike distribution is equal to P(R_t = 2)
    - The price drop distribution is equal to P(R_t = 3)
  - The Gaussian law is equal to P(R_t = 1)
- We perform the K-S test for
  - The subsets
  - And the whole sample

		Probabilities		Moments		Model statistics		KS test *p*-value	
Regime	Model	qii	P(Rt=i)	E(X(t,i))	Var(X(t,i))	+/- IDR	+/- IQR	Regime	Model
		EEX1 (2001-2004) prices							
Base	Vas	0.9841	0.9216	29.7309	23.5258	11.50%	8.75%	0.0012	0.0032
Spike	SLN	0.8128	0.0784	47.5885	288.2319			0.4061	
Base	CIR	0.9908	0.9756	30.2705	28.4517	13.90%	12.48%	0.0000	0.0000
Spike	SLN	0.6304	0.0244	63.3311	676.3723			0.3502	
		EEX1 (2001-2004) log-prices							
Base	Vas	0.9979	0.9972	3.4006	0.0453	29.52%	29.80%	0.0000	0.0000
Spike	SLN	0.2487	0.0028	5.0859	0.1260			0.7371	
Base	CIR	0.9979	0.9973	3.4006	0.0433	28.43%	27.97%	0.0000	0.0000
Spike	SLN	0.2498	0.0027	5.0892	0.1218			0.7506	
		PJM (2001-2004) prices							
Base	Vas	0.9872	0.9075	35.4050	44.2787	-0.36%	9.96%	0.0341	0.0530
Spike	SLN	0.8744	0.0925	63.9720	426.5396			0.4346	
Base	CIR	0.9940	0.9679	36.6248	66.3402	5.40%	20.69%	0.0052	0.0056
Spike	SLN	0.8196	0.0321	79.3603	833.1356			0.2280	
		PJM (2001-2004) log-prices							
Base	Vas	0.9937	0.9441	3.5618	0.0453	5.56%	18.02%	0.0007	0.0012
Spike	SLN	0.8937	0.0559	4.2129	0.0934			0.3219	
Base	CIR	0.9993	0.9973	3.5909	0.0599	9.82%	25.24%	0.0000	0.0000
Spike	SLN	0.7500	0.0027	5.2747	0.0209			0.8760	

Conclusions

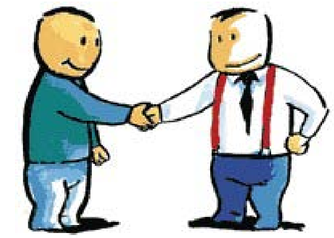

- When fitting MRS models with CIR base dynamics
  - Practically all spikes are identified correctly
  - There are no clusters of wrongly identified spikes
- Compared to Vasicek dynamics the CIR model yields
  - Fewer identified spikes
  - Lower probability of remaining in the spike regime
  - Slightly worse goodness-of-fit
    - Indicated by lower KS-test *p*-values
- None of the 2-regime models gives a satisfactory fit
  - Except for PJM prices
  - Generally this is caused by a bad fit of the base regime

3-regime models revisited

- Perhaps we need a 3rd regime to model the ‘price drops’
- Introduce a 3rd ‘drop’ regime
  - Contrary to the Huisman & Mahieu (2003) model, the price can stay in the ‘excited’ regimes (‘spike’ and ‘drop’)
  - Use a ‘mirror image’ or ‘reflected’ shifted log-normal distribution

EEX1 (2001–2004):

		Probabilities		Moments		Model statistics		KS test *p*-value		KS test *p*-value	
Regime	Model	qii	P(Rt=i)	E(X(t,i))	Var(X(t,i))	+/- IDR	+/- IQR	Regime	Model	Regime	Model
								(wavelet)		(sinusoid)	
		EEX1 (2001-2004) prices									
Base	Vas	0.9138	0.6874	30.9905	11.3089	3.91%	1.21%	0.9830	0.9860	0.5860	0.2497
Spike	SLN	0.8086	0.1102	44.6608	156.3268			0.0679		0.1031	
Drop	SLN	0.8053	0.2024	23.3121	19.9781			0.2609		0.0746	
Base	CIR	0.9794	0.9554	30.3542	23.3666	7.37%	5.26%	0.0908	0.0880	0.0191	0.0191
Spike	SLN	0.3339	0.0318	59.7470	447.3472			0.2290		0.0983	
Drop	SLN	0.8789	0.0128	11.7256	21.7531			0.7765		0.3180	
		EEX1 (2001-2004) log-prices									
Base	Vas	0.9052	0.6635	3.4377	0.0106	5.30%	-1.21%	0.6157	0.4069	0.9430	0.7483
Spike	SLN	0.8058	0.1034	3.7760	0.0531			0.2816		0.1108	
Drop	SLN	0.8105	0.2331	3.1421	0.0498			0.2252		0.0451	
Base	CIR	0.9638	0.9270	3.4085	0.0227	6.07%	1.34%	0.2802	0.2740	0.0371	0.0371
Spike	SLN	0.4775	0.0349	4.0046	0.0816			0.4938		0.4857	
Drop	SLN	0.8097	0.0381	2.7481	0.1528			0.0263		0.3166	

EEX1
(2001-2004)

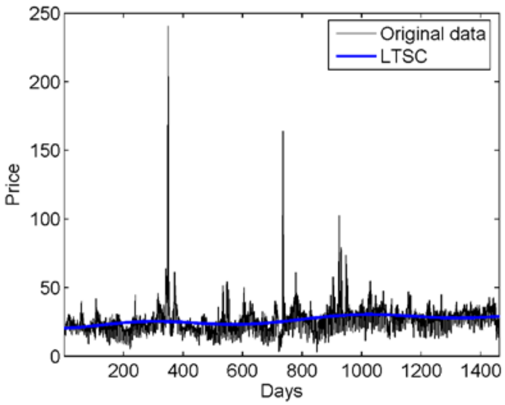

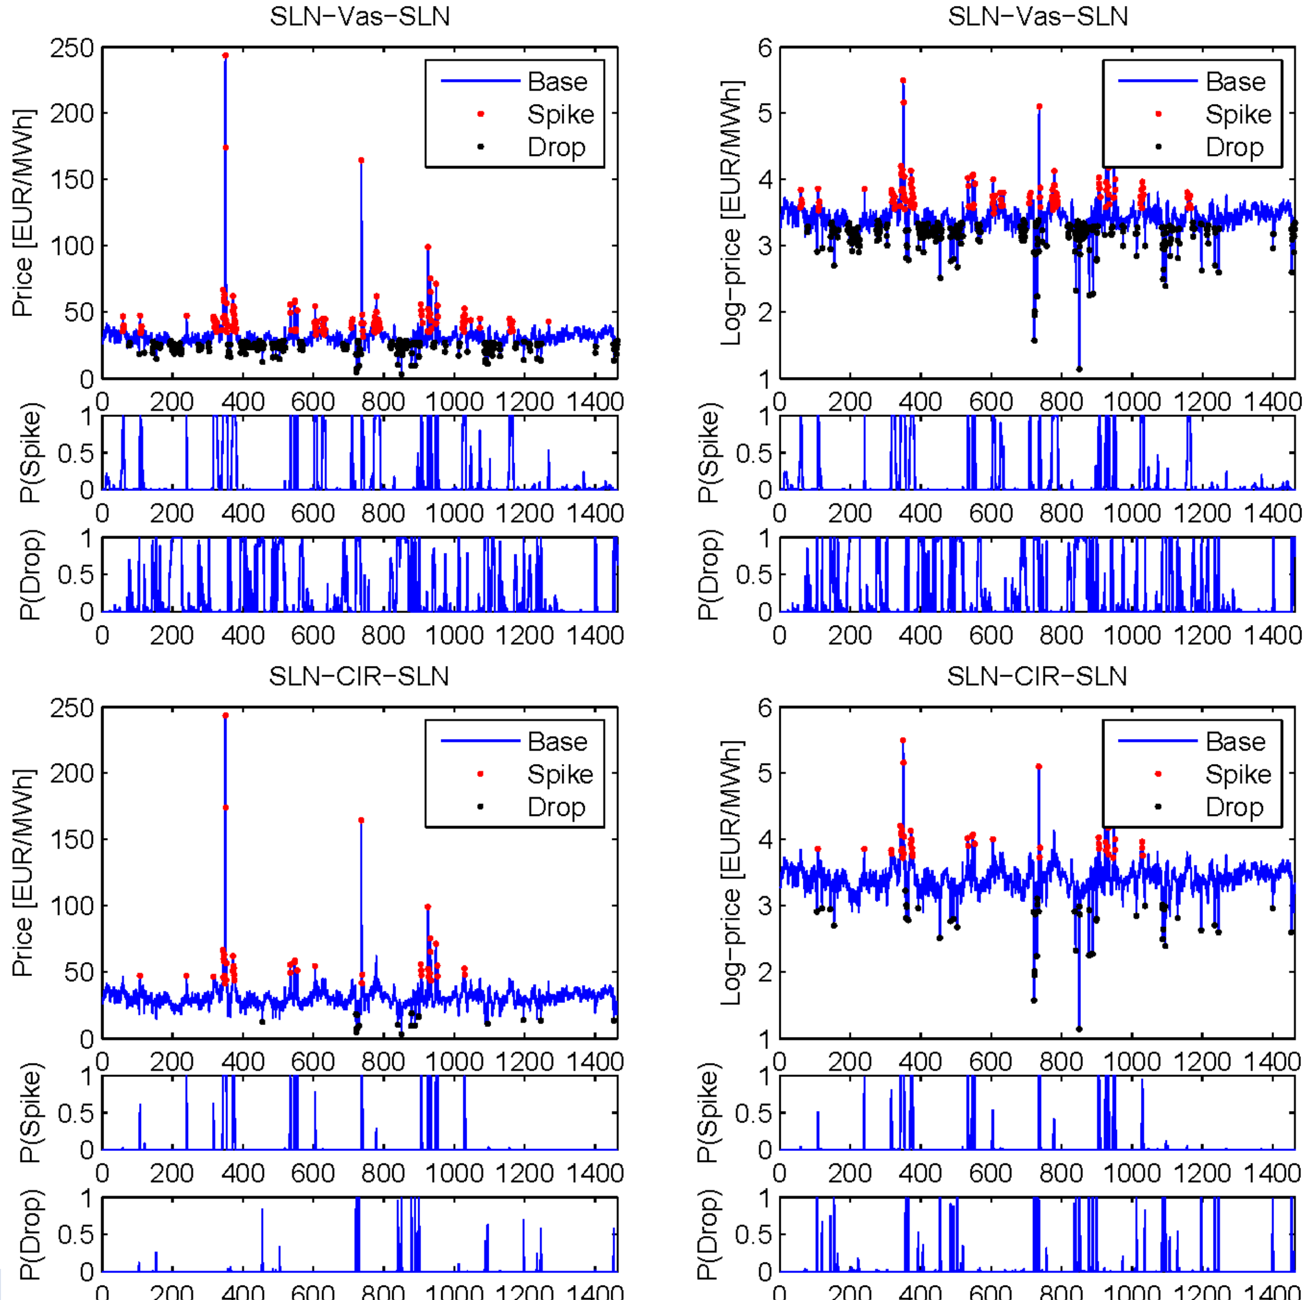

EEX2 (2005-2008):

		Probabilities		Moments		Model statistics		KS test *p*-value		KS test *p*-value	
Regime	Model	qii	P(Rt=i)	E(X(t,i))	Var(X(t,i))	+/- IDR	+/- IQR	Regime	Model	Regime	Model
								(wavelet)		(sinusoid)	
		EEX2 (2005-2008) prices									
Base	Vas	0.9373	0.6849	50.8182	40.9858	0.45%	2.09%	0.8239	0.9059	0.0027	0.0127
Spike	SLN	0.8895	0.1982	74.6810	371.0594			0.2138		0.1788	
Drop	SLN	0.7978	0.1169	33.8736	59.1922			0.6235		0.2520	
Base	CIR	0.9890	0.9759	52.2679	133.2471	6.15%	17.62%	0.0000	0.0000	0.0000	0.0000
Spike	SLN	0.5480	0.0222	117.4088	1154.2644			0.9364		0.9147	
Drop	SLN	0.8761	0.0019	12.7322	37.7523			0.7614		0.9570	
		EEX2 (2005-2008) log-prices									
Base	Vas	0.9464	0.7254	3.9411	0.0201	1.00%	5.59%	0.3010	0.3508	0.0133	0.1323
Spike	SLN	0.9115	0.1483	4.3249	0.0672			0.0724		0.3546	
Drop	SLN	0.7851	0.1263	3.5024	0.0680			0.2302		0.4559	
Base	CIR	0.9760	0.9286	3.9564	0.0380	2.03%	9.65%	0.0161	0.0240	0.0000	0.0001
Spike	SLN	0.5238	0.0210	4.7368	0.0722			0.9680		0.4417	
Drop	SLN	0.7936	0.0504	3.2941	0.1229			0.7128		0.5666	

EEX2
(2005–2008)

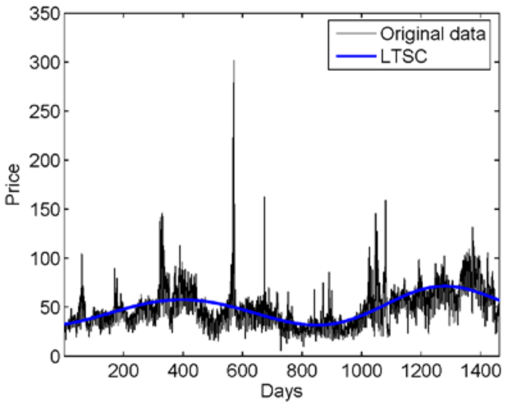

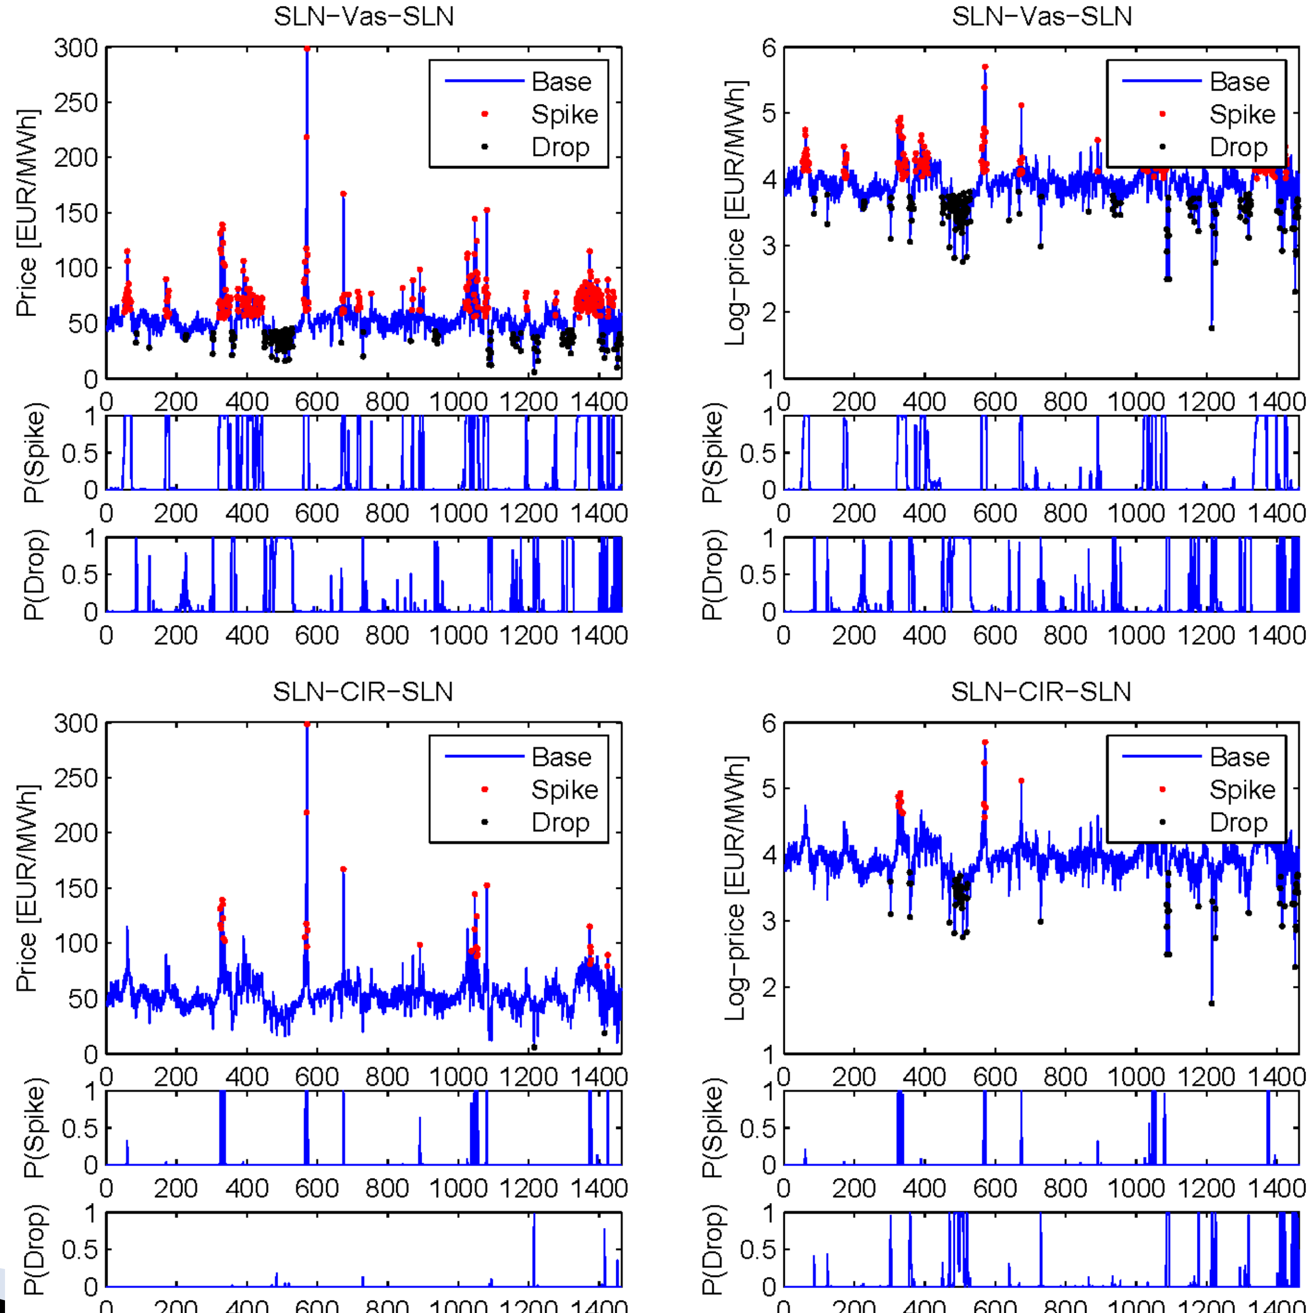

Conclusions

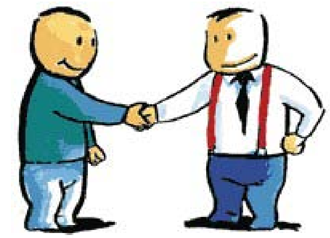

- When fitting 3-regime models
  - Practically all spikes are identified correctly
  - Vasicek base regime dynamics
    - Reasonable goodness-of-fit (high *p*-values)
    - Clusters of moderate prices identified as spikes or drops !
  - CIR base regime dynamics
    - Barely acceptable goodness-of-fit (low *p*-values)
    - No clusters of wrongly identified spikes or drops
- Deseasonalization method makes a difference

… as usual … more work is needed ☺

Post–talk conclusions

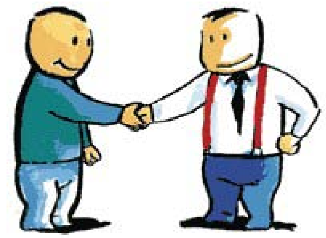

- The **quest** for the **model** is not over
  - Play with base regime dynamics
    - Use heavy tailed innovations in the CIR model
    - Use a different heteroskedastic mechanism
- The **devil** is in **deseasonalization**
  - Use fundamental data to better fit long term seasonality
- MRS models can be used to identify spikes in data
  - Spike identification is dependent on specification of the models for the regimes
    - Compare with other spike identification methods