# Black swans or dragon kings?

## A simple test for deviations from the power law

**Rafał Weron**

(joint work with Joanna Janczura)

Institute of Organization and Management

*Wrocław University of Technology*

# Motivation

- Asked to contribute to the Special Issue:
  - *From black swans to dragon kings.*
    *Is there life beyond power laws?*
  - Forthcoming in EPJ-Special Topics (2012)

- Never heard of 'dragon kings'
- Internal need to define 'dragon kings'
  - Hopefully develop a formal test
  - Check sample datasets

# Agenda

- **Definitions**
  - **Swans and dragons**
- The test
  - CLT-based and exact
- Empirical examples
  - Simulated data
  - CAT claims, financial drawdowns and electricity prices

- Talk based on: J. Janczura, R. Weron (2012) *Black swans or dragon kings? A simple test for deviations from the power law*, European Physical Journal - Special Topics (EPJ ST) 205, 79-93 (doi: 10.1140/epjst/e2012-01563-9)
- Matlab codes available from IDEAS: http://ideas.repec.org/s/wuu/hscode.html

# What is a ‘black swan’?

- An event which
  - Is an unpredictable outlier, beyond the realm of regular expectations
  - Carries an extreme impact
  - Human nature makes it explainable and predictable ... after it has happened
- Name origins
  - Black swans were undocumented until the 18th century

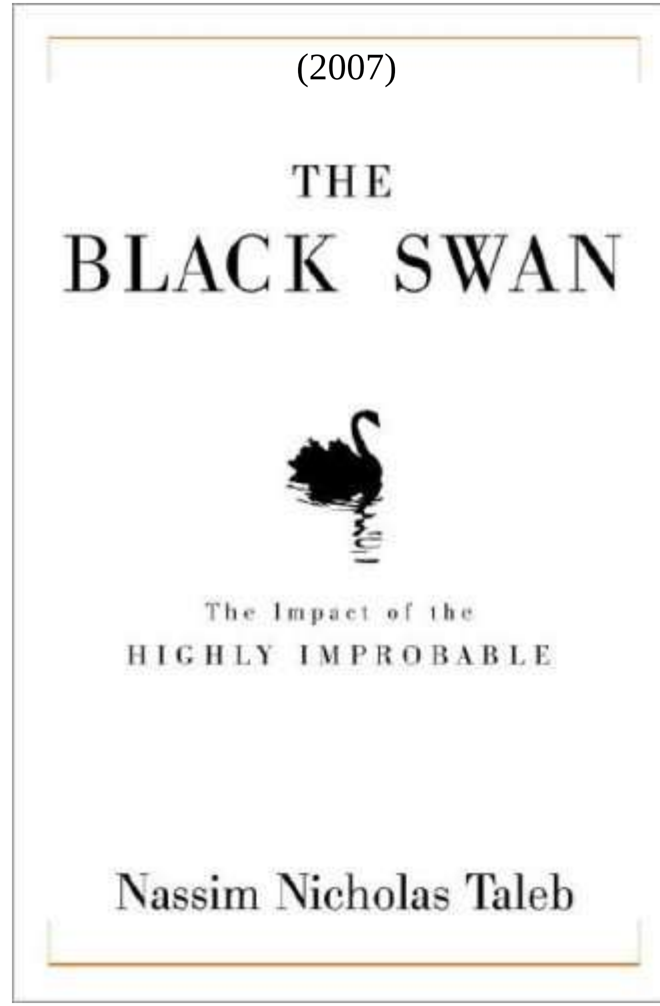

# Black swans and power laws

- Taleb regards almost all major scientific discoveries (PC, internet) and historical events (crash of 1987, 9/11 attack) as black swans
- He makes a distinction between
  - the ‘totally intractable’ black swans and
  - the ‘Mandelbrotian gray swans’, which are ‘tractable scientifically’ ... by means of power laws
  - But the gray swan terminology has not picked up

# Black vs. grey swans

- Why do black swans prevail?
- Power law events pretty well fit Taleb's definition anyway
  - Power law distributions are scale invariant
    - If extreme events are described by a power law distribution there is no way to predict them because nothing can distinguish them from the smaller events
  - … their extreme impacts are beyond the realm of normal expectations

# What is a 'dragon king'?

- Didier Sornette (2009)
  - *A meaningful outlier, which coexists with power laws*
  - Black swan critique
    - *Power laws: there is no way to predict extreme events because nothing can distinguish them from their small siblings*
    - *A great earthquake is just an earthquake that started small … and did not stop*
  - Dragon kings may have properties that make them … **predictable**

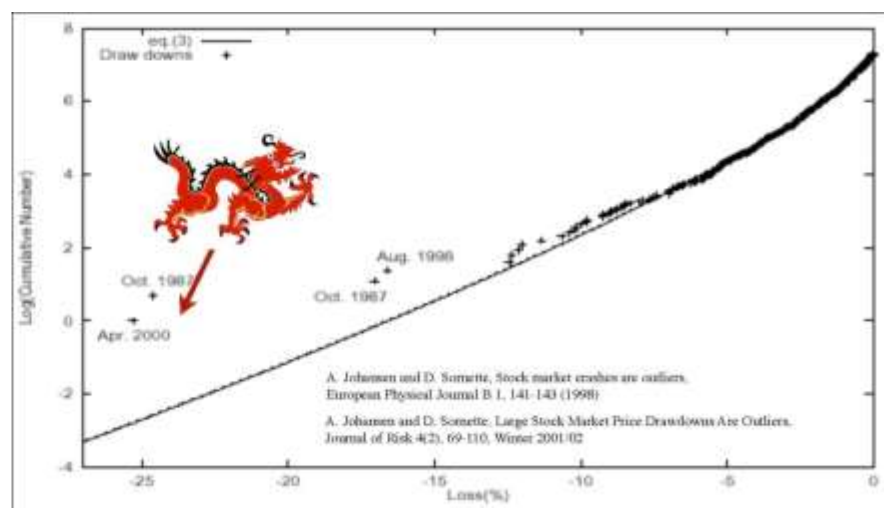

# Sornette's example

## Stretched exponential (Weibull) fit

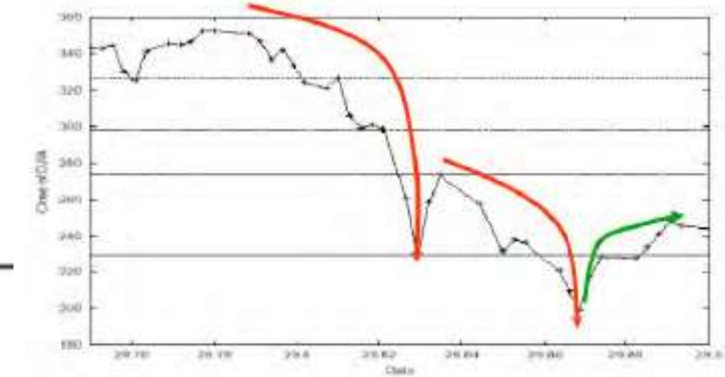

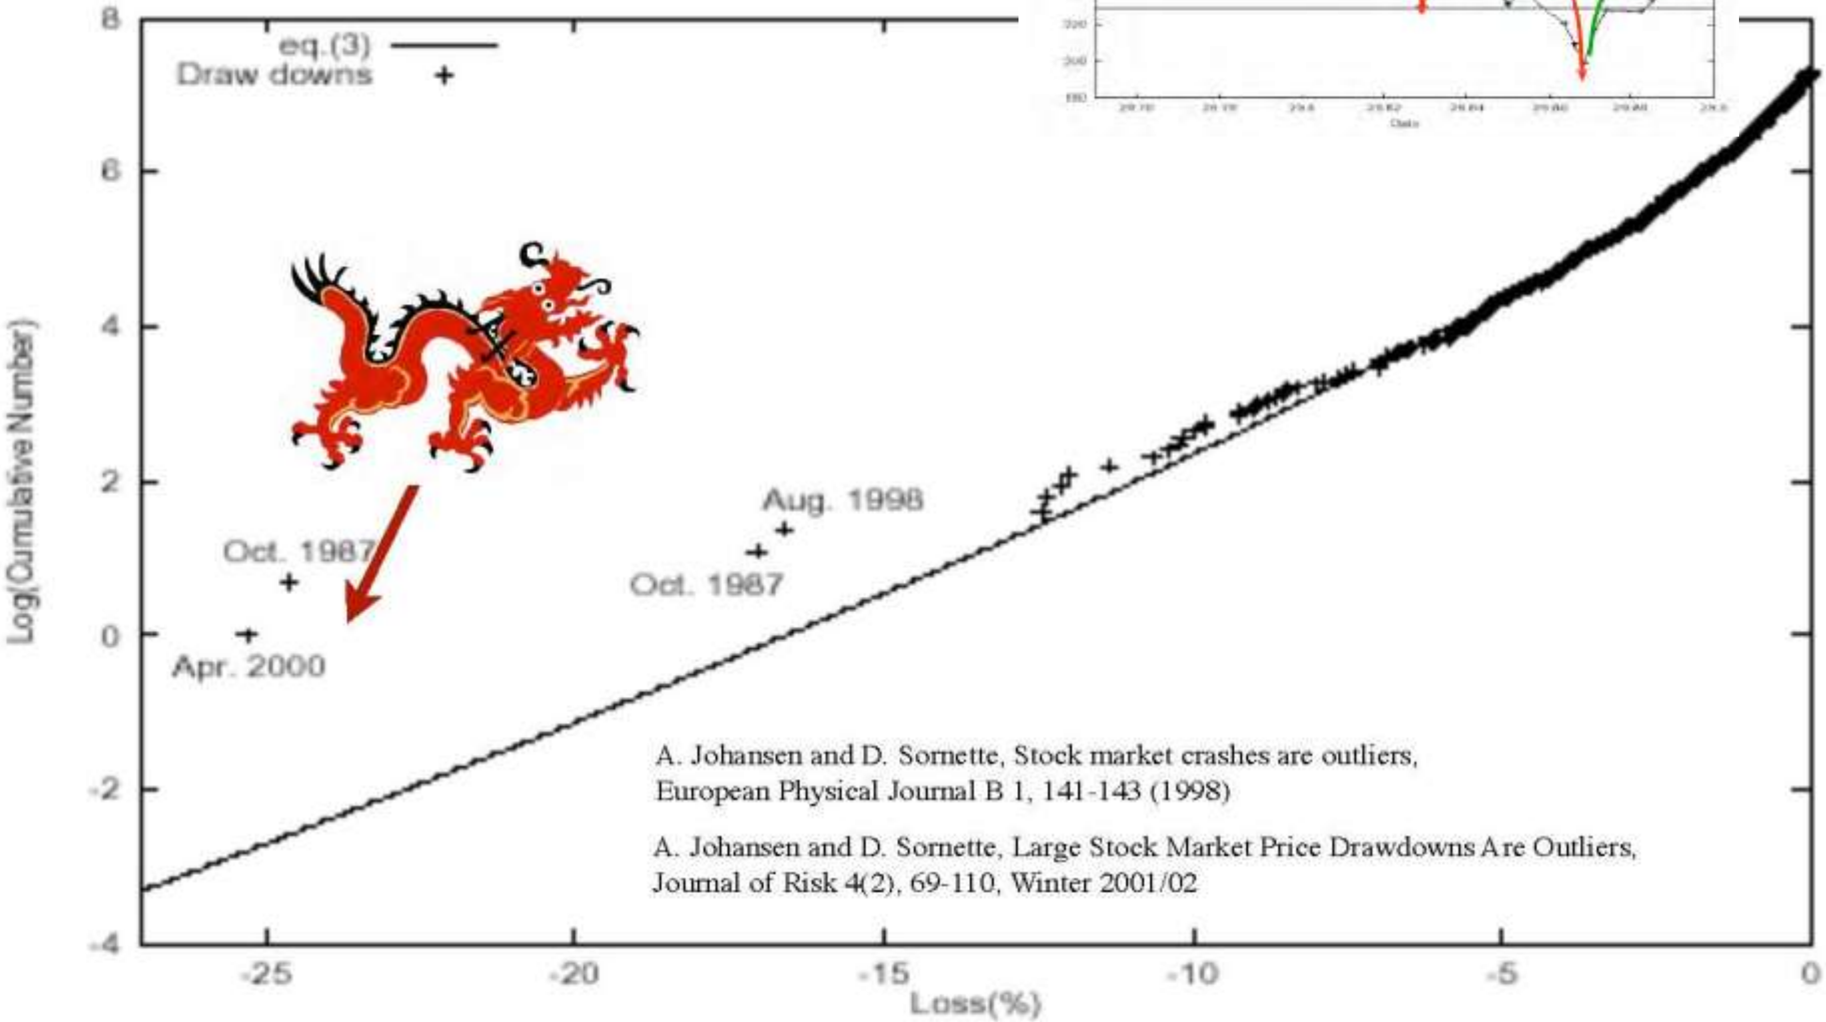

# Agenda

- Definitions
  - Swans and dragons
- **The test**
  - **CLT-based and exact**
- Empirical examples
  - Simulated data
  - CAT claims, financial drawdowns and electricity prices

- Talk based on: J. Janczura, R. Weron (2012) *Black swans or dragon kings? A simple test for deviations from the power law*, European Physical Journal - Special Topics (EPJ ST) 205, 79-93 (doi: 10.1140/epjst/e2012-01563-9)
- Matlab codes available from IDEAS: http://ideas.repec.org/s/wuu/hscode.html

# A simple test for dragon kings

- Recall, the empirical distribution function (edf):

$$F_n(x) = \frac{1}{n} \sum_{i=1}^{n} \mathbb{I}_{\{x_i < x\}}$$

  - defined for a sample of observations ($x_1$, $x_2$, ..., $x_n$)
- The number $Y$ of observations $x_i$ smaller than $x$ is binomially distributed with parameters $n$ and $r = F(x)$:

$$Y = \sum_{i=1}^{n} \mathbb{I}_{\{x_i \leq x\}} \sim B(n, F(x))$$

  - where $F(x)$ is the true cdf

# A simple test cont.

- The first parameter of the binomial law represents the number of trials
  - Here: sample size *n*
- And the second – the success probability in each trial
  - Here: the value of the cdf at point *x*
- Hence, we can write:

$$P\Big(q_{\frac{\alpha}{2}}[n, F(x)] < Y \leq q_{1-\frac{\alpha}{2}}[n, F(x)]\Big) = 1 - \alpha$$

  - where $q_\eta[n, r]$ is the $\eta$-quantile of the binomial distribution $B(n, r)$

# Confidence intervals (CI)

- Construct CI for the edf and an arbitrarily chosen confidence level (1−α)
  - For the left tail

$$P\left(\frac{1}{n}q_{\frac{\alpha}{2}}[n, F(x)] < F_n(x) \leq \frac{1}{n}q_{1-\frac{\alpha}{2}}[n, F(x)]\right) = 1 - \alpha$$

  - And for the right tail

$$P\left(\frac{1}{n}q_{\frac{\alpha}{2}}[n, 1 - F(x)] < 1 - F_n(x) \leq \frac{1}{n}q_{1-\frac{\alpha}{2}}[n, 1 - F(x)]\right) = 1 - \alpha$$

# Confidence intervals (CI)

- Note, that these are pointwise intervals
  - For each $x$, we are $(1-\alpha)\times 100\%$ confident of observing edf($x$) within those limits
- This is not the same as constructing 'confidence bands'
  - Which guarantee that the edf falls within the band for all $x$'s (in some interval)
  - Confidence bands are wider than the curves one obtains by using pointwise CI

# A simple test ...

- Assume power law tails of *F*:

$$F(x) \approx b_1 |x|^{p_1} \quad \text{for} \quad x \to -\infty,$$
$$1 - F(x) \approx b_2 x^{p_2} \quad \text{for} \quad x \to \infty.$$

- Hence, with prob. (1−α) the left tail of the edf should lie in the interval:

$$\left( \frac{1}{n} q_{\frac{\alpha}{2}} [n, b_1 |x|^{p_1}], \frac{1}{n} q_{1-\frac{\alpha}{2}} [n, b_1 |x|^{p_1}] \right]$$

- And with prob. (1−α) the right tail in the interval:

$$\left( \frac{1}{n} q_{\frac{\alpha}{2}} [n, b_2 x^{p_2}], \frac{1}{n} q_{1-\frac{\alpha}{2}} [n, b_2 x^{p_2}] \right]$$

# A simple test ...

- Now, it suffices to
  - Fit a power law to the left or right tail of the edf
  - Plot the respective intervals
- Observations lying outside the curves spanned by the CI are likely to be dragon kings (with prob. $1-\alpha$)

# A CLT-based test

- … but the computation of the CI is not straightforward
  - There are no closed form formulas (not involving special functions) for the inverse of the binomial cdf
- CLT: $F_n(x)$ is asymptotically normally distributed

$$\frac{\sqrt{n}[F_n(x) - F(x)]}{\sqrt{F(x)[1 - F(x)]}} \xrightarrow{d} N(0,1)$$

# A CLT-based test cont.

- Provided that *n* is large enough

$$P\left( z_{\frac{\alpha}{2}} < \frac{\sqrt{n}[F_n(x) - F(x)]}{\sqrt{F(x)[1 - F(x)]}} < z_{1-\frac{\alpha}{2}} \right) \approx 1 - \alpha$$

  - where $z_q$ is the *q*-quantile of N(0,1)
- Hence for the left tail we have

$$P\left( F(x) + \sqrt{\frac{F(x)[1 - F(x)]}{n}} z_{\frac{\alpha}{2}} < F_n(x) < F(x) + \sqrt{\frac{F(x)[1 - F(x)]}{n}} z_{1-\frac{\alpha}{2}} \right) \approx 1 - \alpha$$

- And analogously for the right …

# A CLT-based test cont.

- Assume power law tails of *F*
- Hence, with prob. (1−α) the left tail of the edf should lie in the interval:

$$\left( b_1|x|^{p_1} + \sqrt{\frac{b_1|x|^{p_1}(1-b_1|x|^{p_1})}{n}}\, z_{\frac{\alpha}{2}},\ b_1|x|^{p_1} + \sqrt{\frac{b_1|x|^{p_1}(1-b_1|x|^{p_1})}{n}}\, z_{1-\frac{\alpha}{2}} \right)$$

- And with prob. (1−α) the right tail in the interval:

$$\left( b_2x^{p_2} + \sqrt{\frac{b_2x^{p_2}(1-b_2x^{p_2})}{n}}\, z_{\frac{\alpha}{2}},\ b_2x^{p_2} + \sqrt{\frac{b_2x^{p_2}(1-b_2x^{p_2})}{n}}\, z_{1-\frac{\alpha}{2}} \right)$$

# Exact vs. CLT-based CI

- To construct the CLT-based CI only one standard normal quantile (symmetry) is needed

$$\left(b_1|x|^{p_1} + \sqrt{\frac{b_1|x|^{p_1}(1-b_1|x|^{p_1})}{n}}z_{\frac{\alpha}{2}},\ b_1|x|^{p_1} + \sqrt{\frac{b_1|x|^{p_1}(1-b_1|x|^{p_1})}{n}}z_{1-\frac{\alpha}{2}}\right)$$

- For the exact CI two quantiles of the binomial distribution have to be found for each point x

$$\left(\frac{1}{n}q_{\frac{\alpha}{2}}[n, b_1|x|^{p_1}], \frac{1}{n}q_{1-\frac{\alpha}{2}}[n, b_1|x|^{p_1}]\right]$$

  - For a sample size of 5000 observations, for a power law fitted to the largest 10%-1% values the CLT-based test is over 220 times faster (0.003 vs. 0.665 seconds on an i7-820QM w. Matlab 7.9)

# The tests are ‘distribution free’

- The true distribution can be arbitrary, say, Weibull (also known as stretched exponential)
  - The CI would be computed using the Weibull cdf

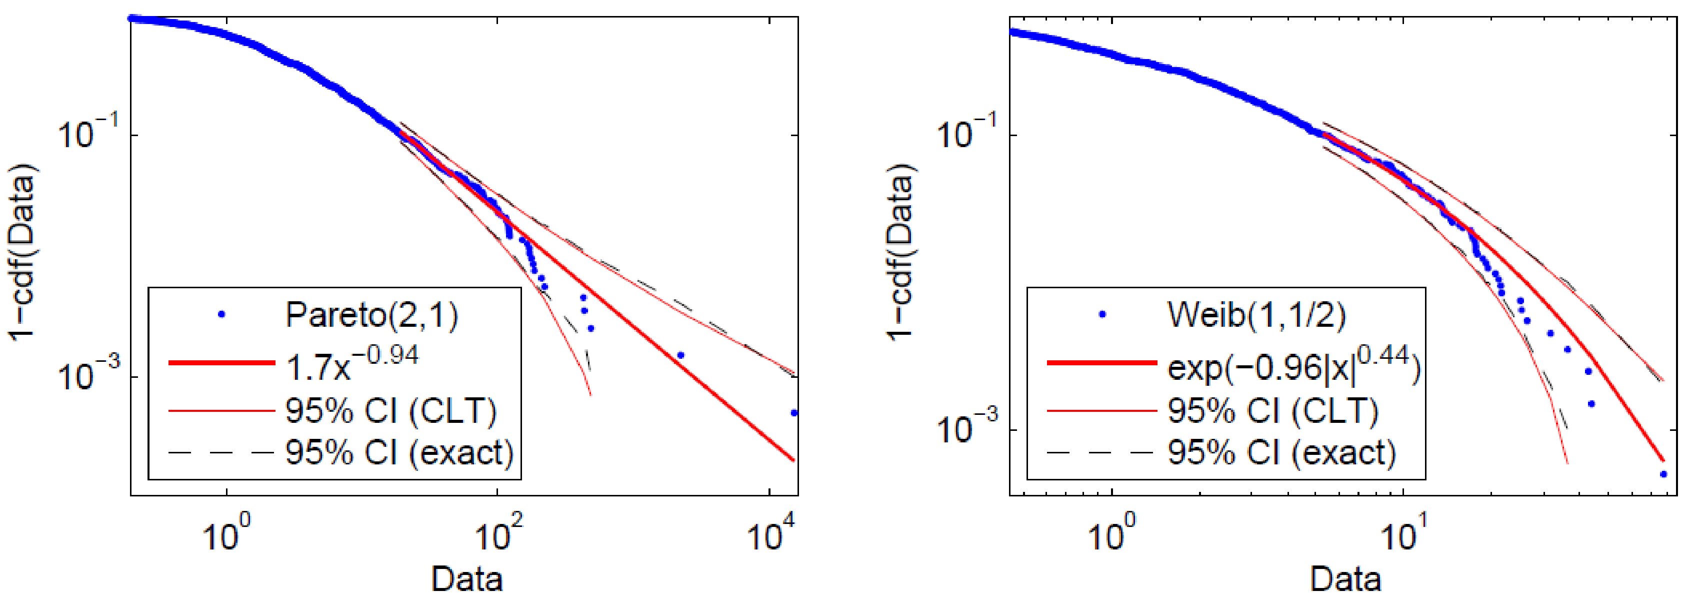

# Agenda

- Definitions
  - Swans and dragons
- The test
  - CLT-based and exact
- **Empirical examples**
  - **Simulated data**
  - CAT claims, financial drawdowns and electricity prices

- Talk based on: J. Janczura, R. Weron (2012) *Black swans or dragon kings? A simple test for deviations from the power law*, European Physical Journal - Special Topics (EPJ ST) 205, 79-93 (doi: 10.1140/epjst/e2012-01563-9)
- Matlab codes available from IDEAS: http://ideas.repec.org/s/wuu/hscode.html

# Power law tails: Exact CI

**Table 1.** Percentage of pointwise outliers with respect to the exact, binomial law-based CI given by formula (5) and $F(x)$ being a power law fitted to the 10%-1% or 25%-2.5% largest observations. For Cauchy and Pareto distributions, outliers with respect to the true power law (implied by the cdf) are also provided. The number of simulated samples is equal to $10^4$.

| CI | Sample size | Cauchy(0,1) True | Cauchy(0,1) Fitted | Pareto(2,1) True | Pareto(2,1) Fitted | Hyp(2,1) Fitted | Weib$(1,\frac{1}{2})$ Fitted |
|---|---|---|---|---|---|---|---|
| | | *Power law fitted to 10%-1% largest observations* | | | | | |
| 90% | 1000 | 10.0% | 4.5% | 9.8% | 4.5% | 24.7% | 16.6% |
| 95% | 1000 | 5.2% | 2.2% | 4.6% | 2.1% | 10.7% | 6.8% |
| 99% | 1000 | 1.1% | 0.3% | 1.0% | 0.3% | 0.7% | 0.2% |
| | | *Power law fitted to 25%-2.5% largest observations* | | | | | |
| 90% | 1000 | " | 12.4% | " | 11.1% | 100.0% | 96.9% |
| 95% | 1000 | " | 6.1% | " | 5.4% | 99.9% | 92.2% |
| 99% | 1000 | " | 1.2% | " | 0.8% | 98.2% | 70.6% |
| | | *Power law fitted to 10%-1% largest observations* | | | | | |
| 90% | 5000 | 9.9% | 10.1% | 10.0% | 11.0% | 99.8% | 98.5% |
| 95% | 5000 | 5.1% | 5.2% | 4.8% | 5.4% | 99.3% | 96.7% |
| 99% | 5000 | 1.1% | 0.7% | 1.0% | 1.1% | 96.3% | 84.8% |

# Power law tails : CTL-based CI

**Table 2.** Percentage of pointwise outliers with respect to the approximate, CLT-based CI given by formula (14) and $F(x)$ being a power law fitted to the 10%-1% or 25%-2.5% largest observations. For Cauchy and Pareto distributions, outliers with respect to the true power law (implied by the cdf) are also provided. The number of simulated samples is equal to $10^4$.

| CI | Sample size | Cauchy(0,1) True | Cauchy(0,1) Fitted | Pareto(2,1) True | Pareto(2,1) Fitted | Hyp(2,1) Fitted | Weib$(1,\frac{1}{2})$ Fitted |
|---|---|---|---|---|---|---|---|
| | | | *Power law fitted to 10%-1% largest observations* | | | | |
| 90% | 1000 | 9.6% | 4.2% | 9.1% | 4.3% | 29.2% | 20.4% |
| 95% | 1000 | 4.8% | 2.1% | 4.1% | 1.7% | 10.1% | 6.1% |
| 99% | 1000 | 1.1% | 0.6% | 0.9% | 0.4% | 0.1% | <0.1% |
| | | | *Power law fitted to 25%-2.5% largest observations* | | | | |
| 90% | 1000 | " | 10.7% | " | 11.2% | 99.1% | 90.9% |
| 95% | 1000 | " | 4.4% | " | 4.7% | 99.8% | 78.1% |
| 99% | 1000 | " | 0.7% | " | 0.5% | 74.3% | 35.0% |
| | | | *Power law fitted to 10%-1% largest observations* | | | | |
| 90% | 5000 | 10.0% | 9.9% | 9.8% | 11.3% | 99.8% | 98.9% |
| 95% | 5000 | 4.6% | 4.3% | 4.7% | 5.1% | 99.5% | 96.5% |
| 99% | 5000 | 1.0% | 1.2% | 1.0% | 0.6% | 93.3% | 78.9% |

# Weibull tails: Exact CI

**Table 3.** Percentage of pointwise outliers with respect to the exact, binomial law-based CI given by formula (5) and $F(x)$ being a stretched exponential (or Weibull) law fitted to the 10%-1% or 25%-2.5% largest observations. For the Weibull distribution, outliers with respect to the tail of the true cdf are also provided. The number of simulated samples is equal to $10^4$.

| CI | Sample size | Cauchy(0,1) Fitted | Pareto(2,1) Fitted | Hyp(2,1) Fitted | Weib$(1,\frac{1}{2})$ True | Weib$(1,\frac{1}{2})$ Fitted |
|---|---|---|---|---|---|---|
| | | *Weibull tail fitted to 10%-1% largest observations* | | | | |
| 90% | 1000 | 66.1% | 63.3% | 10.6% | 10.7% | 11.2% |
| 95% | 1000 | 57.5% | 53.5% | 5.5% | 4.8% | 6.6% |
| 99% | 1000 | 41.5% | 38.6% | 0.8% | 1.2% | 1.9% |
| | | *Weibull tail fitted to 25%-2.5% largest observations* | | | | |
| 90% | 1000 | 97.4% | 97.9% | 59.3% | " | 14.6% |
| 95% | 1000 | 95.8% | 97.0% | 47.6% | " | 8.0% |
| 99% | 1000 | 90.6% | 92.7% | 25.8% | " | 2.4% |

# Weibull tails : CTL-based CI

**Table 4.** Percentage of pointwise outliers with respect to the approximate, CLT-based CI given by formula (14) and $F(x)$ being a stretched exponential (or Weibull) law fitted to the 10%-1% or 25%-2.5% largest observations. For the Weibull distribution, outliers with respect to the tail of the true cdf are also provided. The number of simulated samples is equal to $10^4$.

| CI | Sample size | Cauchy(0,1) Fitted | Pareto(2,1) Fitted | Hyp(2,1) Fitted | Weib$(1,\frac{1}{2})$ True | Weib$(1,\frac{1}{2})$ Fitted |
|---|---|---|---|---|---|---|
| | | *Weibull tail fitted to 10%-1% largest observations* | | | | |
| 90% | 1000 | 62.1% | 58.3% | 11.5% | 9.6% | 10.6% |
| 95% | 1000 | 56.4% | 52.4% | 4.9% | 4.8% | 5.8% |
| 99% | 1000 | 44.8% | 41.3% | 0.6% | 1.1% | 2.4% |
| | | *Weibull tail fitted to 25%-2.5% largest observations* | | | | |
| 90% | 1000 | 96.7% | 97.5% | 63.2% | " | 14.0% |
| 95% | 1000 | 95.3% | 96.6% | 47.7% | " | 8.0% |
| 99% | 1000 | 91.8% | 94.1% | 19.0% | " | 2.6% |

# Agenda

- Definitions
  - Swans and dragons
- The test
  - CLT-based and exact
- Empirical examples
  - Simulated data
  - **CAT claims, financial drawdowns and electricity prices**

- Talk based on: J. Janczura, R. Weron (2012) *Black swans or dragon kings? A simple test for deviations from the power law*, European Physical Journal - Special Topics (EPJ ST) 205, 79-93 (doi: 10.1140/epjst/e2012-01563-9)
- Matlab codes available from IDEAS: http://ideas.repec.org/s/wuu/hscode.html

# Catastrophe claims

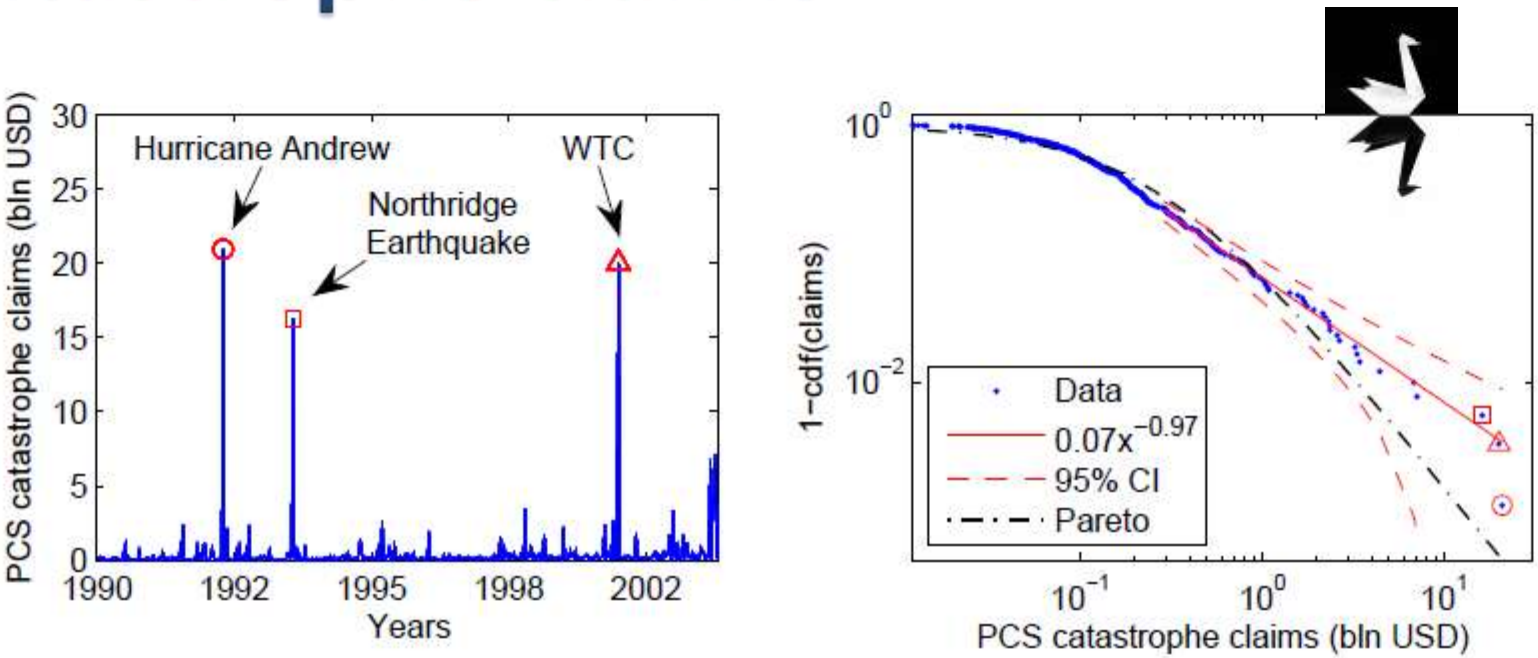


**Fig. 2.** *Left panel:* PCS catastrophe loss data, 1990-2004. The three largest losses in this period were caused by Hurricane Andrew (24 August 1992), the Northridge Earthquake (17 January 1994) and the terrorist attack on WTC (11 September 2001). *Right panel:* Right tail of the empirical distribution of claim sizes. The three largest losses do not deviate significantly from the fitted power law. However, two out of three seem to be outliers with respect to the Pareto law fit.

# Financial ‘drawdowns/ups’

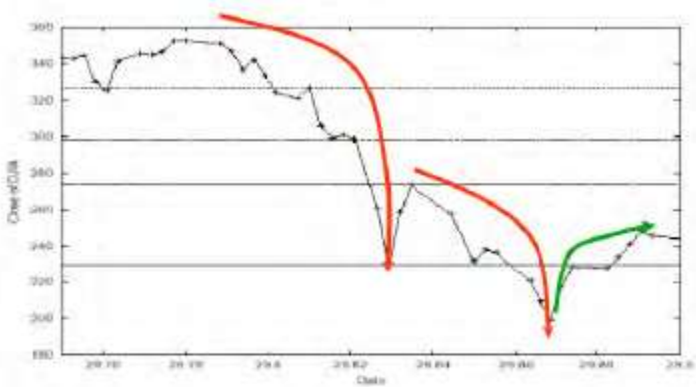

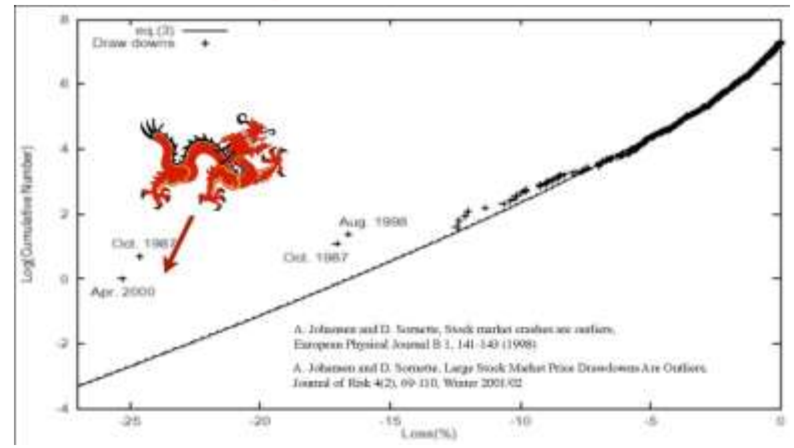


5.2.1971- 31.5.2000

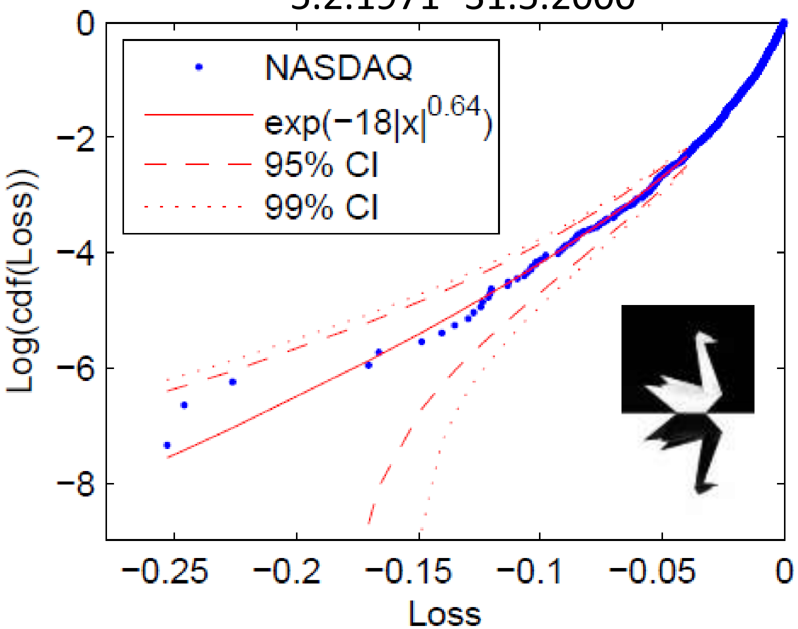


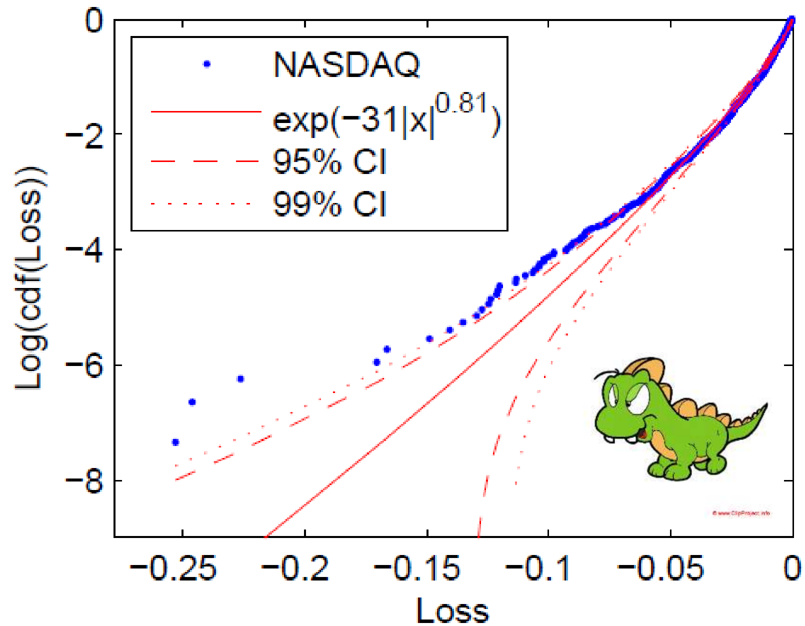

# Electricity spot prices

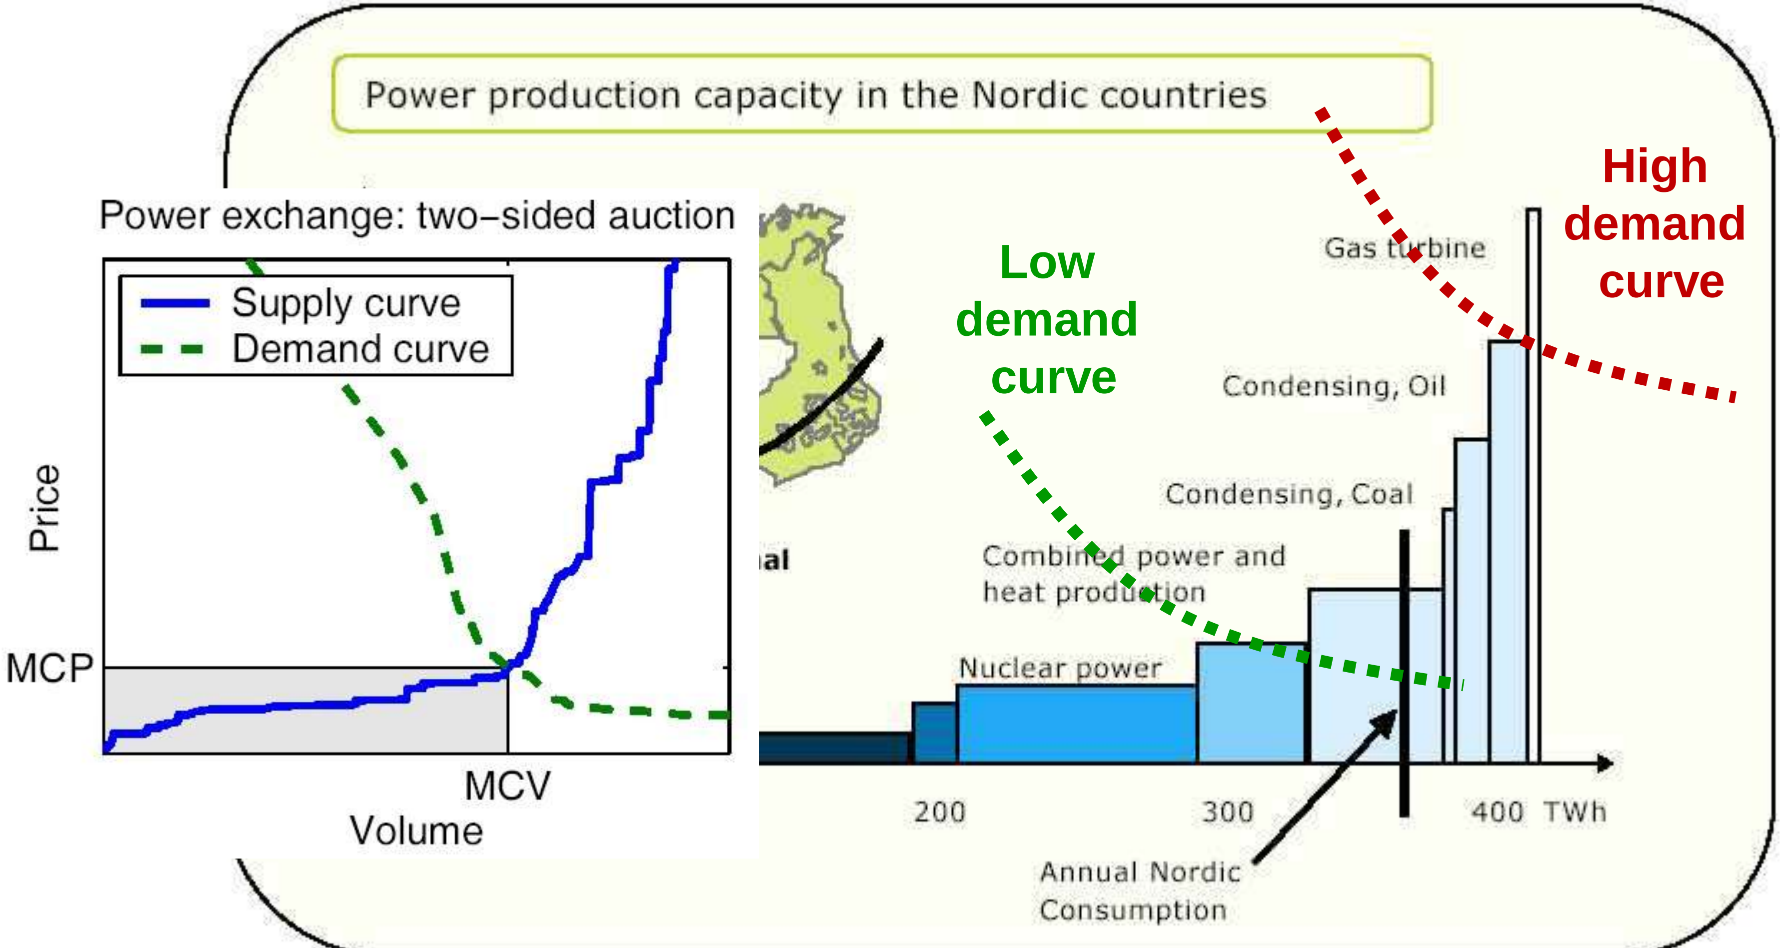

# Price spikes … are transient

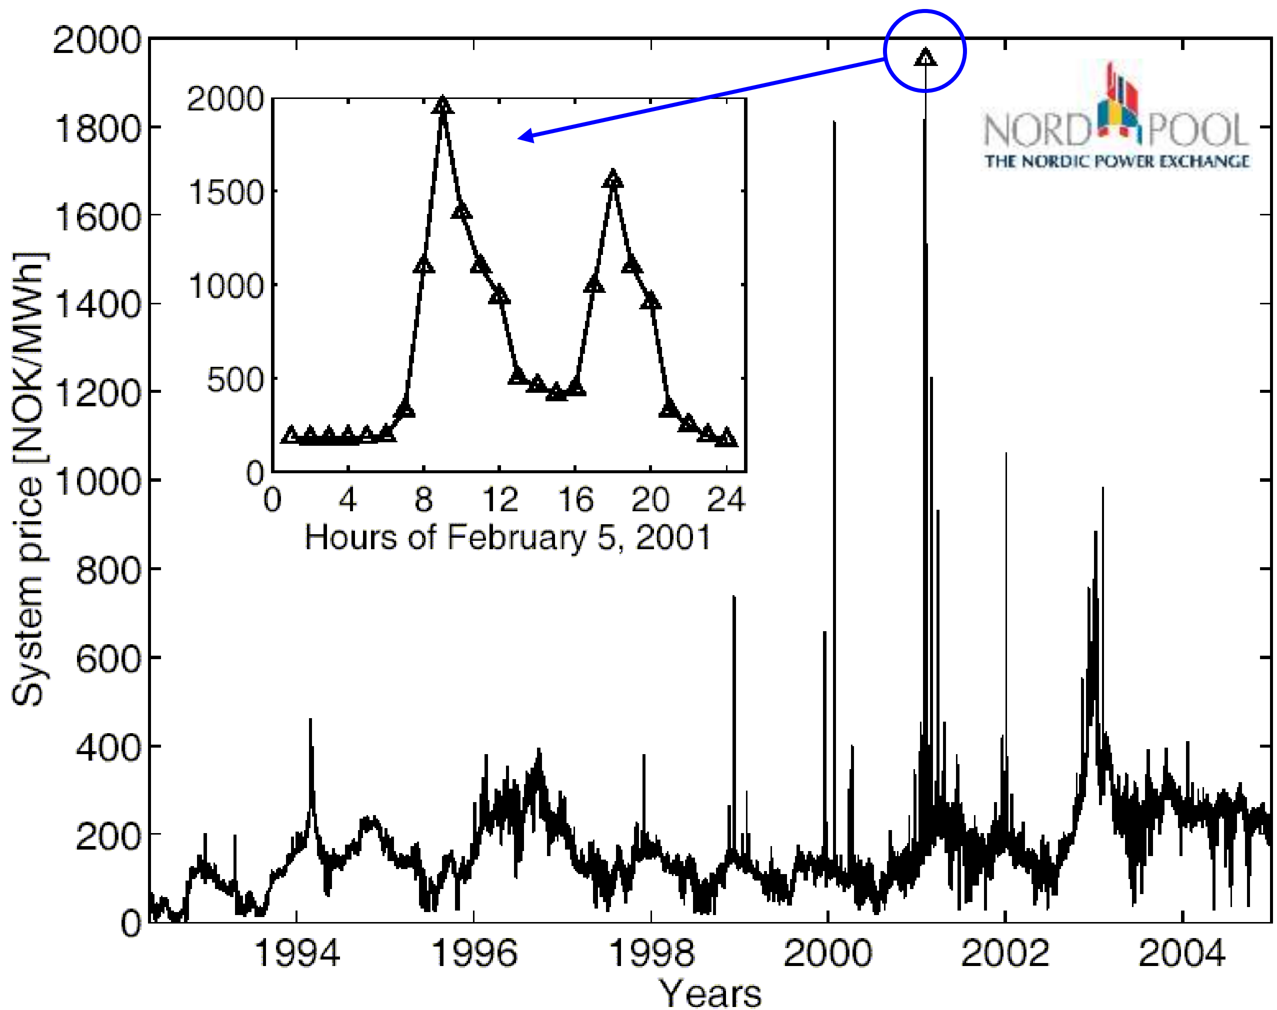

# German (EEX) and Australian (NSW) spot prices

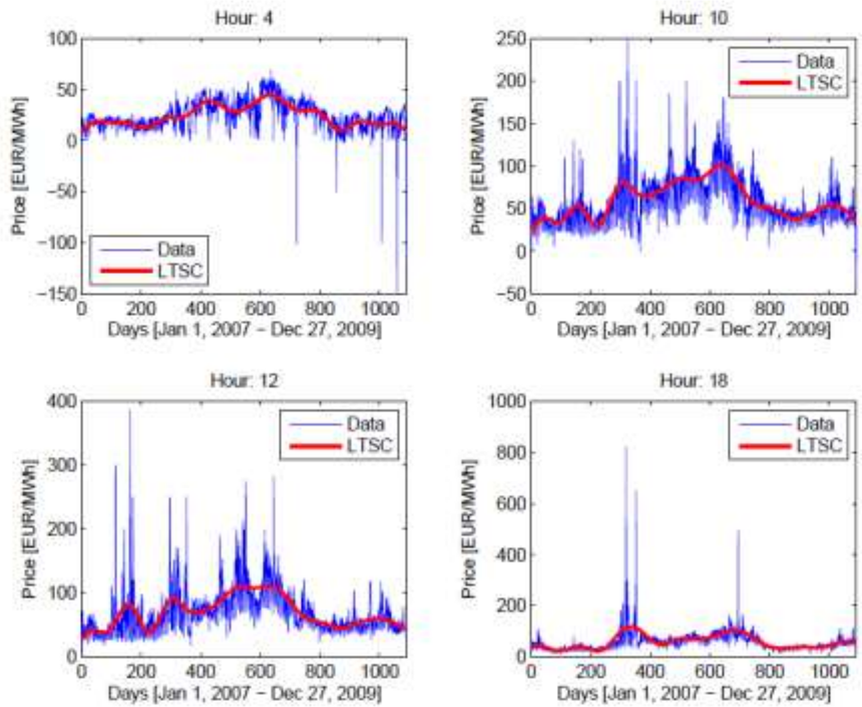


Fig. 3. Electricity spot prices and the estimated LTSCs in the German EEX market for four sample hours and 1092 days in the period Jan. 1, 2007 – Dec. 27, 2009. Note, the different scales in the four panels.

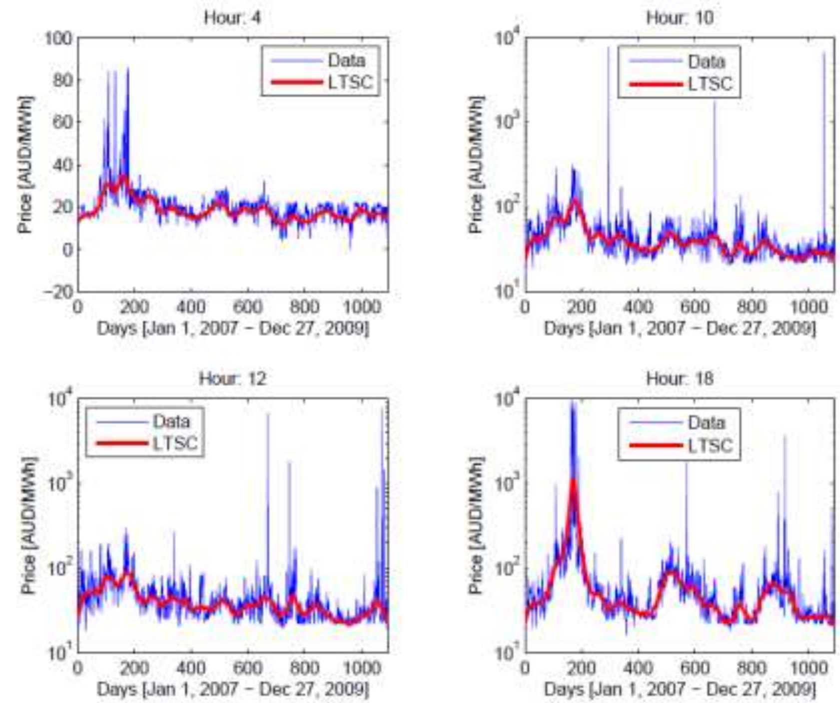


Fig. 4. Electricity spot prices and the estimated LTSCs in the Australian NSW market for four sample hours (in fact, half-hourly intervals starting at full hours) and 1092 days in the period Jan. 1, 2007 – Dec. 27, 2009. Note, the semilogarithmic scale for hours 10, 12 and 18. The few extremely spiky prices would render the 'normal' prices invisible on a linear scale.

# EEX: edf right tails

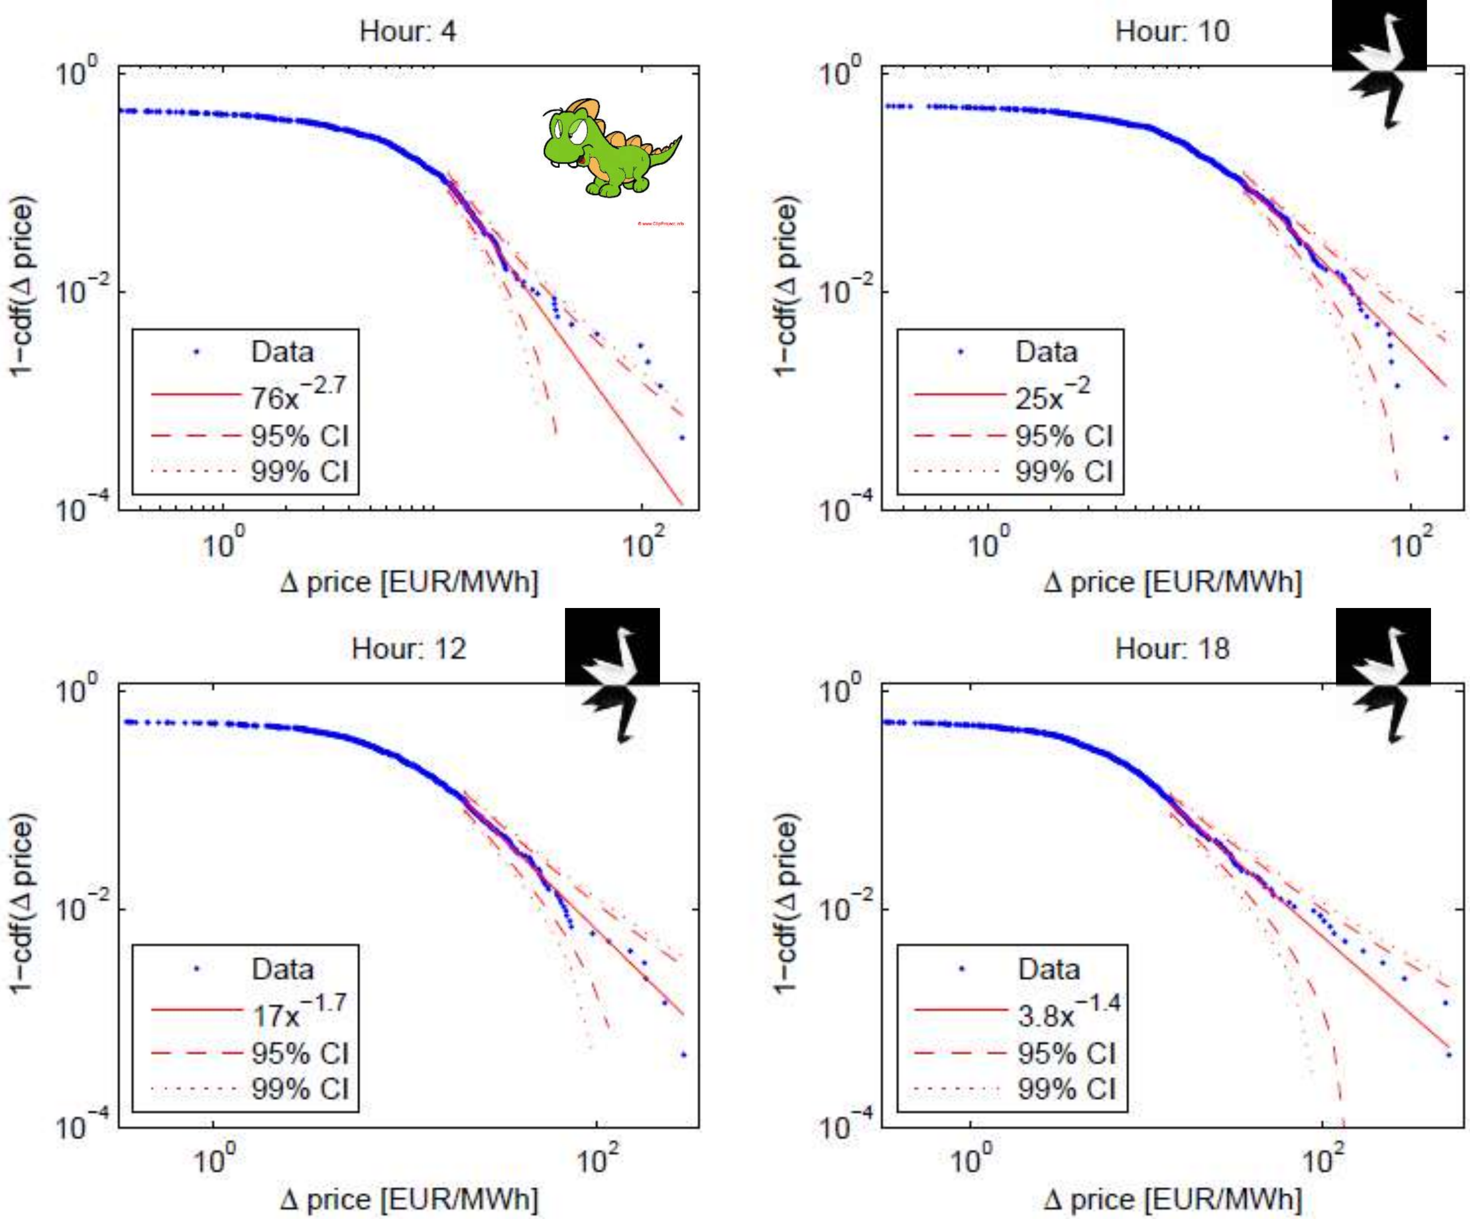


**Fig. 5.** Left (*four upper panels*) and right (*four lower panels*) tails of the empirical distribution of electricity spot price changes, computed for EEX prices depicted in Figure 3. The solid lines represent the fitted power law tails (to the lowest or highest 1%-10% of observations). The dashed and dotted curves indicate the 95% and 99% CI, respectively.

# NSW: edf right tails

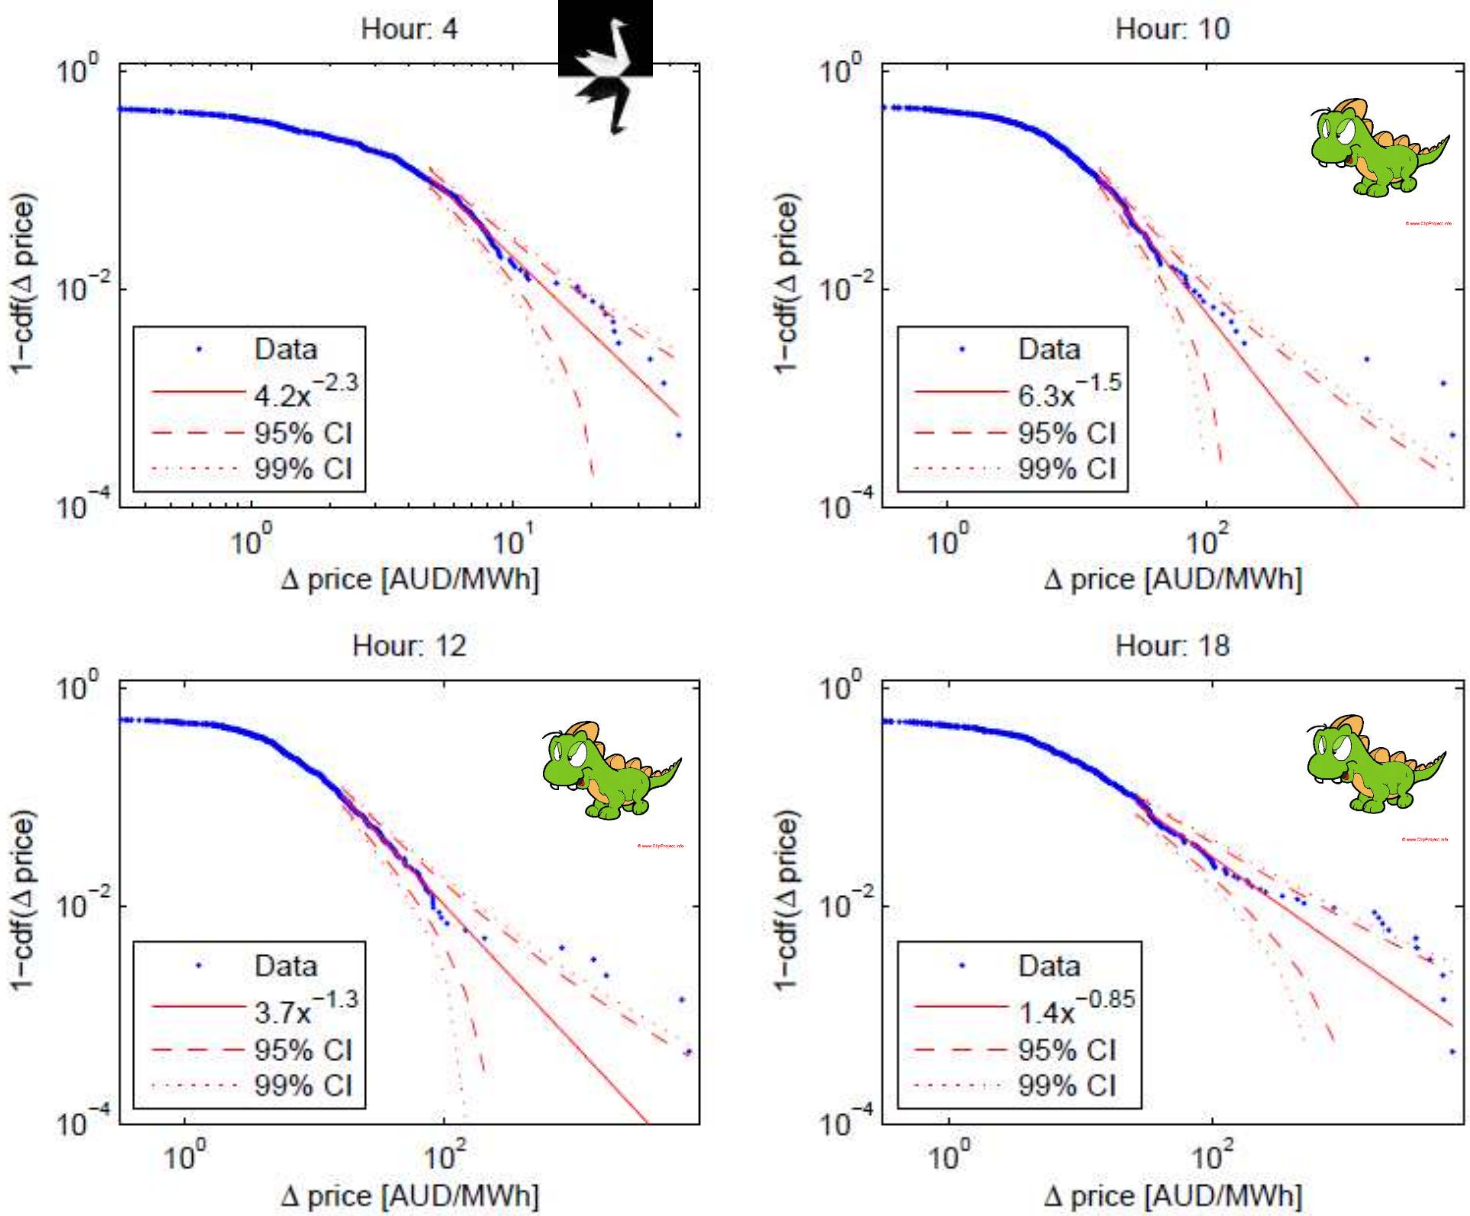


**Fig. 6.** Left (*four upper panels*) and right (*four lower panels*) tails of the empirical distribution of electricity spot price changes, computed for NSW prices depicted in Figure 4. The solid lines represent the fitted power law tails (to the lowest or highest 1%-10% of observations). The dashed and dotted curves indicate the 95% and 99% CI, respectively.

# Conclusions

- Yes, there are **outliers to power laws**
- We can call them **dragon kings**
- … but can we **predict** them

# Can we predict the electricity price spikes?

- Compute the ratio of demand forecast *D(t,T)* and predicted generation capacity *C(t,T)*

$$\rho(t,T) = \frac{D(t,T)}{C(t,T)}$$

- The ratio may take values > 1
- When the ratio is
  - 'small' - it is less likely for spikes to appear
  - 'large' - it is likely to observe price spikes

# Predicting the spikes ...

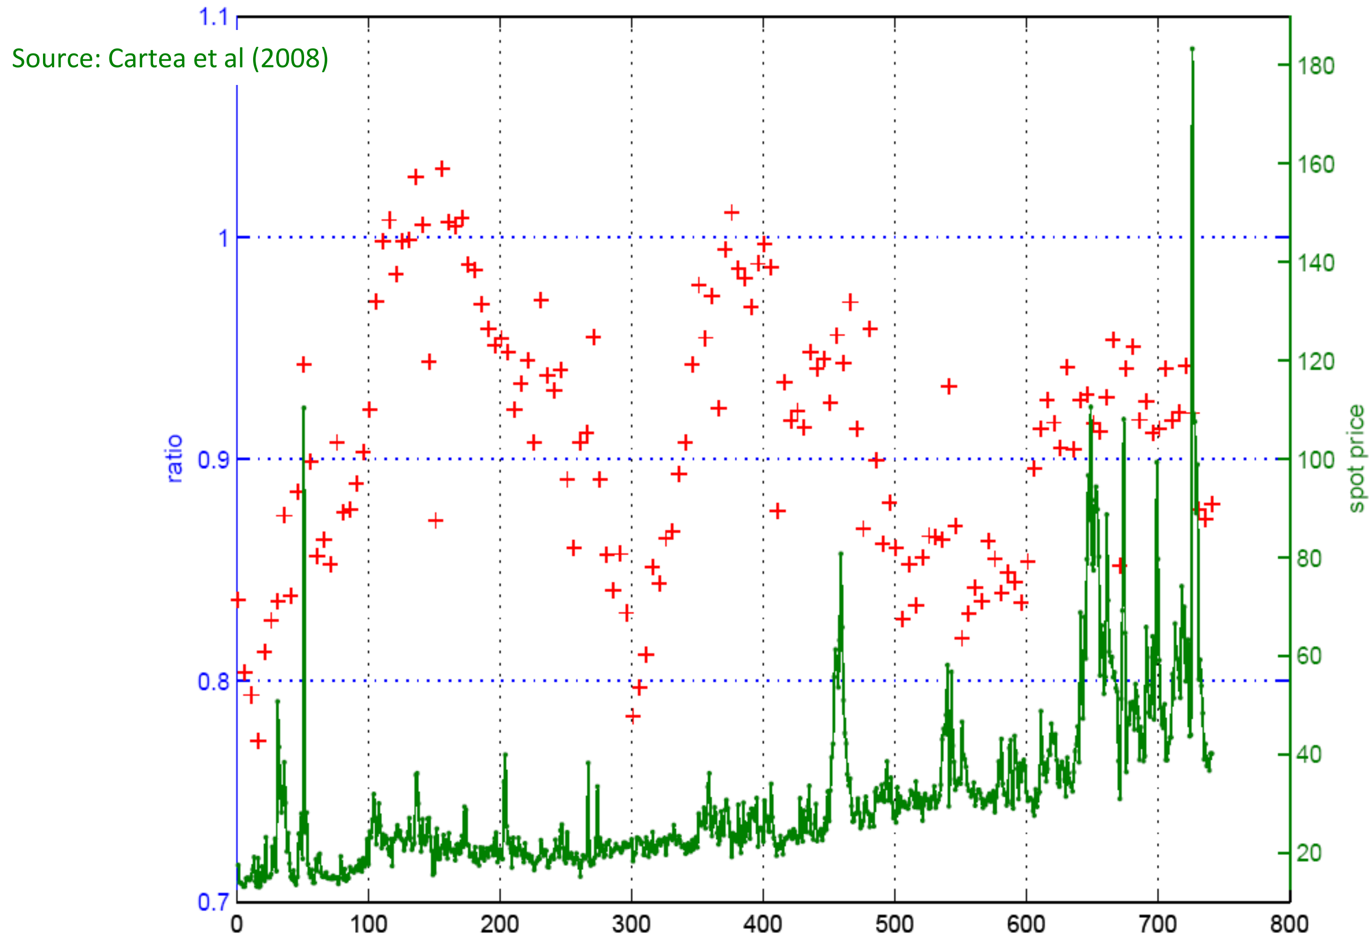

# Predicting the spikes ...

| Bin | #spikes | #weeks | Spike numbers |
|---|---|---|---|
| [0.77893, 0.79185) | 0 | 2 | n.a. |
| [0.79185, 0.80476) | 0 | 2 | n.a. |
| [0.81768, 0.83059) | 0 | 1 | n.a |
| [0.83059, 0.84351) | 1 | 3 | 31 |
| [0.84351, 0.85642) | 0 | 9 | n.a |
| [0.85642, 0.86934) | 0 | 7 | n.a |
| [0.86934, 0.88225) | 0 | 13 | n.a |
| [0.88225, 0.89517) | 0 | 12 | n.a |
| [0.89517, 0.90808) | 0 | 7 | n.a |
| [0.90808, 0.92100) | 3 | 7 | 267 , 699, 726 |
| [0.92100, 0.93391) | 3 | 14 | 543 , 649, 661 |
| [0.93391, 0.94682) | 3 | 15 | 51, 459, 674 |
| [0.94682, 0.95974) | 2 | 7 | 203, 274 |
| [0.95974, 0.97265) | 0 | 16 | n.a. |
| [0.97265, 0.98557) | 0 | 7 | n.a. |
| [0.98557, 0.99848) | 0 | 7 | n.a. |
| [0.99848, 1.01140) | 0 | 5 | n.a. |
| [1.01140, 1.02430) | 0 | 6 | n.a. |
| [1.02430, 1.03720) | 1 | 2 | 136 |

Source: Cartea et al (2008)