Current estimates of mean-reversion speeds for energy commodities like natural gas and power may be strongly biased toward zero. *Cliff Parsons* explains this bias, derives a formula for it in simpler cases, and gives evidence of it for estimates on US natural gas prices

Explaining bias in mean-reversion speed estimates for energy prices

★ In the natural gas and power industries, daily spot prices are presumed to be governed by mean-reversion. As argued in Mastrangelo (2007), we believe mean-reversion derives from daily supply and demand, which are either fairly constant or mean-reverting, respectively. And since storage in these industries is either limited or non-existent for carrying over supply to higher demand times, natural gas and power spot prices tend to mean-revert with demand.

Evidence of this mean-reversion exists: as natural gas and power prices typically trade within upper and lower bounds and volatility term-structures of their forward prices are downward sloping. The former indicates bounded variances over time, which is consistent with mean-reversion and the latter indicates that the effects of price shocks are expected to dissipate over time, which occurs with mean-reversion.

Estimating the speed of mean-reversion is important since the speed determines the strength of expected spot price trends, and, hence, the value of certain spot-price derivatives.[1]

However, previous speed estimates on US natural gas and power prices indicate speeds are usually slight and often insignificant. Ghazi & Sivothayan (2007) questions whether mean-reversion exists at all in current energy prices. We show that the usual estimates of mean-reversion for prices on commodities like natural gas and power can be biased heavily towards low and insignificant values by injudicious choices of estimation techniques and sampling.

Ordinary least squares (OLS) and maximum-likelihood are the standard techniques for estimating mean-reversion speeds and are usually applied to fit a simple first order autoregressive (AR(1)) model using large samples of spot prices on power, natural gas or oil/products. We explain that this approach tests not only for the hypothesis of mean-reversion, but also for the implicit (and apparently unintended) hypothesis that only a constant long-run mean governs the spot process. We show that a violation of the implicit hypothesis results in low and usually insignificant mean-reversion speed estimates, even when true mean-reversion is extremely high. Specifically, a sample in which the long-run mean truly varies may imply that no strong reversion exists towards the one constant mean implied from the estimation's intercept term. And since natural gas and power prices are either seasonal or appear to possess stochastic long-run means, the hypothesis of a constant long-run mean is usually violated in samples.

We contribute to the literature by characterising the nature of and deriving a formula for this bias in simpler cases; we also give a stark example of it.

Previous studies of mean reversion

In previous studies, Pilipovic (1998), Clewlow & Strickland (2000) and Eydeland & Wolyniec (2003) find low to moderate mean-reversion speed estimates on large samples of various energy commodities. Pilipovic (1998) introduces a harmonic function in her estimation to capture the seasonality of different energy commodities; however, her speed estimates are still quite low. She estimates mean-reversion from short-term forward prices (not spot prices), which one may argue have no mean-reversion in expectation, and we are unsure if her harmonic function correctly captures the true seasonality in her samples. Ghazi & Sivothayan (2007) documents findings of no significant mean-reversion in large samples on crude oil and natural gas prices. However, they do document significant mean-reversion in these commodities over sub-samples in which price levels are fairly stable: that is, the long-run mean appears to vary little within the sub-samples.

1. One such derivative is a natural gas storage lease. The holder of such a lease may buy and sell natural gas to inject into or withdraw from storage, respectively. By buying and selling at daily prices, the lease-holder attempts to profit from daily (mean-reverting) price trends.

Three studies that find high mean-reversion are Knittel & Roberts (2001), Benth & Benth (2004) and Pilipovic (2007). Knittel & Roberts (2001) finds high and significant estimates of mean-reversion speed in hourly power prices; however, we are unsure of how the samples in this study do or do not violate the implicit hypothesis mentioned above.

In Benth & Benth (2004), the authors find high and significant mean-reversion in natural gas spot prices, but not for Brent crude oil prices. In their estimations, the authors use a harmonic function to account for the varying long-run mean in their sample of natural gas prices, thus avoiding violating the implicit hypothesis. However, they do not fit such a harmonic to their oil prices, and the graph of their oil price sample appears to contain a varying long-run mean, perhaps leading to their low mean-reversion speed estimates. Pilipovic (2007) de-seasonalises spot prices for various energy commodities and finds high mean-reversion for power prices (which contrasts from her previous work), but not for natural gas prices.

In the past, we have used a different estimation method, one that controls for a stochastic long-run mean within samples, and have found evidence that the mean-reversion speed in US natural gas prices is very high: over 70. For that method, historical daily and forward prices were obtained and used in a two-factor tree model (one of the two factors shocks the long-run mean) to calculate expected trading values and simulate actual trading values for certain spot-price dependent derivatives. The mean-reversion speed and other estimated parameters were chosen to best fit the actual-to-expected trading values. Detailing this method further is beyond this paper's scope.

A formula for bias in mean-reversion speed estimates

We now recreate the standard model used to estimate mean-reversion speeds. Afterwards, we assume a price sample containing a simple time-varying long-run mean and derive a formula for OLS estimation bias based on it. To develop the standard model, assume the following for the spot-price process:

$$\frac{dS_t}{S_t} = a\left(\ln(L) - \ln(S_t)\right)dt + \sigma dz_t \qquad (1)$$

where

- S_t = the spot price for delivery at time t
- L = the process long-run mean
- a = the process mean-reversion speed
- σ = the process instantaneous volatility
- z_t = a Brownian motion

Converting to log-prices and solving gives the standard model used in estimating mean-reversion:

$$\ln(S_{t+\Delta t}) = \alpha + \beta \ln(S_t) + \varepsilon_{t+\Delta t} \qquad (2)$$

where:

$$\alpha = (1-\beta)\left(\ln(L) - \frac{\sigma^2}{2a}\right)$$

$$\beta = e^{-a\Delta t}$$

$$\varepsilon_{t+\Delta t} \sim \text{independent, } N\left(0, \sigma^2 \frac{1-e^{-2a\Delta t}}{2a}\right)$$

The model in equation (2) appears in such literature as Pilipovic (1998), Clewlow & Strickland (2000), Knittel & Roberts (2001), Eydeland & Wolyniec (2003), and Benth & Benth (2004). If we have a sample of daily spot prices, then we let $\Delta t = 1/365$ (one day expressed in years) and commence with estimating α and β. The estimate of β is then inverted to obtain the estimate for the mean-reversion speed, a. However, the underlying hypothesis tested in equation (2) is not obvious and bears on the before-mentioned bias. The true hypothesis tested is a conjunction of two hypotheses: Does mean-reversion govern the process, and does the process have a constant long-run mean? The former hypothesis is explicit while the latter is implicit.

We suspect that testing the latter hypothesis is unintended by most authors. Rejection of either hypothesis leads to low and possibly insignificant mean-reversion speed estimates, and rejection of only the latter, implicit hypothesis leads to the bias. For example, if samples include periods in which spot prices revert around higher means at times and lower means at other times, then mean-reversion speed estimates will be low since spot prices do not appear to be strongly reverting around the one mean embedded in the estimate of α. We now derive a formula for this bias under the assumptions that our sample is governed by one long-run mean for the first several observations, by a different long-run mean for the remaining observations, and that OLS estimation is employed.

We assume a sample of daily spot prices with size $n = n_1 + n_2$, where the process for the first $n_1 > 0$ prices is governed by long-run mean L_1, while the process for the other $n_2 > 0$ prices is governed by long-run mean L_2. We also assume that $L_1 < L_2$ and that the processes for both sub-samples are governed by the same mean-reversion speed, a, and volatility, σ. This setup[2] and equations (1) and (2) imply that two intercepts exist in reality: α_1 corresponding to L_1 and α_2 corresponding to L_2. If we do not account for the time-varying long-run mean in our sample, then our OLS estimate of β (which by (2) is inverted to obtain the estimate of the mean-reversion speed) is as follows:

$$\hat{\beta} = \frac{\sum_{i=1}^{n}\left(Y_i - \bar{Y}_0\right)\left(Y_{i-1} - \bar{Y}_{-1}\right)}{\sum_{i=1}^{n}\left(Y_{i-1} - \bar{Y}_{-1}\right)^2} \qquad (3)$$

2. Extending this setup to three or more long-run means becomes apparent as we proceed.

where:

$$\hat{\beta} = \text{the OLS estimate of } \beta$$

$$Y_i = \ln(S_i),\ i = 0, \ldots, n$$

$$\bar{Y}_0 = \frac{\sum_{i=1}^{n} Y_i}{n}$$

$$\bar{Y}_{-1} = \frac{\sum_{i=0}^{n-1} Y_i}{n}$$

Using equation (2) to substitute for Y_i and $\bar{Y}_0$, the right-hand side of equation (3) becomes:

$$\hat{\beta} = \frac{\left(\begin{array}{l} \sum_{i=1}^{n_1}\left(\alpha_1 + \beta Y_{i-1} + \varepsilon_i - \bar{\alpha} - \beta\bar{Y}_{-1} - \bar{\varepsilon}\right)\left(Y_{i-1} - \bar{Y}_{-1}\right) \\ + \sum_{i=n_1+1}^{n}\left(\alpha_2 + \beta Y_{i-1} + \varepsilon_i - \bar{\alpha} - \beta\bar{Y}_{-1} - \bar{\varepsilon}\right)\left(Y_{i-1} - \bar{Y}_{-1}\right) \end{array} \right)}{\left(\sum_{i=1}^{n_1}\left(Y_{i-1} - \bar{Y}_{-1}\right)^2 + \sum_{i=n_1+1}^{n}\left(Y_{i-1} - \bar{Y}_{-1}\right)^2 \right)} \quad (4)$$

where:

$$\bar{\varepsilon} = \sum_{i=1}^{n} \varepsilon_i / n$$

$$\bar{\alpha} = \frac{n_1}{n}\alpha_1 + \frac{n_2}{n}\alpha_2$$

Combining terms and taking expectations of both sides of equation (4) gives:

$$E\left(\hat{\beta}\right) = \beta + E\left(\frac{\sum_{i=1}^{n_1}\left(\alpha_1 - \bar{\alpha}\right)\left(Y_{i-1} - \bar{Y}_{-1}\right) + \sum_{i=n_1+1}^{n}\left(\alpha_2 - \bar{\alpha}\right)\left(Y_{i-1} - \bar{Y}_{-1}\right)}{\sum_{i=1}^{n_1}\left(Y_{i-1} - \bar{Y}_{-1}\right)^2 + \sum_{i=n_1+1}^{n}\left(Y_{i-1} - \bar{Y}_{-1}\right)^2} \right) + E\left(\frac{\sum_{i=1}^{n}\left(\varepsilon_i - \bar{\varepsilon}\right)\left(Y_{i-1} - \bar{Y}_{-1}\right)}{\sum_{i=1}^{n_1}\left(Y_{i-1} - \bar{Y}_{-1}\right)^2 + \sum_{i=n_1+1}^{n}\left(Y_{i-1} - \bar{Y}_{-1}\right)^2} \right) \quad (5)$$

The third term on the right-hand side of equation (5) is insignificant; if the second term were also insignificant, then $\hat{\beta}$ would be a fairly unbiased estimator of β. However, in this case it is significant, and we can see this by examining its numerator. The first sum in the numerator tends to be positive since the products being summed are typically both negative: α_1 is less than $\bar{\alpha}$ by assumption, and Y_{i-1} tends to be less than $\bar{Y}_{-1}$ over the first n_1 observations. By similar reasoning, one sees that the two terms in the products of the numerator's second sum are typically both positive. And since the denominator is always positive, the whole second term on the right-hand side of equation (5) tends to be positive under our assumptions.

The amount of bias can be quite large, and we can see this by transforming equation (5) a little further. Extrapolating from equation (2), we see that:

$$\alpha_i = (1-\beta)\left(\ln(L_i) - \frac{\sigma^2}{2a}\right) \text{ for } i = 1, 2.$$

Let $U_i = \ln(L_i),\ i = 1, 2.$

Then $\bar{U} = \left(\frac{n_1}{n}U_1 + \frac{n_2}{n}U_2\right)$ and

$$\left(\alpha_i - \bar{\alpha}\right) = (1-\beta)\left(U_i - \bar{U}\right) \text{ for } i = 1, 2.$$

Substituting these expressions into the second term of equation (5) gives the following:

$$E\left(\hat{\beta}\right) = \beta + (1-\beta)E\left(\frac{\left(\begin{array}{l} \sum_{i=1}^{n_1}\left(U_1 - \bar{U}\right)\left(Y_{i-1} - \bar{Y}_{-1}\right) \\ + \sum_{i=n_1+1}^{n}\left(U_2 - \bar{U}\right)\left(Y_{i-1} - \bar{Y}_{-1}\right) \end{array} \right)}{\left(\sum_{i=1}^{n_1}\left(Y_{i-1} - \bar{Y}_{-1}\right)^2 + \sum_{i=n_1+1}^{n}\left(Y_{i-1} - \bar{Y}_{-1}\right)^2 \right)} \right) \quad (6)$$

Finally, if the deviations of U from its mean are approximately equal to the deviations of Y from its mean over each of the two sums in the numerator, then the expectation on the right-hand side of (6) is approximately 1.0, and the estimate of β is approximately:

$$E\left(\hat{\beta}\right) \approx \beta + (1-\beta) = 1 \quad (7)$$

For example, the bias approaches $(1-\beta)$ for a sample having prices propagated with a large, true mean-reversion speed, coupled with large differences in long-run means between sub-samples of the sample, and where prices propagated with each mean form a significant proportion of the sample. Such a scenario makes the U and Y_i deviations close. We give an example of this in the next section.

The approximation in (7) shows that higher, true mean-reversion speeds, which have true betas both positive and much lower than 1.0, may be counteracted with greater bias, making the bias especially insidious. This bias, in which β is truly less than one while $\hat{\beta}$ is biased upward towards one, is opposite of the unit-root bias.[3] The above derivations confirm our suspicions that, even if true mean-reversion speed is very high, insignificant mean reversion speed estimates may result from violating the before-mentioned implicit hypothesis.

3. The unit-root bias occurs for processes like equation (2) where β is close to one and least-squares is employed: The estimate $\hat{\beta}$ is biased downward. Greene (1993) summarises this bias.

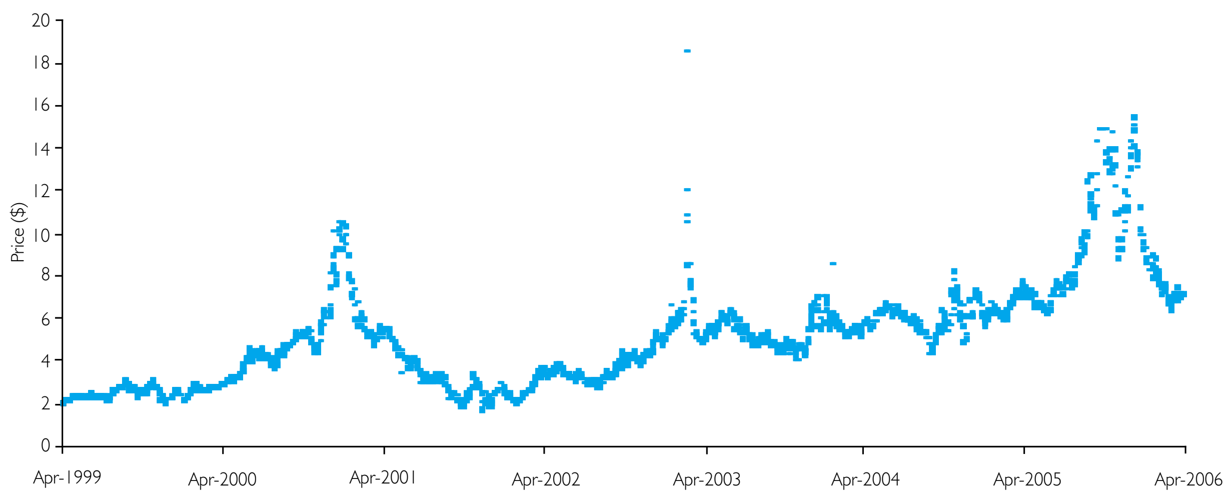

F1. Daily delivery: Henry Hub

Prices for daily delivery at Henry Hub, Louisiana USA, for the period April 1, 1999 to April 3, 2006

Source: Norman's Historical Data

Examples of bias in mean-reversion speed estimates

Spot price processes for energy commodities like natural gas and power appear to have long-run means that blend seasonality with stochastic behaviour. Figure 1 illustrates this for prices on daily-delivery US natural gas traded at Henry Hub, Louisiana. These prices possess volatilities that can easily exceed 80% per year; however, these prices tend to revert towards a long-run mean level that wanders within a band of prices. Seasonality tends to push this level higher in winter than in summer due to US demand for natural gas being higher in winter. The data plot in figure 1 shows that the long-run mean for US natural gas is more complicated than what is assumed in the usual estimation models, which can lead to the bias explained previously. We give evidence of this bias later in this section when using US natural gas prices, but for now, we illustrate the large bias formulated in (7) with a simple example.

We generated 200 independently and normally distributed ε terms under the assumptions that $a = 365$, $\sigma = 50\%$, and $\Delta t = 1/365$ (one-day). Using those terms and assuming $L = \$5.00$ for the first hundred observations, $L = \$8.00$ for the second hundred observations, and $S_0 = \$4.80$, we generated 200 daily spot prices. Thus, $n_1 = n_2 = 100$. Our chosen mean-reversion speed is high by most measures to illustrate the bias more clearly, but the price level, volatility, and long-run mean spread are realistic regarding US natural gas spot prices as seen in figure 1.

Next, we performed three OLS regressions: one on the first half of the sample, one on the second half, and one on the whole sample. The results for the estimates of a ($\hat{a}$), given in table 1, show that although strong mean-reversion is present, the presence of a greatly varying long-run mean in the sample can bias the mean-reversion speed estimate strongly towards zero, as our derivations leading to the approximation in (7) show.

We examine figure 2 for intuition about this example. This figure shows a hypothetical sample in log-price space from a price process that varies around a lower long-run mean part of the time, and a higher long-run mean the rest of the time. Such a process would produce two clusters of outcomes, similar to those shown in the figure, in which each cluster tends to lie around a corresponding regression line. The process would also cause each cluster to center around the 45-degree line from the origin where both the x

T1. An example of OLS mean-reversion speed bias

Source: based on data generated by the author for illustration purposes

Sample	$\hat{a}$	Null: $a = 0$
First half	270.68	reject
Second half	354.70	reject
Whole sample	3.02	do not reject

Null hypothesis was tested at the 95% level. Prices were propagated with a mean-reversion speed of 365.0, a long-run mean of \$5.00 for the first sample half, and a long-run mean of \$8.00 for the second sample half. The whole-sample estimate illustrates the downward bias caused by a varying long-run mean within the sample

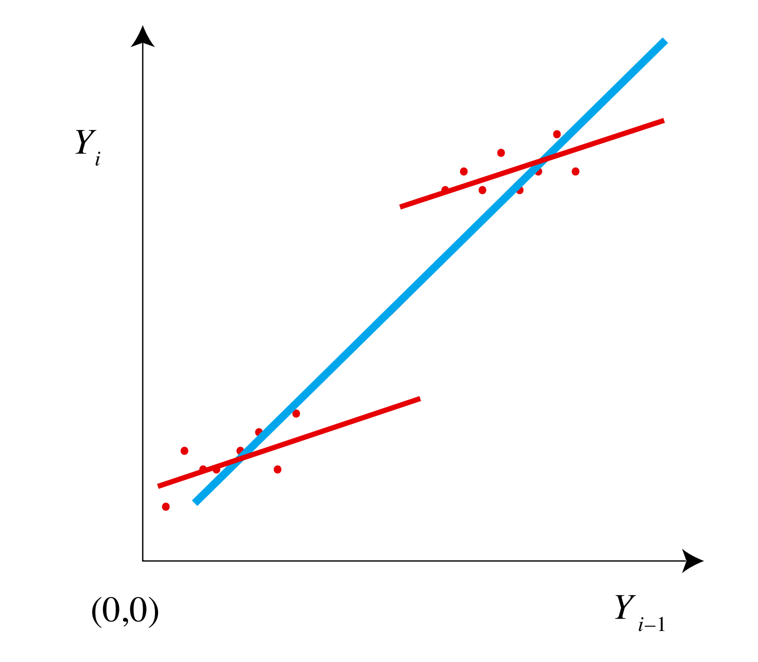

F2. Example of price dispersion caused by a varying long-run mean
The graph illustrates a sample where mean-reverting outcomes are governed by one of two long-run means at any one time. Each of the two thin lines fits outcomes associated with their respective long-run mean. The thick line is a 45-degree line from the origin and would closely coincide with an OLS fit to the whole sample *Source: based on data generated by the author for illustration purposes*

indicates that previous outcomes and their subsequent outcomes tend to remain on the same side of the corresponding long-run mean; the slope being less than one indicates that subsequent outcomes are closer to the corresponding long-run mean than previous outcomes in expectation. The formula for equation (2) implies these relationships since the slope coefficient essentially weights the previous log-price and log-long-run mean for forecasting the next log-price.

A simple OLS regression on such a sample tends to fit the sample's two clusters of outcomes with a line close to the 45-degree line from the origin. Thus, $\hat{\beta}$ for this fit is close to one, which implies a mean-reversion speed estimate of close to zero. This result is what our derivations leading to the approximation in (7) show.

We now apply this intuition to actual data. US natural gas prices indicate the presence of a varying long-run mean – for example, the seasonal winter/summer price spread, which can cause multiple clusters of outcomes in large samples. The graph of these outcomes in the log-price space of figure 2 typically shows a 'cloud' of points along the 45-degree line from the origin. Therefore, when testing for mean-reversion, we must sample in a way that controls for a varying long-run mean.

and *y* co-ordinates approximately equal the corresponding log-long-run mean. The positive slope of each cluster's regression line is consistent with the price process and

We choose a simple approach: Use small samples such as 21 days. By using such small samples, we hope the long-run mean within the samples varies little.

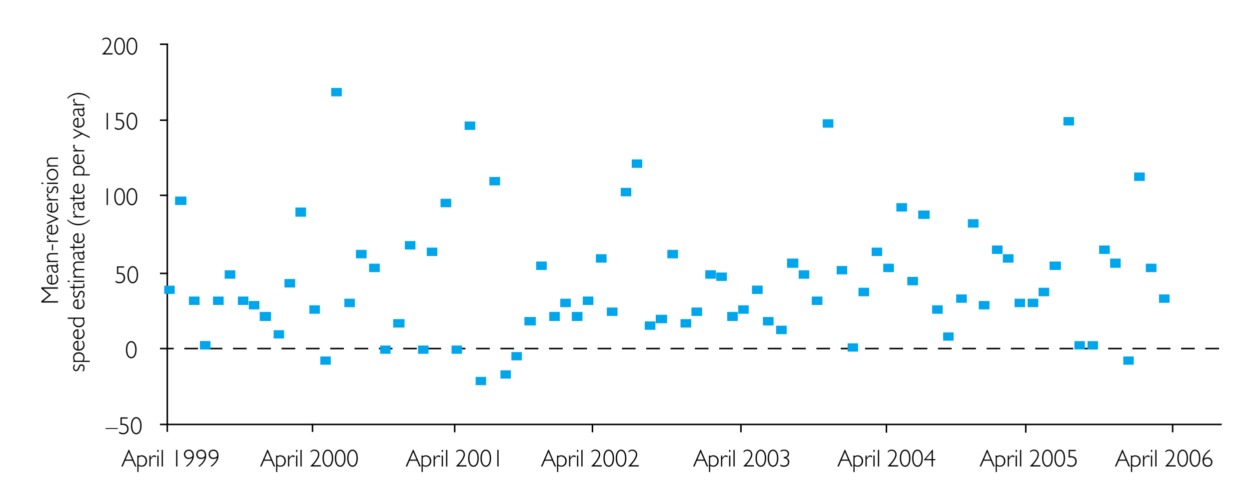

F3. Speed estimates
Mean-reversion speed estimates by month from 1999 to 2006 *Source: estimated using Norman's Historical Data*

The price data are for Henry Hub, span 1999 to 2006, and are gas-daily prices.[4] Two arbitrary samples of size 21 were taken for each of the seven years: one sample from the summer, and one from the winter. Of the 14 mean-reversion speed estimates, eight were over 90.0, three of these were significant at the 95% level, and seven of these were significant at the 90% level.[5] Finding such significant results on such small samples is noteworthy, especially considering how 'noisy' small samples on natural gas prices are. Only three of the 14 estimates were less than 50.0, and these were statistically insignificant. Alternatively, the whole sample of 1,485 prices produced a mean-reversion speed estimate of only 3.40, which was just significant at the 95% level and was much lower than our estimates derived from smaller samples. Such behaviour is consistent with assumptions underlying the derivation of (7) and suggests that US natural gas mean-reversion speeds could be much higher than previously documented.

Finally, and as an ancillary exercise, we examined our estimates for seasonality. We estimated mean-reversion speeds per month using prices for each calendar month in our seven-year sample. The results are plotted in figure 3 and show that the estimates do not appear to have any obvious seasonality. Also, and consistent with the results just above, 72 of the 84 monthly estimates were higher than the whole-sample estimate of 3.40, with the average of the 84 estimates being 43.5.

Conclusion

We have shown that, by using standard estimation techniques, mean-reversion speed estimates can be biased greatly towards zero if samples include prices propagated with a varying long-run mean. Such samples for natural gas and power prices are common since these prices are typically seasonal, and, at the very least, appear to have stochastic long-run means. We characterise and derive a formula for this bias under simple conditions.

Sampling and estimation techniques must change when estimating mean-reversion on energy prices that are influenced by seasonality or varying long-run means. We suggest, as does Ghazi & Sivothayan (2007), the employment of two-factor mean-reverting models for testing, at the very least. And although our research gives some direction on how current techniques should change, we leave such innovations to further research.

Acknowledgements

The author especially thanks Charles Riedhauser and Thanassis Tjavaras for many helpful comments on this paper. The author also thanks Rafael Mendible, Chris Richards, Wei Zhang and Dayne Zimmerman for discussions and comments regarding this paper. ER

Cliff Parsons, is president of WTM Energy Software; 6422 Cottonwood Park Lane, Houston, TX 77041, USA. www.wtmenergy.com

Emails: Inquire@WTMenergy.com; cliff.parsons@sbcglobal.net

References

Benth F and J Benth, 2004
The normal inverse gaussian distribution and spot price modelling in energy markets
International Journal of Theoretical and Applied Finance, Vol. 7, No. 2, pp 177–192

Clewlow L and C Strickland, 2000
Energy derivatives: pricing and risk management
Lacima Publications

Eydeland A and K Wolyniec, 2003
Energy and power risk management
John Wiley & Sons

Ghazi D and S Sivothayan, 2007
Has mean reversion had its day?
Energy Risk, April 2007, pp 50–53

Greene W, 1993
Econometric analysis
2nd edition, Macmillan, chapter 19

Knittel C and M Roberts, 2001
An empirical examination of deregulated electricity prices
Working paper, School of Management, Boston University, October

Mastrangelo E, 2007
An analysis of price volatility in natural gas markets
US Energy Information Administration, August

Pilipovic D, 1998
Energy risk: valuing and managing energy derivatives
McGraw-Hill

Pilipovic D, 2007
Energy risk: valuing and managing energy derivatives
2nd edition, McGraw-Hill

4. All prices used in this paper come from Norman's Historical Data: http://www.normanshistoricaldata.com/

5. The caveat in estimating parameters and measuring significance levels using OLS on smaller samples is that errors should be distributed normally, independently, and identically through time in order to achieve better results.