Correlation and Cointegration in Energy Markets

Professor Carol Alexander

ISMA centre, Reading University

www.ismacentre.reading.ac.uk

Successful risk management requires real understanding of the nature of volatility and correlations between financial markets, and the problems inherent in calculating statistical estimates of these quantities. Whilst volatilities are based on the variances of individual returns distributions, correlations depend on characteristics of the joint distributions between two related markets. This extra dimension adds a great deal of uncertainty to correlation risk measures. In fact, whilst it seems reasonable to assume that individual return processes are stationary, so that volatilities do exist, it is by no means always the case that two returns processes will be jointly stationary. So unconditional correlations may not even exist.

Of course it is always possible to calculate a number that supposedly represents correlation, but often these numbers change considerably from day to day, a sign that the two returns processes are not jointly stationary. It is unfortunate that some standard correlation estimation methods induce an apparent stability that is purely an artefact of the method, and the true nature of underlying correlations is obscured.

The first objective of this chapter is to review the different approaches to measuring energy correlations, pointing out their advantages and limitations. Although correlation has become the ubiquitous tool for measuring comovements in asset returns, its limitations are substantial.

Firstly, data are detrended before the analysis and this precludes any possibility to investigate long-term common trends in asset prices. Correlation based hedges may

require frequent rebalancing because there is nothing in the computation of hedge ratios to guarantee that the hedge is tied to the underlying over the longer term. Secondly correlation, being essentially a static measure, cannot reveal any dynamic causal relationships. Indeed it is possible that hedges could be based on spurious correlations, that appear high even though there is no underlying causal relationship, particularly when some of the rather misleading correlation measures are employed.

The advanced management of correlation risks should take account of any lead-lag relationships, for example the 'price discovery' between spot and futures, but correlation is not an adequate tool for such analysis.

Applications of correlation to energy markets include computation of hedge ratios, the pricing of multi-asset options, and the risks of trading spreads and commodity baskets. These applications are discussed with reference to the crude oil and natural gas markets. We use daily data on WTI crude oil spot and near futures prices from 1st July 1988 to 26th February 1999, and the NYMEX sweet crude prices from 1 to 12 months from 4th February 1993 to 24th March 1999. For the natural gas market we use the NYMEX prompt month future with Gas Daily time series for natural gas from 10th April 1992, and the Kansas City "Western" natural gas contract from 18th December 1995, to 3rd March 1999.

The last part of this chapter introduces a method of measuring comovements between markets that overcomes some of the limitations of correlation. Cointegration, a methodology that has become standard practice in econometrics during the last decade, is now showing itself to be a very useful tool for hedging financial assets.

Cointegration refers not to comovements in returns, but comovements in asset prices. If spreads are mean-reverting, asset prices are tied together in the long-term by a common stochastic trend. Hedges based on cointegration may deviate from the

underlying in the short-term, but are tied to the underlying by a long-run equilbrium relationship and hence require less frequent rebalancing.

Cointegration is a two step process: first any long-run equilibrium relationships between prices are established, and then a dynamic correlation model of returns is estimated. This 'error correction model', so-called because short-term deviations from equilibrium are corrected, reveals the dynamic causalities that must be present in a cointegrated system. The chapter concludes with empirical testing of a cointegration model for crude oil spot and futures.

Statistical Measures of Correlation

Returns to financial assets, the relative price changes, are well approximated for short holding periods by the difference in log prices. For ease of exposition we consider daily returns $R_t = \log P_t - \log P_{t-1}$ although correlation methods apply equally well to different returns frequencies, provided data are synchronous.

It is reasonable to assume that returns are generated by a stationary stochastic process. That is:

$E(R_t)$, the unconditional mean, is a finite constant;

$V(R_t)$, the unconditional variance, is a finite constant ;

$COV(R_t, R_{t-s})$, the unconditional autocovariance, depends only on the lag s.

The mean-reversion property of stationary series is well known. A stationary process is mean-reverting, not in the sense of a mean-reverting term structure, but mean-reverting over time. They can never drift too far from their mean because of the finite variance.

The speed of mean-reversion is determined by the autocovariance: mean-reversion is quick when autocovariances are small, and slow when autocovariances are large. At one end of the spectrum we have 'white noise', when returns are independent so $COV(R_t, R_{t-s})$ is zero, and mean-reversion is instantaneous. At the other extreme $COV(R_t, R_{t-s}) = V(R_t)$ so autocorrelations are unity, there is no mean-reversion, and returns are not stationary.

Two stationary returns processes R_1 and R_2 are jointly covariance stationary if $COV(R_{1,t}, R_{2,t-s})$ depends only on the lag s. In particular the contemporaneous covariance $COV(R_{1,t}, R_{2,t})$ is a constant, irrespective of the time at which it is measured.

For jointly stationary returns we may define a contemporaneous cross correlation as

$$Corr(R_{1,t}, R_{2,t}) = COV(R_{1,t}, R_{2,t}) / \sqrt{(V(R_{1,t}) V(R_{2,t}))} \qquad (1)$$

Or in alternative notation $\rho = \sigma_{12} / \sigma_1 \sigma_2$. So the extension of the constant volatility assumption explained in chapter * (ref to chapter on volatility) to constant unconditional cross correlations requires joint stationarity.

This is quite an heroic assumption, except in special circumstances. So it should be clear from the outset that the computation of unconditional correlations may be a meaningless exercise. Nevertheless it is standard practice, so one focus of this chapter is to point out the dangers of using such measures when they are not, in fact, valid.

If it exists the unconditional correlation is one number, ρ, that is the same throughout the process. Correlations always lie between +1 and –1. High positive values indicate that the returns move together in the same direction, and high negative values indicate that they tend to move in opposite directions. Orthogonal or uncorrelated returns have zero correlation.

Any differences between estimates of ρ at different times arise from differences in samples. The smaller the sample the bigger these differences, because sampling errors are inversely proportional to sample size. But when returns have a high degree of joint stationarity correlation estimates should not jump around too much even for small sample sizes. On the other hand if correlation estimates are highly unstable this is a sure sign of non-joint stationarity. So, whilst it is always possible to calculate a number based on some formula or model for correlation, it does not always make sense to do so.

It may also be that correlations appear quite high for a long period, even when they are spurious. For example contemporaneous data on live hog spot prices and crude oil prompt futures may be available from NYMEX, and correlations estimates could be calculated. But it is probable that little underlying causal relation exists between live hogs and crude oil, except perhaps in transportation costs. Their returns are unlikely to be jointly stationary, but correlation calculations according to some methods might result in apparently high and stable correlations.

The next section shows how the most common correlation estimates of all, the equally weighted 'historic' correlations, will have apparent stabilities that are, in fact, just an artefact of the estimation method. Using a more appropriate correlation model, such as an exponentially weighted average or a GARCH model, would reveal greater instabilities in correlation, particularly if returns are not jointly stationary.

But then, if unconditional correlations do not exist because returns are not jointly stationary, and if conditional correlations are jumping around all over the place, what can be done to hedge correlation risk? In the absence of a correlation futures contract based on equally weighted averages it may be better to look for alternative measures of comovements between assets. For example, it would be possible to base hedging

strategies on the cointegration error correction models of asset prices that are introduced at the end of this chapter.

EQUALLY WEIGHTED MOVING AVERAGES

Unbiased estimates of unconditional correlation are usually calculated by a weighted moving average, with either equal or exponential weighting. A standard method based on (1) is to estimate variance as a weighted average of squared returns and to divide the covariance, estimated as a similarly weighted average of cross products of returns, by the square root of the product of the variances.

Consider for example the WTI crude oil spot[1] and NYMEX near futures prices shown in figure 1. They should be very highly correlated because the major price changes arise from supply constraints, such as the Gulf crisis in 1990/1991, rather than demand fluctuations. This is in marked contrast to the natural gas market that is discussed in box 1.

However, crude oil spot-futures correlations are also quite variable, because the market has oscillated between backwardation and contango during the course of the decade (see box 2). Fluctuations in the convenience yield arise as perceptions of inventory and financing costs change, and these perceptions are governed by micro and macro economic factors that can vary considerably between the different players in the oil market (refer to excellent discussion in chapter on forward curve).

---

[1] Spot is the 1st Month Cushing until the future's expiry, then the 2nd Month Cushing until the 25th, then back to the 1st Month Cushing

Figure 1: Daily Average Spot and Future Prices for WTI Crude Oil

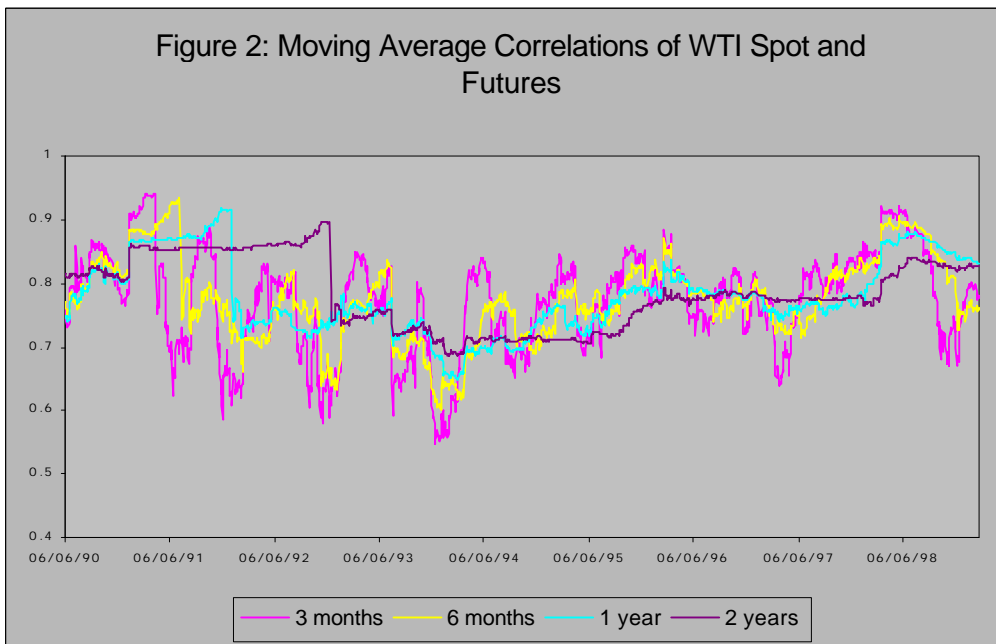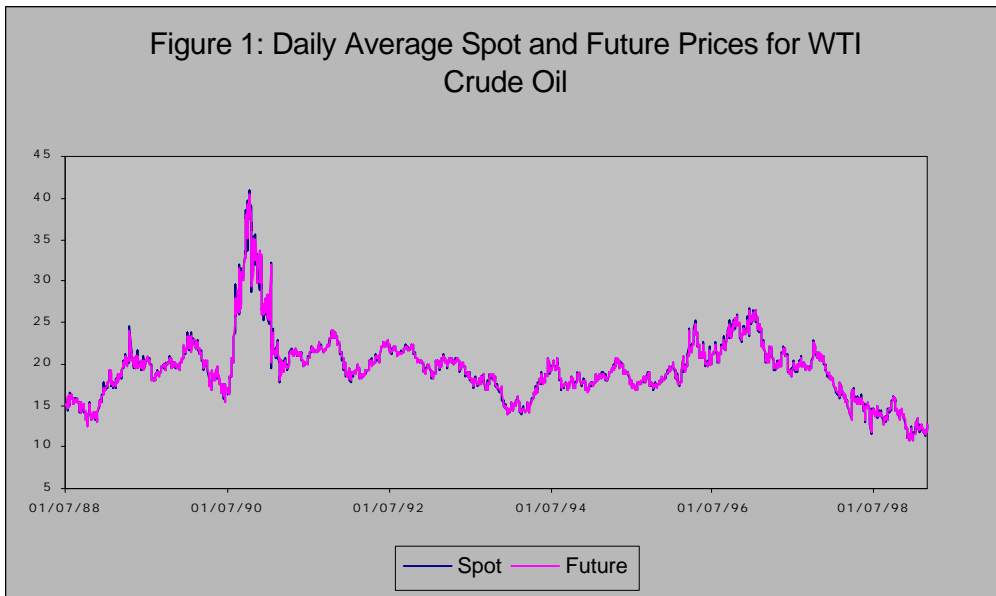Figure 2: Moving Average Correlations of WTI Spot and Futures

Figure 2 shows correlations between spot and futures crude oil prices calculated using equal weighting over 3 months, 6 months, 1 year and 2 years. The longer the averaging period, the more stable correlations appear to be. This is because the

pronounced effect on correlations that always follows an extreme event in the markets will last for exactly n-days, where n is the length of the averaging period.

For example on 17th January 1991 when spot and future prices dropped from about 32$ to about 22$ overnight with the outbreak of war in the Gulf, equally weighted correlations increased substantially by an amount in inverse proportion to the length of average. On 18th January 1991 the 3 month correlation rose from 0.8 to 0.91, staying above 0.9 until 17th April 1991 when it jumped down from 0.94 to 0.83. The 2 year correlation jumped from 0.81 to 0.86, staying at around this level for exactly 2 years, long after the other averages had returned to more realistic levels. But nothing special happened on 17th April 1991, or on 18th July 1991 or 1 year or 2 years after the outbreak of the Gulf war. The sharp declines in correlation measures on these dates are just an artefact of the estimation method.

These 'ghost' effects of extreme events on correlation are less intense but longer lasting, as the averaging period increases. So if equal weighting is to be applied for measuring correlations and hedge ratios there is a case for ignoring extreme events. Otherwise they can bias estimates for a long time after the event occurred.

EXPONENTIALLY WEIGHTED MOVING AVERAGES

One of the advantages of using exponential rather than equal weighting is that shocks to correlation die out exponentially, at a rate determined by the smoothing constant. Exponential smoothing takes the form

$$(1-\lambda)\sum_{i=1}^{\infty}\lambda^{i-1}\,x_{t-i-1}$$

where $0<\lambda<1$. Another time series is created, the exponentially weighted moving average (EWMA) that is 'smoother' than the original. The degree of smoothing is determined by the size of the smoothing constant λ.

To calculate an exponentially weighted correlation, take three EWMAs with the same value of λ. First calculate each of the two returns variances by smoothing the squared returns in each market. Then calculate the EWMA covariance, so that x is the cross product of returns, and finally divide this by the square root of the product of the two variances.

Exponential weighting is a simple method of measuring correlations that has advantages over both the 'historic' equally weighted averages and the more technical GARCH models that are introduced in the next section. But the big question with exponential weighting is which value of λ should be used? The larger is λ, the smoother the correlation becomes because observations far in the past still effect the current average. The smaller is λ, the more responsive the correlation to daily moves in the markets.

There is no one best method for optimising the value of λ. The examples in this chapter all take the RiskMetrics λ value of 0.94. In fact the exponentially weighted moving average (EWMA) with $\lambda = 0.94$ has a half-life of about 30 days, so its variability is similar to that of the 30 day equally weighted moving average (see figure 3). The main difference between the two methods is that the equally weighted measure has 30 day 'ghost' effects, whereas shocks die out exponentially in the EWMA. It would therefore be more realistic to base spot-futures or forward curve arbitrage on exponentially rather than equally weighted correlation estimates.

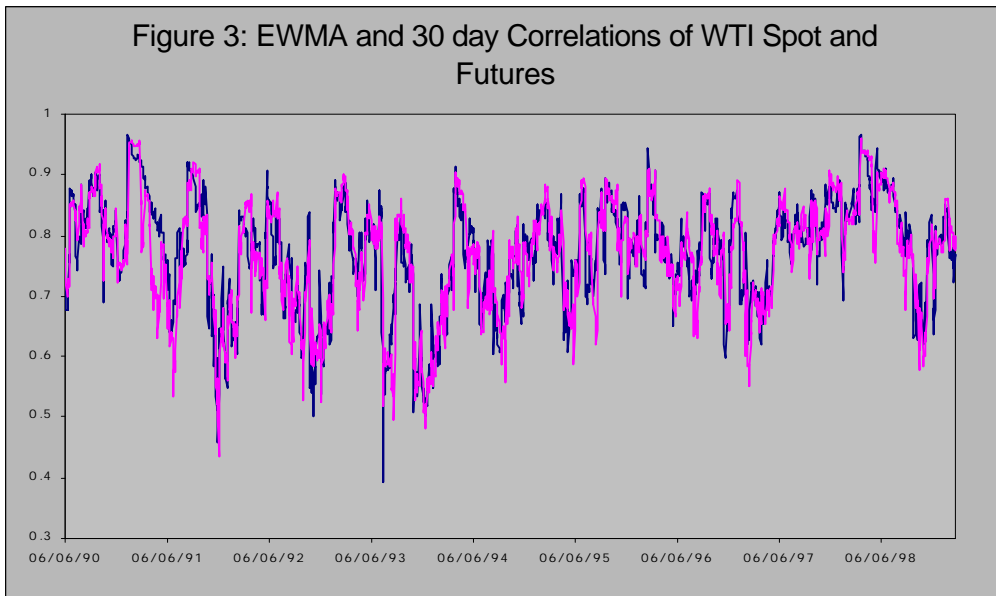Figure 3: EWMA and 30 day Correlations of WTI Spot and Futures

Figure 3 shows the exponentially weighted correlation between spot and prompt future prices of crude oil using the RiskMetrics smoothing constant $\lambda = 0.94$.[2] It is quite evident from figure 3 that these estimates are more unstable then those in figure 2. But in both figures it is the same correlation that is being estimated.

BOX 1: The Correlation between Natural Gas Spot and Prompt Future

Consider hedging natural gas spot prices as indicated by the Gas Daily index, by the NYMEX future for delivery over the next month, shown in figure 4.

---

[2] A review of the RiskMetrics data and alternative covariance matrices may be found in Alexander (1996).

Figure 4: Natural Gas Prices
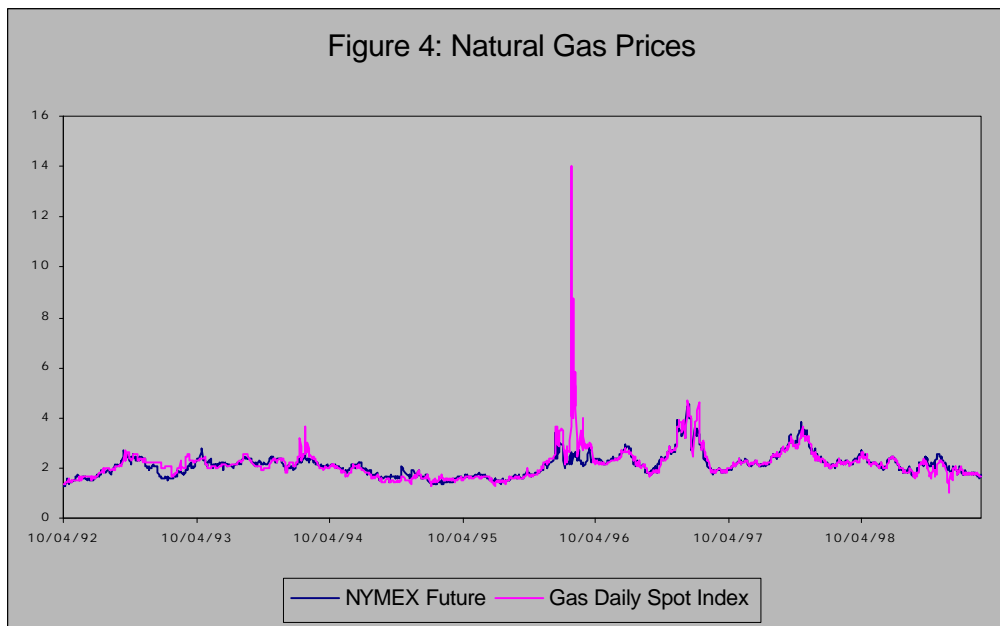
Natural gas storage facilities play a crucial role in balancing supply and demand in North America, and there are substantial seasonal effects. In the summer months excess production is injected into storage and in the winter months the storage gas is withdrawn to supply any excess load. In cold winters, when demand typically exceeds production, one would expect to see Gas Daily prices rise sharply during periods of extreme cold. Futures prices may also rise because depleting storage may raise future price expectations.

However, because storage costs are high, there can be a substantial decoupling of these two prices under the proper conditions. In some cases, Gas Daily prices may react much more strongly than the futures prices, as was seen when the gas daily prices spiked in February 1996. The net effect is that correlations between spot and futures are quite variable, and may be rather low when storage is filled to capacity.

The winters of 1997/98 and 1998/99 were very mild. During the most recent Autumn you could hardly give away spot gas during some weeks, since storage was completely full. This is apparent in the downward price spikes in Gas Daily prices last

11

October and November. With these two warm winters the exponentially weighted moving average (EWMA) correlation between spot and futures prices has decreased, as shown in figure 5. Recently it has become negative. Storage was filled to capacity during the Autumn months of 1997 and 1998, so futures prices have responded little to daily demand fluctuations.

Figure 5: EWMA Correlation in Natural Gas Spot and Futures
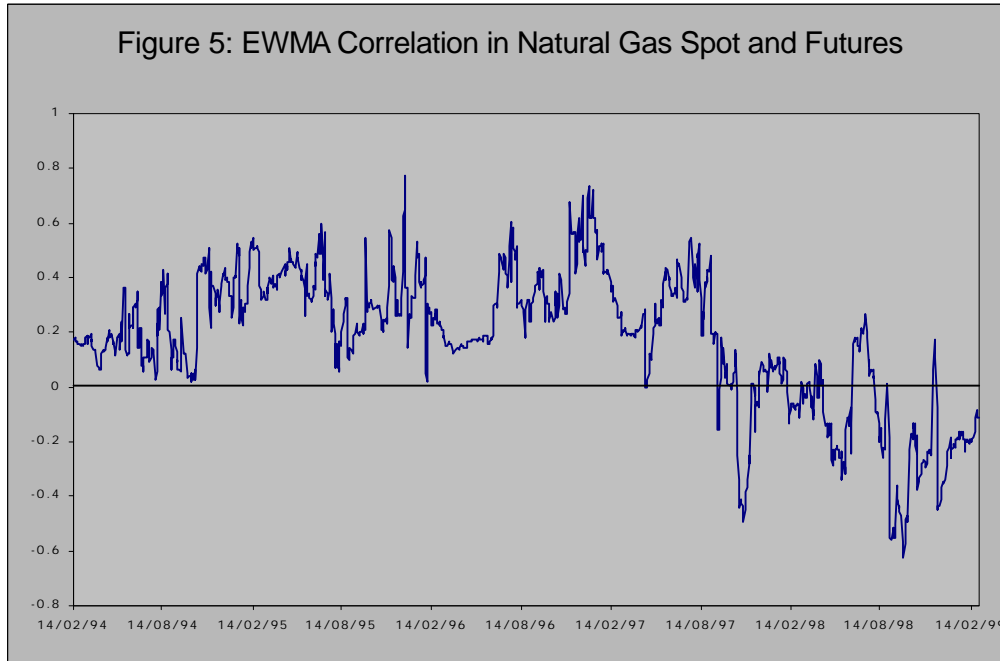
Exponentially and equally weighted moving average correlation measures are all estimating the same thing, the constant unconditional correlation that is assumed by the weighted average model. So how can it be that so many different results are obtained from estimating correlation on the same data? A generic problem with correlation estimation is that wildly different results may be obtained, depending on the time period of the data and the estimation method used.

For any given model, be it an equally weighted average over a fixed number of days, or an exponentially weighted average with some value for λ, the variation in results at

different time periods can only be attributed to 'noise' arising from differences in samples. There is nothing else in these models to explain time variation in correlation.

When correlation estimates are found to be unstable the constant correlation assumption that underlies weighted average correlation measures should be questioned. Conditions in many energy markets may not be conducive to joint stationarity. For example the correlations between NYMEX sweet crude futures of various maturities can be very unstable, particularly when there are big differences in trading volume and open interest (box 2).

Readers interested in statistical tests for joint stationary are referred to standard econometrics texts such as Hamilton (1994). They are based on the eigenvalues of coefficient matrices from a vector autoregressive model: if they lie within the unit circle the series are jointly covariance stationary.[3] But it should not really be necessary to go into such details: if correlations are found to be unstable, returns are unlikely to be jointly stationary. In this case it is not correct, strictly speaking, to use a model based on constant correlation, such as the weighted average models that have been discussed above.

Box 2: Correlation in the Crude Oil Term Structure

The NYMEX 'Sweet' crude prices from 1 to 24 months are illustrated in Figure 6. An immediate feature to note is that the volatility of futures prices decreases with maturity, the 1 month future being the most volatile and the 24 month future the least. This is true for all term structures, because short term expectations are more volatile than long term expectations and open interest and trading volume are concentrated at the low end. Slicing through the data gives the futures term structure on any given day, and both downwards and upwards sloping term structures are evident as the market oscillates between backwardation and contango.

Figure 6: NYMEX Sweet Crude Prices

---
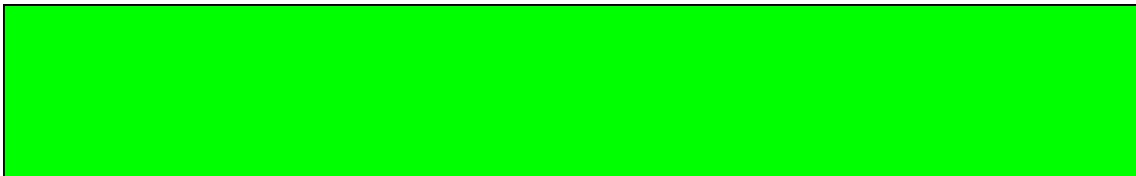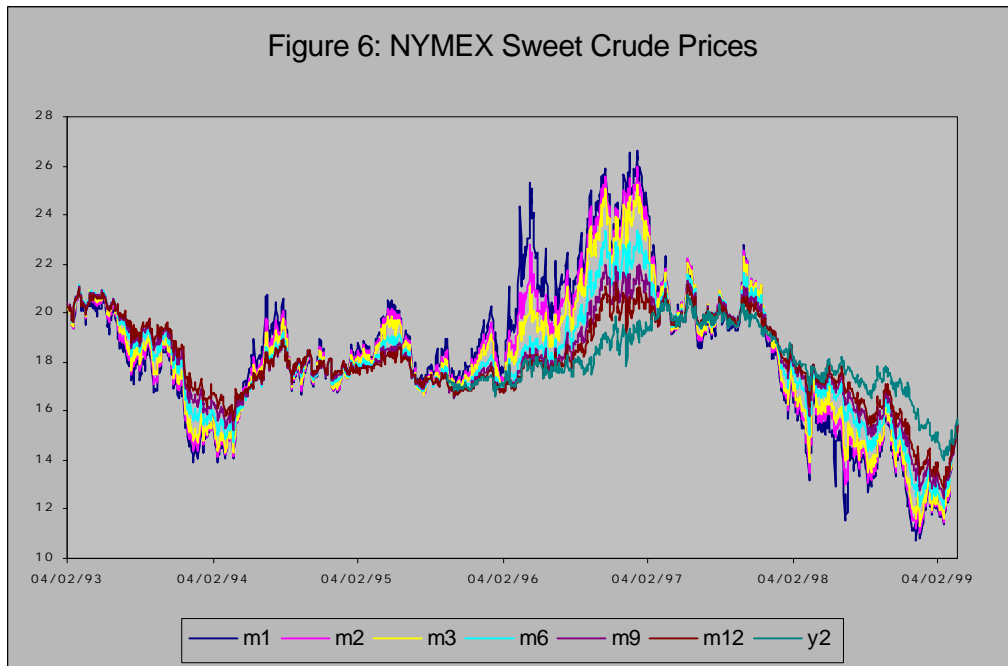
[3] Hence the term 'unit root' to indicate non-stationarity.

Figure 7: NYMEX Sweet Crude EWMA Correlations
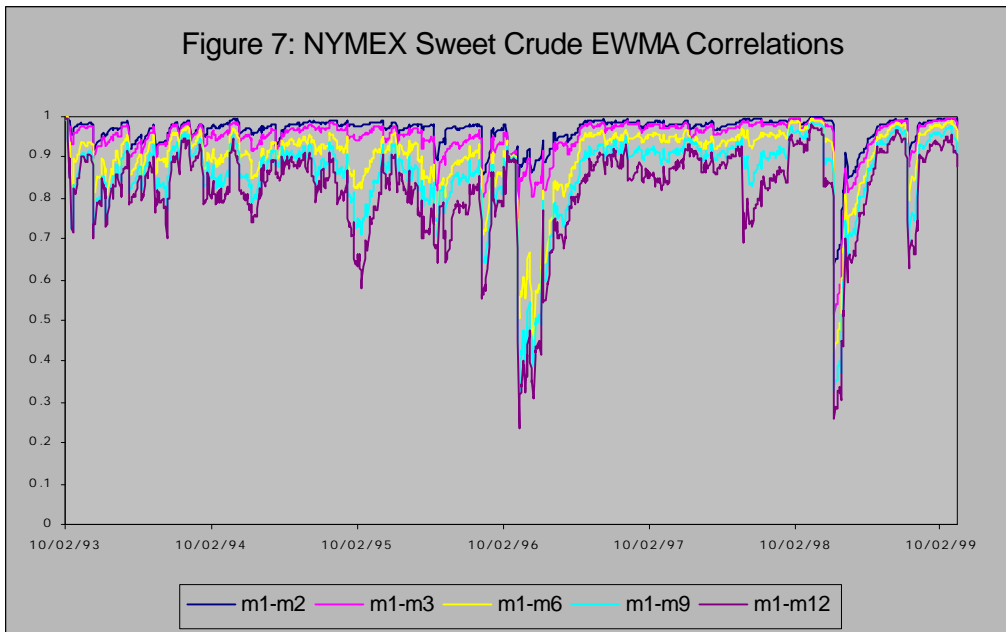
Correlations of the 1 month future with futures of other maturities are illustrated in figure 7, calculated using an EWMA with $\lambda = 0.94$. They decrease as the spread increases, as would be expected. But what is interesting about these correlations is their variability over time, a feature that is most pronounced in the less liquid longer maturities.

For example on 13[th] March 1996 the 1mth – 6mth correlation stood at 0.9, the 1mth – 2mth correlation at 0.92. One week later the 1mth - 2mth correlation had fallen to 0.8 but the 1mth - 6mth had dropped substantially, to 0.57. Even greater changes in correlations can occur with longer maturities, because most of the trading is on the shorter futures.

Interestingly, the relatively stable periods in correlation do not necessarily occur when the term structure itself is flat and/or stable. For example, correlations were relatively high and stable during the entire year from October 1996 to October 1997. But during this time both strong and weak backwardation, and weak contango situations

occurred. On the other hand correlations were quite unstable in 1998 and the early part of 1999, although the market experienced a strong contango throughout.

The expected price rises occurred at the end of the data period, in March 1999 when OPEC producers restricted supply. At the time of writing the contango has been replaced by a wobbly but flat term structure. Uncertain expectations about the direction of the next price move have had the effect of increasing volatility and decreasing correlation at the end of the sample.

GARCH

Generalised autoregressive conditional heteroscedasticity (GARCH) models extend the constant 'unconditional' model to time varying 'conditional' (or instantaneous) parameters of returns distributions. Although GARCH estimates are based on a time-varying correlation model and EWMA estimates are not, they are quite similar. In fact the EWMA is a simple integrated GARCH model without a constant term.

When it comes to modelling volatility, GARCH models have many advantages over EWMA. In particular:

(a) GARCH parameters are estimated independently and optimally using maximum likelihood. In the EWMA model the reaction and persistence coefficients, λ and $1-\lambda$ respectively, are constrained to sum to one. Also there is no one optimal technique for estimating the parameter of an EWMA;

(b) The GARCH stochastic volatility model gives convergent term structure forecasts, whereas the EWMA model assumes constant volatility. The EWMA is just an estimate, not a forecast. It is assumed that current levels of volatility will persist forever (the 'square root of time' rule), which is rather unrealistic.

GARCH may also be used to estimate and forecast correlation, but with much less success. The multivariate GARCH model introduced in chapter * may be used to obtain conditional correlation estimates that supposed to be time-varying, because they are based on the assumption of stochastic rather than constant correlation. So the variation in correlation through time is part of the model, and not just ascribed to 'noise', as it is with weighted average models.

But whereas GARCH volatility models are easy enough to implement, multivariate models often experience convergence problems. This is because the likelihood function becomes very flat and difficult to optimise as the number of parameters increases. The bivariate BEKK model has 11 parameters compared to the 3 parameters of a symmetric univariate GARCH, so often parameters are imposed.

For example the diagonal vech multivariate GARCH described in chapter * has only 9 parameters, but it assumes that the cross market effects are zero which is not very realistic. The multivariate GARCH volatilities and correlations that are estimated depend very much on the parameterisation chosen, and it is extremely difficult to determine which is the best GARCH model.

Given the uncertainty in correlation estimates, and the difficulty with doing anything else but assume the current correlation estimate is the forecast, the advantages of multivariate GARCH are nothing like as clear as those of univariate GARCH. Readers that wish to find out more about the subject have a huge literature to choose from. See for example the surveys in Bollerslev, Chou and Kroner (1992), Bollerslev, Engle and Nelson (1994) and Alexander (1998).

To summarise the correlation models introduced here, the simple weighted average methods are easy to implement and are recommended for use in different circumstances, and for different reasons. Long term equally weighted averages can provide a good indication of the average correlation over a large number of months,

and preferably years. But for short term correlation, exponentially weighted moving averages are recommended because they are similar to GARCH estimates and do not suffer from the 'ghost' features of equally weighted averages following extreme market events.

Some Applications of Correlation

Several of the chapters in this book describe the type of products being traded in energy markets, why they are traded and by whom, and the growth in volumes on such markets. So there is no need (and no space) to reiterate these discussions here: instead this section focuses on the way correlation effects the pricing of hedges, spreads and multi-asset options.

What should be done if a trade is based on a correlation that does not materialise? Even when energy markets are sufficiently liquid to admit correlation hedging, these correlations may be too unstable for hedging to be effective. In some cases the short term equally weighted averages or exponentially weighted averages jump around all over the place, and the long term averages will be very misleading for some time after an extreme event in the markets. It is therefore very important to conduct a thorough empirical testing of any model in which derivative prices are affected by correlations, to assess how realistic is the correlation measure.

SPREADS

A spread is a first order correlation product, so called because correlation has a direct influence on price through the volatility of the spread. Spread volatility, which is based on the formula $\sigma^2_{x-y} = \sigma^2_x + \sigma^2_y - 2\rho\sigma_x\sigma_y$ is lowest when underlyings are highly correlated.

For example consider the spread between the Kansas City "Western" KCBOT, and the NYMEX natural gas future contracts. The NYMEX has always traded at a premium, as shown by the positive spread in figure 8.

When the KCBOT was first introduced both volume and volatility were very low relative to the NYMEX. But as the KCBOT contract trading volume increased so did their relative volatility and the spread decreased substantially. In recent years their correlation has been very high and stable: although the KCBOT closes later, trading on this contract after the NYMEX has closed is still very thin. Consequently spread volatility is now relatively low and this affords a certain degree of predictability in the spread.

Figure 8: NYMEX - KCBOT Spread and EWMA Relative Volatility
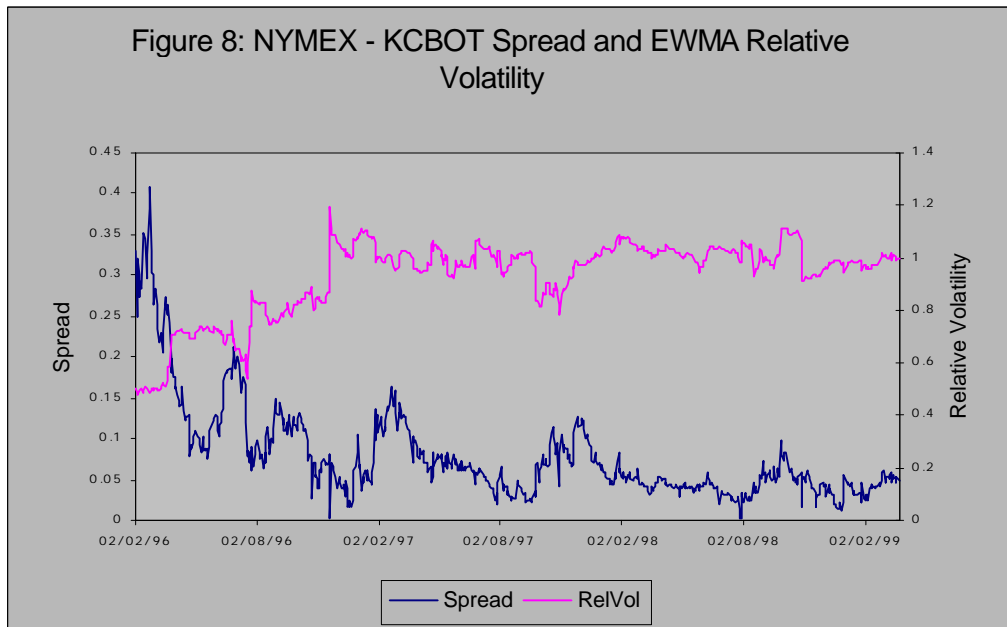
However close the contracts appear to be, caution should always be exercised when equally weighted correlation measures are used. The exponentially and equally weighted correlations shown in figure 9 were indeed very similar during the second half of the data period. But in the earlier days of the KCBOT, the use of equally weighted correlations created some misleading measures.

A small decoupling on a single day will effect the 30-day correlation for exactly 30 days, as for example following the sharp decrease in the spread and increase in relative volatility on 26th March 1996. Then 30 days later, on 9th May 1996, the 30-day correlation fell from 0.94 to 0.74, although nothing particular happened in the market on that day. It is just an artefact of the equal weighting of historical data. These effects are not apparent with the exponentially weighted correlation measure.

Figure 9: EWMA and 30-day Correlation of NYMEX-KCBOT Prompt Futures
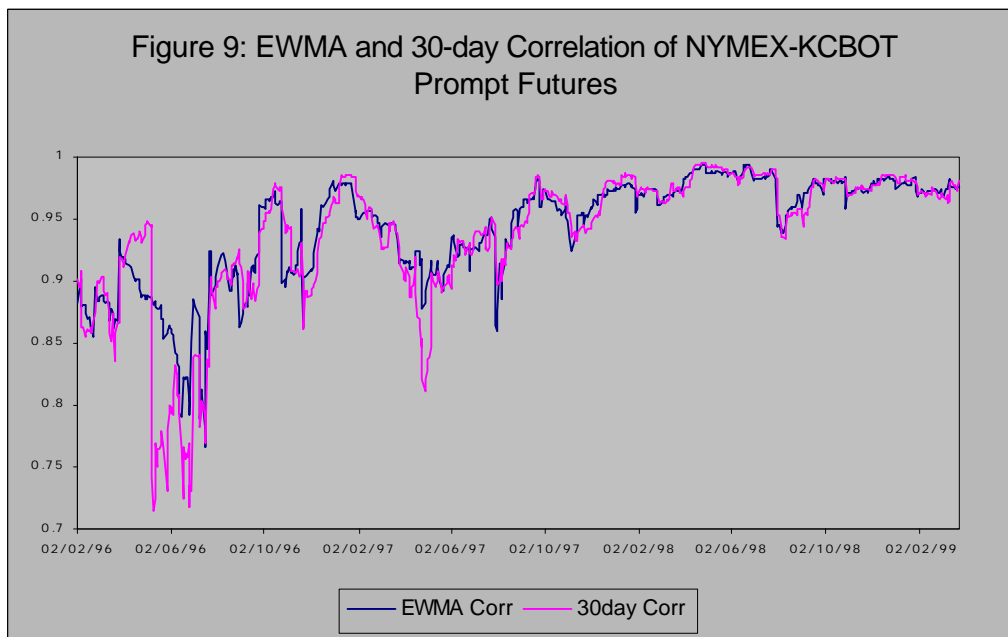
Traders that are fully hedged over a long period of time often wish to take advantage of short term profits or losses arising from movements in the spread. Such trades are usually based on the assumption that the spread is a stationary (mean-reverting) process. Of course many spreads *are* stationary, time spreads and crude spreads in particular. In that case not only will spread trading be relatively predictable, but also spread options will be reasonably cheap since the high correlation between prices serves to decrease spread volatility.

But in some energy markets there is empirical evidence that certain spreads are non-stationary. For example, on crack spreads that are heavily traded by refiners, speculators and arbitrageurs alike, market efficiency may result so that spreads become random walks (see figure 8 in chapter ** on the oil market).

But when spreads are not stationary the unconditional variance is infinite, the standard formula for spread volatility will not be valid, and univariate statistical models of non-stationary spreads will have little forecasting power. However it may be possible to use cointegration between related spreads to build an error correction model for trading, as explained in the next section.

HEDGING

Hedging is a very uncertain activity in energy markets, given the unique supply and demand structures that frequently decouple the spot price from prices of futures and forwards, and given the extreme volatility and unstable correlations that are inherant in these markets. Some producers may be unwilling to hedge at all, but there are still very many players, from end-users to speculators and arbitrage traders, that have created a large and growing demand for energy derivatives.

When hedging x with a single product y, hedge ratios are defined multiplying correlation by the relative volatility:

$$COV(R_{x,t}, R_{y,t})/\ V(R_{y,t}) = Corr(R_{x,t}, R_{y,t})\sqrt{(V(R_{x,t})/\ V(R_{y,t}))} \qquad (2)$$

Or in alternative notation $\beta = \sigma_{xy}/\sigma_y^2 = \rho\ (\sigma_x/\sigma_y)$. So when the correlation is close to 1, the hedge ratio is the relative volatility of the underlying with respect to the hedge.

Since hedge ratios depend on correlation, they may also display features that are purely artefacts of the model used to measure correlations.

For example if equally weighted averages are used, one would see 'ghost' effects in hedge ratios just as one does in correlation. These are evident in figure 10, which shows exponentially weighted and 30 day equally weighted hedge ratios for hedging the Western gas Kansas City contract with NYMEX prompt month futures.

Figure 10: EWMA and 30-day Hedge Ratios for NYMEX-KCBOT Prompt Futures
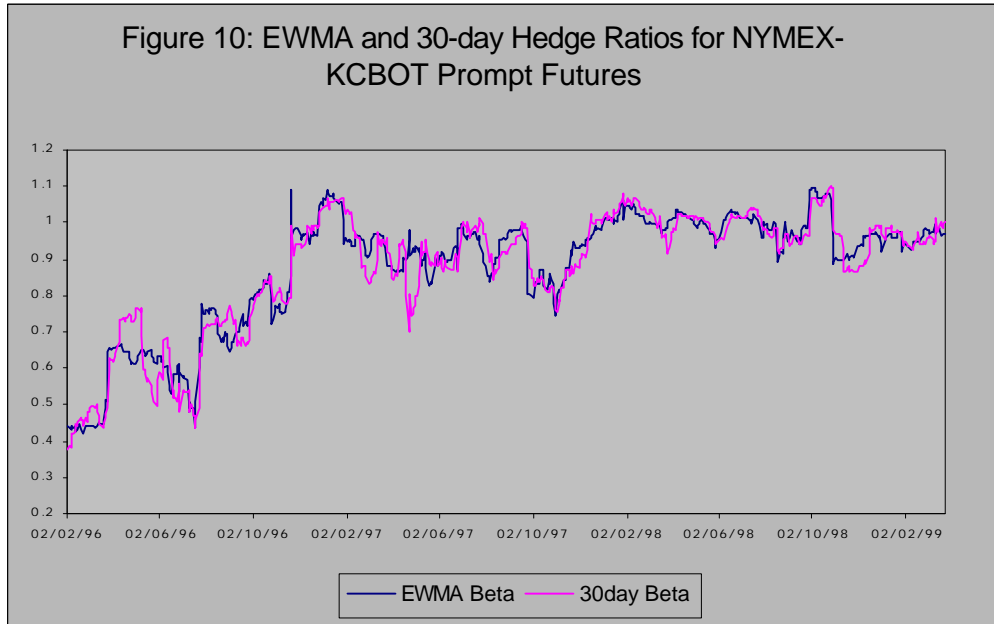
MULTI-ASSET OPTIONS

There are many products on related energy markets that are tailor made for end-users and producers alike. Energy producers that are exposed to many commodities will hedge revenues with basket options that are cheaper than buying options on individual markets. 'Best of' options allow end users to purchase energy supply at either the gas price or the oil price (say), which ever is better[4]. Long term swaptions and options on related markets diversify the risks from hedging all costs with derivatives based on a single market. Currency protected products allow the purchaser to hedge all foreign exchange risk, and derivatives for end-users can be based upon several indices.

The prices of these products depend to a greater or lesser extent on cross market correlations. Many of these derivatives are second order correlation products, so called because correlation has a lesser effect on price, affecting it only through changes in discount rates rather than directly through volatility.

For example, the price of a currency protected derivative depends on the 'quanto' correlation between the underlying and the exchange rate, but only in so far as it changes the discount rate in the Black-Scholes formula.

Although these quanto correlations are likely to be low, they can also be very unstable. But the instability of cross market correlations is not so much of an issue with second order products. It is the first order correlation products that will be very difficult to price when correlation is unpredictable.

Consider for example a basket option on crude oil and natural gas near month futures. The basket option should be cheaper than buying separate options on each underlying, because basket volatility is related to the volatility of individual options as $\sigma^2_{x+y} = (\sigma_x + \sigma_y)^2 - 2(1-\rho)\sigma_x\sigma_y$. Thus the basket volatility is less than the sum of individual volatilities unless $\rho = 1$.

Figure 11 shows equally weighted correlation measures between NYMEX prompt futures on crude oil and natural gas. The longer term correlations in figure 11 are very small, in the region of 0.1 to 0.2, so long term basket options on natural gas and crude oil should be relatively cheap. But, amongst other factors, differences in settlement dates and procedures across different markets produce highly unstable short term correlations, as for example in the 30 day correlation in figure 11. So even though they may be cheaper, prices of short term basket options will be subject to great variability.

---

[4] Price depends on correlation through the volatility of the price ratio (not the ratio of price volatilities)

Figure 11: Equally Weighted Correlation Measures of Natural Gas and Crude Oil
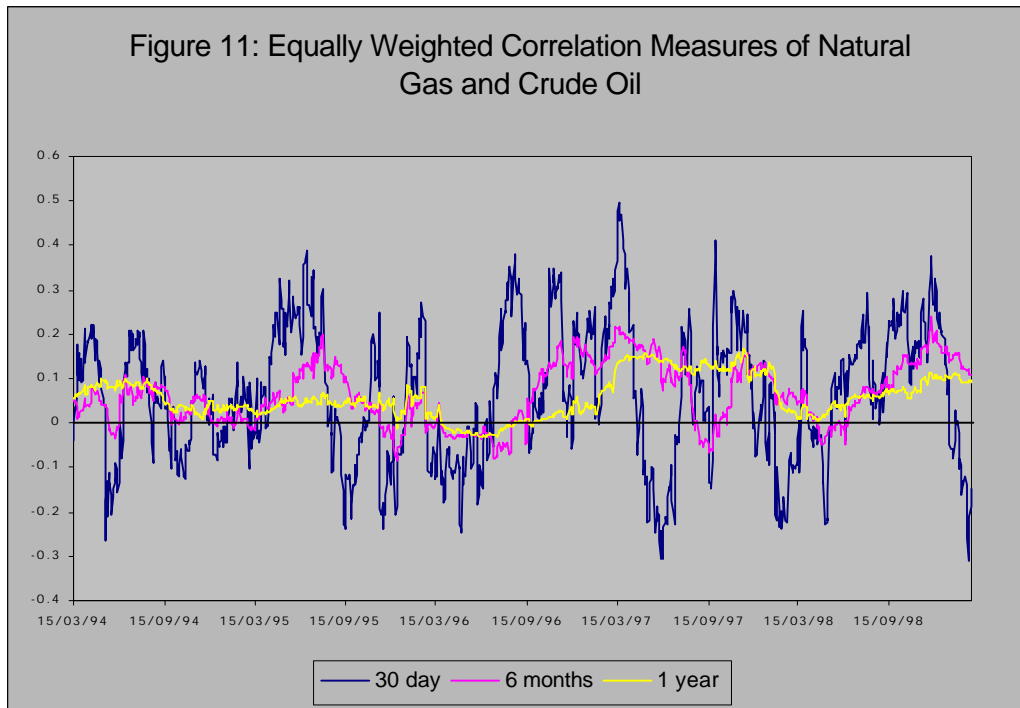
Introducing Cointegration

The classic paper on cointegration by Engle and Granger (1987) engendered a revolution in applied economic methods. Cointegration has emerged as a powerful technique for investigating common trends in multivariate time series, and provides a sound methodology for modelling both long run and short run dynamics in a system.

Although models of cointegrated financial time series are now relatively common place in the literature their importance has, until very recently, been mainly theoretical. This is because the traditional starting point for risk management is a correlation analysis of returns, whereas cointegration is based on the prices themselves.

In standard risk-return models the price data are differenced before the analysis is even begun, which removes a-priori any long-term trends in the data. Of course these trends are implicit in the returns data, but in correlation models any decision based on long-term common trends in the price data is excluded.

Cointegration extends the traditional model of correlation to include a preliminary stage in which the multivariate price data are analysed for long run equilibria called 'cointegrating vectors'. Then in a dynamic correlation model called the 'error correction model' (ECM) the causal flows between returns are investigated.

Lack of space requires that only a simplistic version of the theory is presented here, but the more general theory is more than adequately covered in standard econometrics texts such as Campbell, Lo and MacKinlay (1997), Hamilton (1994) and Hendry (1995, 1996).

Two price processes are cointegrated if there is a linear combination of these prices that is stationary, and any such linear combination is called the 'cointegrating vector'. The cointegrating vector is a spread, often taken to be a difference in log prices so that the error correction model is based on returns. So generally speaking when spreads are stationary prices are cointegrated. Of course prices may deviate in the short term, and correlations may be low at times, but they are 'tied together' by a long term common trend because of the mean-reversion in the spread.

Spot and futures prices are cointegrated when the basis is mean-reverting. Related commodity prices may be cointegrated if costs of carry are well behaved, but that is not always the case. Cointegration arises naturally in many other financial markets: equities within an index, along or between yield curves, in currency systems and between international market indices. The interested reader may consult Alexander (1999) and the references there for more details.

Box 3 analyses cointegration and error correction models in the crude oil market. This serves both to introduce some basic concepts in cointegration and to analyse how dynamic relationships can behave in energy markets. Although spot and prompt future prices are taken in this analysis, similar methods could be applied to any points in the term structure, to related non-stationary crack spreads, and indeed to any energy markets that have common stochastic trends.

Before commencing cointegration analysis the non-stationarity of data should be established. There are many statistical tests for stationarity described in the voluminous econometric literature on unit root tests. The test described in Phillips and Perron (1988) is perhaps the most appropriate given the fat-tailed nature of energy markets. But for simplicity only a basic test is described here.

The augmented Dickey-Fuller (ADF) statistic is based on a regression of Δx on a constant, one lag of Δx and one lag of x, where Δ denotes the first difference. The t-ratio on the lag of x is the ADF statistic, which has a 5% critical value of –2.88 and 1% critical value of –3.5. If the ADF exceeds the α% critical value then x is stationary at the α% level. (see Alexander and Johnson, 1994, Dickey and Fuller, 1979 and MacKinnon, 1994). Results of ADF tests on crude oil price data are shown in box 3.

The next step in cointegration is to establish that a 'cointegrating vector' exists between related price series. This is a linear combination of non-stationary prices that is stationary.

Box 3: Cointegration and Error Correction in Crude Oil Prices

The first step in cointegration is to check that the price data are non-stationary. In major equity, currency and fixed income markets daily prices are almost always governed by integrated processes, if not even a random walk. However unusual

demand and supply constraints influence the very uncertain prices and extreme variability that is a characteristic of many energy markets, and it is by no means certain that prices will be non-stationary.

The augmented Dickey-Fuller statistics from WTI crude oil log spot prices are shown in figure 12. Each regression is based on 5 years of data, a somewhat arbitrary choice but not a choice that hugely influences the qualitative nature of results. The regressions are rolled over the whole data period, each time recording the ADF, giving the data in figure 12. Although it exceeds the 5% critical value marked for a brief period in 1996 (an artefact of the Gulf War 5 years earlier) it stays well above the 1% level in recent years. So we conclude that WTI crude oil prices have significant non-stationary.

Figure 12: ADF for WTI Crude Oil Log Prices
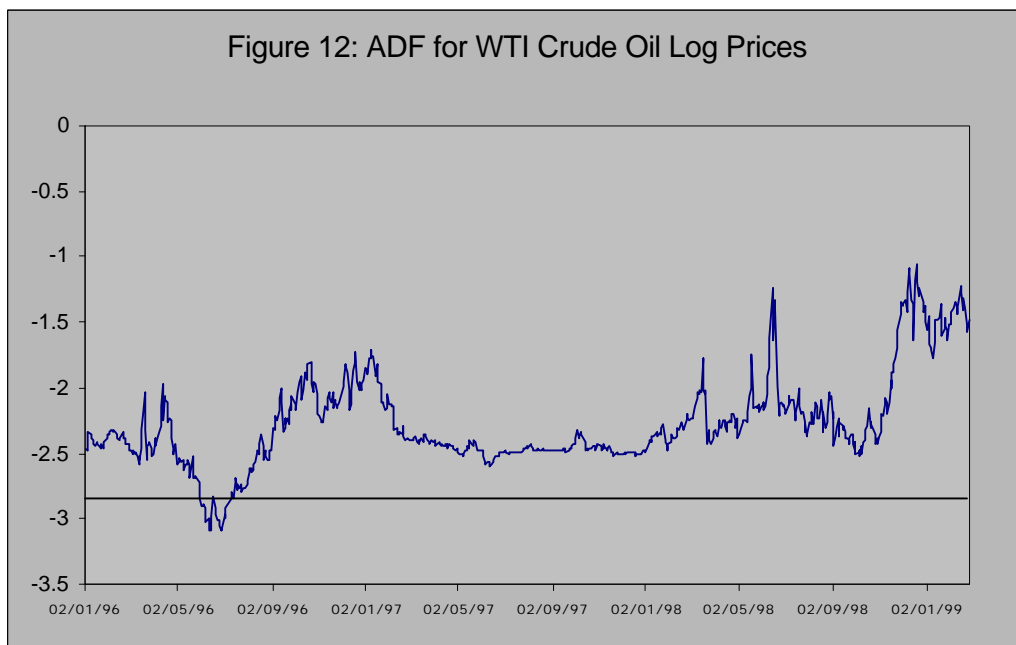
Spot and futures prices are cointegrated if there is a stationary linear combination. Whilst there may be more stationary combinations than the basis $z = \log F - \log S$, there is no doubt that the basis is stationary and it is more intuitive to use this as the cointegrating vector in error correction models.

Error correction models of the form (3) are estimated using ordinary least squares. Then simple t tests on the significance of the coefficients in rolling regressions show how the lead-lag relationship between spot and futures prices evolves over time. Figure 13 shows t-stats on α_2 (future to spot), α_3 (basis to spot), β_1 (spot to future) and β_3 (basis to future) for the crude oil data using rolling regressions on a 4 year window.

Figure 13: T Statistics in Error Correction Models
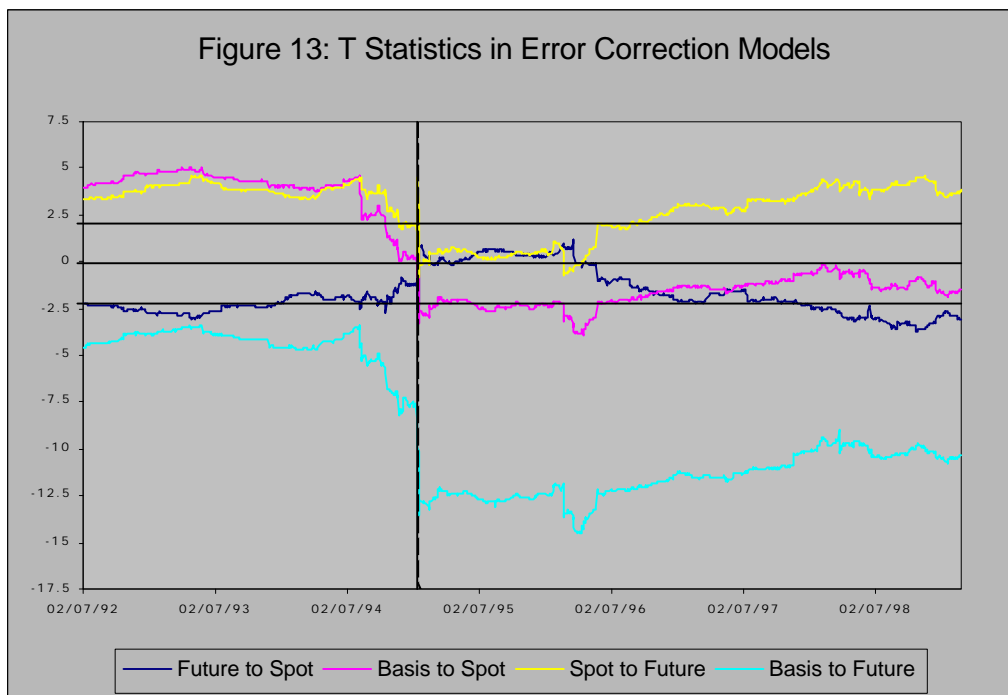
Note that after the structural break on 17th January 1995, when the dramatic fall in prices on 17th January 1991 drops out of the data, part of the error correction mechanism broke down. The coefficient α_3 is no longer positive. However the t-stats on β_3 are very large indeed, and negative as they should be, so the error correction mechanism is currently working through changes in futures prices.

Figure 13 gives a very clear message that it is futures and not spot prices that are being driven: there are very significant causalities from both spot prices and the basis into futures prices on the next day.

It is not surprising that futures are not good forecasts of spot prices in the crude oil market. In fact in any energy market, demand fluctuations produce an immediate response in spot prices because of the inelastic supply curve. The subsequent effect on inventory levels changes the convenience yield, but it may take time for futures prices to respond. It is the spot price, and in the crude oil market the basis even more so, that predicts futures prices. However spot prices are difficult to predict, which is to be expected since demand fluctuations are governed by so many unpredictable quantities.

In the simple case that only two series are considered, one would perform a regression of one log price y on the other log price x and then test the residual for stationarity. If the residuals indicate that the error process is indeed stationary, then the cointegrating vector is $z = y - \beta x$ where β is the regression coefficient. This is the method proposed by Engle and Granger (1987). For more details on this and other methods see Alexander (1999).

The final step is to estimate an error correction model (ECM) on returns, which may have quite a complex lag structure. If only the first lags are used the ECM takes the simple form:

$$R_x = \alpha_0 + \alpha_1 R_x(-1) + \alpha_2 R_y(-1) + \alpha_3 z(-1) + \varepsilon_s$$

(3)

$$R_y = \beta_0 + \beta_1 R_x(-1) + \beta_2 R_y(-1) + \beta_3 z(-1) + \varepsilon_f$$

The ECM is the mechanism that ties cointegrated series together in the long run. It takes its name from the fact that α_3 is positive and β_3 is negative. Thus if the cointegrating vector z is above (below) its equilibrium value, next period price of x will tend to increase (decrease), the price of y will tend to decrease (increase) and both serve to reduce the size of cointegrating vector.

The ECM may be used to analyse 'Granger' causality, which must be present in cointegrated series. Granger causality means that turning points in one series lead turning points in the other (Granger, 1988). If the ECM coefficients α_2 and/or α_3 are significant, there is a dynamic causality from y to x. If the coefficients β_1 and/or β_3 are significant, there is a causality from x to y. Granger causality between crude oil spot and futures is investigated in box 3.

In many financial markets there may be a case for 'price discovery' where futures prices lead spot prices (for example see Bopp and Sitzer 1987, Schroeder and Goodwin 1991, Schwartz and Laatsch 1991, Schwartz and Szakmary 1994, Wang and Yau, 1994).

But the relationship between spot and futures prices in energy markets is quite unique, as is demonstrated by the cointegration analysis of crude oil prices in box 3. For example in equity indices it is the future that is actually traded - completely the other way around to energy markets where only spot prices govern the physical delivery.

The information that is revealed in a dynamic correlation model, such as an error correction model, should be very useful for short term trades in and between many other energy markets. All the methods described here may be implemented in statistical packages such as Excel, so readers are invited to experiment.

Acknowledgement

Bibliography

Alexander, C. and A. Johnson (1994) "Dynamic Links" *RISK* 7:2 pp 56-61

Alexander C.O. (1996) 'Evaluating the use of RiskMetrics™ as a risk measurement tool for your operation : What are its advantages and limitations' *Derivatives Use, Trading and Regulation* **2** No. 4

Alexander, C (1998) "Volatility and correlation: measurement, models and applications" in *Risk Management and Analysis (2nd ed.) Volume 2: Measuring and Modelling Financial Risk* Edited by C. Alexander, John Wiley and Sons

Alexander, C. (1999) "Optimal hedging using cointegration" *Philosophical Transactions of the Royal Society (forthcoming, August 1999)*

Bollerslev, T, R.Y. Chou and K.F. Kroner (1992) "ARCH modeling in finance". *Journal of Econometrics* **52** pp 5-59.

Bollerslev, T. R.F. Engle and D.B. Nelson (1994) "ARCH models" in *Handbook of Econometrics* volume 4 Edited by R.F. Engle and D.L. McFadden, North Holland.

Bopp A.E. and S. Sitzer (1987) "Are petroleum prices good predictors of cash value?" *Journal of Futures Markets* 7 pp 705-719.

Campbell, J.Y., A.W. Lo and A.C. MacKinley (1997) *The Econometrics of Financial Markets* Princeton University Press.

Choi, I. (1992) "Durbin-Hausmann tests for a unit root" *Oxford Bulletin of Economics and Statistics* 54:3 pp. 289-304.

Chowdhury, A.R. (1991) "Futures market efficiency: evidence from cointegration tests" *The Journal of Futures Markets* 11:5 pp 577-89. John Wiley and Sons Inc., 1991.

Covey, T. and D.A. Bessler (1992) "Testing for Granger's full causality". The *Review of Economics and Statistics* pp 146-153.

Dickey, D.A. and W.A. Fuller (1979) "Distribution of the estimates for autoregressive time series with a unit root" *Journal of the American Statistical Association* 74 pp. 427-9

Engle, R.F. and C.W.J. Granger (1987) "Co-integration and error correction: representation, estimation, and testing". *Econometrica* 55:2 pp 251-76.

Granger, C.W.J. (1988) "Some recent developments on a concept of causality". *Jour. Econometrics* <u>39</u>, pp. 199 - 211.

Hamilton, J.D. (1994) *Time Series Analysis* Princeton University Press

Hendry, D.F. (1986) "Econometric modelling with cointegrated variables: an overview". *Oxford Bulletin of Economics and Statistics* 48:3 pp 201-12.

Hendry, D.F. (1995) *Dynamic Econometrics* Oxford University Press

Mackinnon, J.G. (1994) "Approximate asymptotic distribution functions for unit-root and cointegration tests" *Journal of Business Economics and Statistics* 12:2, April 1994, pp. 167-176.

Phillips, P.C.B. and P. Perron (1988) "Testing for a unit root in time series regressions" *Biometrika* 75 pp. 335-46

Schroeder , T C and B.K. Goodwin (1991) "Price discovery and cointegration for live hogs" *Journal of Futures Markets* 11 : 6 pp685-696

Schwarz, T.V. and F.E. Laatsch (1991) "Dynamic efficiency and price discovery leadership in stock index cash and futures market" *Journal of Futures Markets* 11 : 6 pp669-684.

Schwarz, T.V. and A.C. Szakmary (1994) "Price discovery in petroleum markets: arbitrage, cointegration, and the time interval of analysis". *Journal of Fututres Markets* 14:2, pp 147-167.

Wang, G. H. K. and J. Yau (1994) "A Time Series Approach To Testing For Market Linkage: Unit Root And Cointegration Tests" *Journal of Futures Markets* 14:4