Vincent Grégoire, Christian Genest and *Michel Gendron* show how forecasts for crude oil and natural gas prices can be improved by modelling the dependence between them. Movements in individual returns are described by time-series models, their dependence is captured by a copula, and Monte Carlo simulations are used to forecast prices

Using copulas to model price dependence in energy markets

★ Dependence between prices is crucial in many aspects of financial risk management, such as multi-asset derivative pricing and portfolio hedging. This is especially so when modelling price behaviour in energy markets because these commodities have become increasingly intertwined in recent years. For instance, both natural gas and crude oil are used to generate electricity and heating. Natural gas is used also to extract oil from tar sands.

Alexander (2005) studied dependence between prices for futures on crude oil and natural gas. She concluded that it is strong and cannot be modelled correctly by a bivariate Normal distribution. More precisely, when prices for these commodities are fitted by time-series models, the distributions for the error terms are found to be asymmetric. Moreover, crude oil and natural gas log-returns exhibit non-linear dependence. What could we do to take these factors into account?

In this paper, we use copulas to study the relationship between prices for futures on crude oil and natural gas. We model the log-returns on each commodity individually as time series, and we account for the dependence between them by fitting various families of copulas to the error terms. To help us select the best copula, we perform a range of goodness-of-fit tests. Finally, we use Monte Carlo simulations to derive a joint predictive distribution for the prices of crude oil and natural gas.

Copula models

Through a detailed description of the steps involved in building a copula model, our goal is to illustrate the generality and the flexibility of this approach. As we shall see, multivariate Normality is only one option in a wide range of copula-based models that can capture the critical features of financial data, such as asymmetry, non-linear dependence or heavy-tail behaviour.

Let X and Y be continuous random variables, with distribution functions $F(x) = \Pr(X \le x)$ and $G(y) = \Pr(Y \le y)$, respectively. Following Sklar (1959), there exists a unique function C such that:

$$\Pr(X \le x, Y \le y) = C(F(x), G(y)) \tag{1}$$

where $C(u, v) = \Pr(U \le u, V \le v)$ is the distribution of the pair $(U, V) = (F(X), G(Y))$ whose margins are uniform on [0, 1]. The function C is called a copula. As argued by Joe (1997) or Nelsen (1999) among others, C characterises the dependence in the pair (X, Y).

When the joint distribution of (X, Y) is unknown, we can model it by assuming specific parametric forms for F, G and C in (1). For example, we might take X and Y to be exponential with different means $\kappa, \lambda > 0$. This is achieved using:

$$F(x) = 1 - e^{-x/\kappa}, \qquad G(y) = 1 - e^{-y/\lambda} \tag{2}$$

for all $x, y > 0$. For simplicity, suppose also that we choose to model the dependence by the Farlie-Gumbel-Morgenstern or FGM copula:

$$C(u, v) = uv + \theta uv(1-u)(1-v). \tag{3}$$

Here θ in [–1, 1] is a parameter to be determined. Replacing u by $F(x)$ and v by $G(y)$ in equation (3), we have a copula model for (X, Y):

$$\begin{aligned}\Pr(X \le x, Y \le y) &= F(x)G(y) + \theta F(x)\,G(y)\{1-F(x)\}\{1-G(y)\} \\ &= \left(1-e^{-x/\kappa}\right)\left(1-e^{-y/\lambda}\right)\left(1+\theta e^{-x/\kappa}e^{-y/\lambda}\right).\end{aligned}$$

Although the joint distribution seems complex at first sight, a closer look reveals that it is in fact a simple function of F and G.

A distinct advantage of copula modelling is that distributions F, G and C in (1) can be chosen independently of one another. Depending on the circumstances, we could decide to replace (2) by a statement that F is log-Normal and G is a Pareto distribution, say. This could be done while keeping C in the FGM family of copulas. But for fixed F and G, we could choose also to select C from another class of copulas. Popular choices include the Archimedean, extreme-value, and meta-elliptical families of copulas described by Joe (1997), Nelsen (1999) and Fang *et al* (2002).

Copulas are a powerful modelling tool because they can fit a large range of dependence structures. Copula modelling techniques are widely recognised in statistics, biostatistics and actuarial science. In the past 10 years, copulas have also become increasingly popular in finance, where they have found applications in derivatives pricing, credit risk and portfolio management, and so on. Introductions to copula modelling in statistics, finance and quantitative risk management have been written by Genest & Favre (2007), Cherubini *et al* (2004) and McNeil *et al* (2005). Copulas can be used also to fit data with complex time-varying patterns; see Patton (in press) for applications in econometrics and finance. We shall see below that it is fairly simple to infer the best-fitting copula and derive more realistic predictions than under the assumption that the joint distribution of (X, Y) is Normal. Assuming bivariate Normality amounts to taking both F and G as Normal, and restricting C to a specific parametric class of copulas called the Normal copulas.

Oil and gas market dependence over time

Consider prices of crude oil and natural gas, based on one-month-ahead futures contracts traded on Nymex. This data is available from Bloomberg's financial information services. The log-returns for both series are plotted in figure 1 from July 1, 2003 to July 19, 2006. The top chart shows the log-returns as a function of time for the front-month light sweet crude oil futures contract, which is widely reported as a proxy for the cost of imported crude oil. Displayed in the bottom chart are log-returns as a function of time for a futures contract on 10,000 million British thermal units (MMBtu) of natural gas to be delivered at Henry Hub, a pricing point on the natural gas pipeline located in Erath, Louisiana, USA.

A rolling gap occurs in each series on the day of the month when the contract expires. Part of the variation observed on these days is the daily movement, but part is simply due to the difference in value for each commodity between the two contracts. For the period considered, these rolling gaps were judged to be sufficiently small to be ignored.

It is clear from figure 1 that the log-returns on crude oil and natural gas vary over time, but whether they are also interdependent is more difficult to tell. Before we can settle this issue, we must account for autocorrelation and heteroscedasticity in the marginal series.

Models for the marginal series

We performed standard Box-Pierce and Ljung-Box tests on log-returns and their squared values to detect the presence of autocorrelation and heteroscedasticity in the series. For crude oil, we found that neither of these tests is significant at the 5% level ($P > 0.05$). Consequently, crude oil log-returns $O_1, O_2, \ldots$ observed from July 2003 to July 2006 can be viewed as a random sample from some distribution F with mean $\mu = 1.220 \times 10^{-3}$ and standard deviation $\sigma = 1.912 \times 10^{-2}$.

However, the tails of F are fatter than those of a Normal distribution. Standard goodness-of-fit tests confirm that the error term $X_t = (O_t - \mu)/\sigma$ is not Normal; its behaviour is more accurately represented by a Student distribution with $\psi = 13.745$ degrees of freedom.

When we applied Box-Pierce and Ljung-Box tests to natural gas log-returns N_t, we observed strong autocorrelations at various lags between their squared values. We concluded that the null hypothesis of homoscedasticity should be rejected ($P < 0.05$ at lag 1, $P < 0.001$ at lag 10). Generalised autoregressive conditional heteroscedasticity (Garch) methodology is often used to model heteroscedasticity in financial series. Using the inference techniques described by Bollerslev (1986), we found that the

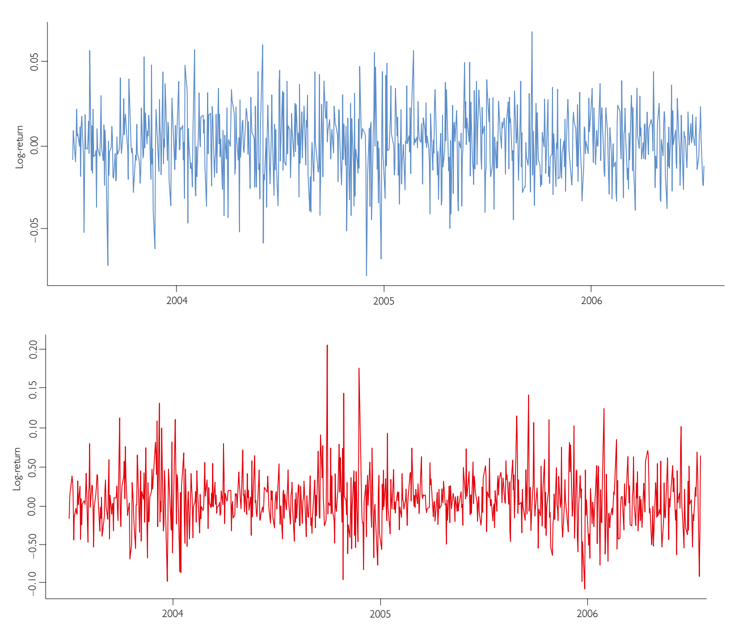

F1. Futures: crude oil versus natural gas
Log-returns for one-month-ahead futures on crude oil (top) and natural gas (bottom) *Source: Bloomberg*

following Garch(1,6) model is an adequate representation of the variation in the natural gas log-return series:

$$N_t = 1.281\times10^{-4} + Y_t h_t^{1/2},$$
$$h_t = 2.138\times10^{-4} + 6.206\times10^{-2} N_{t-1}^2 + 3.663\times10^{-3} N_{t-2}^2 + 5.862\times10^{-2} N_{t-4}^2 + 9.319\times10^{-2} N_{t-5}^2 + 9.472\times10^{-2} N_{t-6}^2 + 0.535 h_{t-1}.$$

In this model, the error terms $Y_1, Y_2, \dots$ form a random sample from some distribution G. Typically, the latter is assumed to be Normal but here, this hypothesis is unrealistic. A better choice is given by the skewed t-distribution introduced by Azzalini & Capitanio (2003). Thus we assumed that $Y_t = \xi + V^{-1/2} Z$, where ξ is a location parameter, V is a chi-square random variable with ν degrees of freedom and Z is a skew-Normal variate independent of V. The distribution of Z involves parameters Ω and α controlling dispersion and shape, respectively. In our case, we found

$$Y_t \sim \text{skewed} - t \left(\xi = -0.622,\ \Omega = 0.986,\ \alpha = 1.014,\ \nu = 6.737\right). \quad (4)$$

Accounting for dependence between the series

The models given above provide adequate descriptions of the daily movements in crude oil log-returns O_t, and natural gas log-returns N_t, taken individually. However, it may be that the prices for these commodities exhibit some dependence. Accounting for this dependence, if any, would lead to more realistic predictions for prices of both crude oil and natural gas. From a statistical viewpoint, we must address two questions: (i) Are the variables X_t and Y_t independent? (ii) If not, what would be an appropriate joint distribution for the pair (X_t, Y_t)? This is where copulas come in handy. If $X = X_t$ and $Y = Y_t$ are independent, then their joint distribution can be expressed in the form (1) with:

a) F: the Student distribution with ψ=13.745 degrees of freedom;
b) G: the skewed t-distribution specified in (4);
c) $C(u,v) = uv$ for all u, v in [0,1].

However if there is dependence between X and Y, then we need another copula C to 'glue' the marginal distributions F and G together. When F and G are specified as above, copulas are the single most convenient way of combining them.

We can address issues (i) and (ii) simultaneously using the empirical copula (Deheuvels, 1979). This is a consistent estimation C_n of the true underlying copula C. Several variants of Deheuvels' definition exist; the most common version assigns a weight of $1/n$ to each pair (u_t, v_t), where

$$u_t = \frac{1}{n+1} \times \text{ rank of } X_t \text{ among } X_1, \dots, X_n,$$
$$v_t = \frac{1}{n+1} \times \text{ rank of } Y_t \text{ among } Y_1, \dots, Y_n.$$

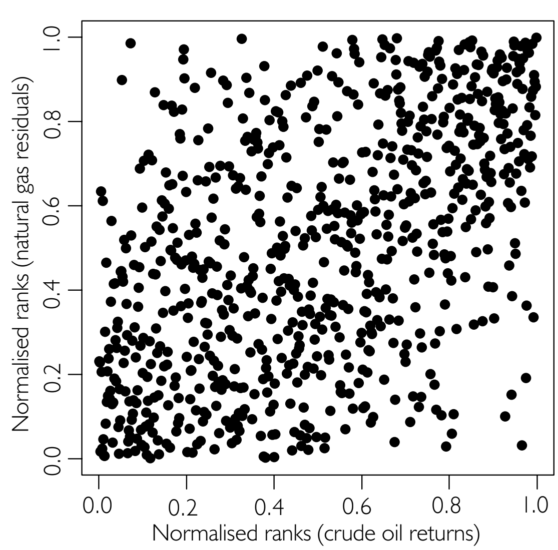

F2. Dependence between crude and natural gas
Pairs of normalised ranks showing the dependence between crude oil and natural gas

Figure 2 shows a graph of the pairs (u_t, v_t) for the n=756 observations at hand. It is clear from it that there is positive dependence between u_t and v_t. This means that whatever their marginal distributions, X_t and Y_t have a tendency to vary together. In other words, the margins F and G identified earlier are not connected by the independence copula $C(u, v) = uv$.

To confirm this judgement, we can compute the correlation between the values of u_t and v_t. This measure of dependence is called Spearman's rho rather than Pearson's correlation, simply to emphasise that it is computed from the ranks of the data, not from the data itself. Here we find ρ_n=0.508, which is significantly different from zero ($P < 0.001$). This conclusion is based on the statistic $(n-1)^{1/2}\rho_n$, which is standard Normal for large n under the null hypothesis of independence.

Kendall's tau is another rank-based measure of dependence often used in this context. To compute τ_n, a line is drawn to connect the pairs (u_1, v_1) and (u_2, v_2). If the slope is positive, we count 1; if the slope is negative, we count –1 instead. We repeat the process for all other choices of distinct pairs (u_s, v_s) and (u_t, v_t). There are $m = n(n-1)/2$ such choices and τ_n is simply the sum of all 1s and –1s, divided by m. (Adjustments are necessary if the slope is 0 or infinite, but this shouldn't occur if the data are continuous and measured with precision.)

Spearman's rho and Kendall's tau emphasise different aspects of the dependence. Both are valid, although in practice, it is often the case that $\tau_n < \rho_n$ (Capéraà & Genest, 1993). Here τ_n = 0.349 is smaller than ρ_n=0.508 but significantly different from zero nonetheless (P<0.001). The test is based on the statistic $\{(9/2)\, n(n-1)/(2n+5)\}^{1/2}\tau_n$, which is standard Normal for large n under the null hypothesis of independence.

Selection of the copula: estimation

Having concluded that the pairs (X_t, Y_t) are associated, we must find a suitable copula C to describe this dependence. This copula should be such that if n pairs $(U_1,V_1),\ldots,(U_n,V_n)$ were drawn from it repeatedly, then on average the resulting graph should look like figure 2. A close match should also be observed between empirical measures of dependence ρ_n, τ_n and the parameters they estimate, as per the following:

$$\rho(C) = -3 + 12\iint C(u,v)\,\mathrm{d}v\mathrm{d}u,$$
$$\tau(C) = -1 + 4\iint C(u,v)\,c(u,v)\,\mathrm{d}v\mathrm{d}u, \quad (5)$$

where $c(u,v)=\partial^2 C(u,v)/\partial u\partial v$ is the density of C (assuming it exists). Spearman's rho and Kendall's tau range from –1 to 1 and their population values vanish in case of independence; see Chapter 5 of Nelsen (1999) for derivations of the above formulas and additional details. Note that the range for $\rho(C)$ and $\tau(C)$ can be restricted when C belongs to certain parametric classes. For example, $\rho(C_\theta)=\theta/3$ spans $[-1/3,1/3]$ and $\tau(C_\theta)=2\theta/9$ is limited to $[-2/9,2/9]$ for the FGM copula. This model is clearly inadequate here because we cannot get $\rho(C_\theta)=0.508$ or $\tau(C_\theta)=0.349$ for any value of θ in $[-1,1]$. Besides, taking $|\theta| > 1$ in (3) yields an invalid distribution.

To assist us in the choice of C, we considered six families of copulas commonly used in finance. They are listed in table 1, along with their parameter range. The first three are Archimedean (Joe, 1997; Nelsen, 1999) and the last two are meta-elliptical (Fang *et al*, 2002). Plackett's copula is in a class of its own.

We can see from table 1 that the Student copula actually involves two parameters: θ and γ. In this model, the parameter θ governs dependence, while the degrees of freedom, γ, modify the shape of the copula for a given level of dependence θ. In our study, we allowed γ to vary between 2 and 30.

For each model in table 1, we estimated θ by inversion of Kendall's τ. Specifically, the theoretical value $\tau(\theta)$ was computed using (5) for each possible copula, and the equation $\tau(\theta)=\tau_n$ was solved to obtain θ_n. Table 2 lists expressions for θ as a function of τ for four of the six families, along with the estimates obtained by solving the equation $\tau(\theta)=\tau_n$. For example, $\tau(C_\theta)=\theta/(\theta+2)$ for Clayton's copula and hence $\theta=2\tau/(1-\tau)$; from $\tau=\tau_n=0.349$ we get $\theta=\theta_n=1.07572$. For the Frank and Plackett copulas, τ cannot be computed analytically; therefore, numerical interpolation was used.

Under weak regularity conditions on the copula family, the inversion of Kendall's tau is known to yield a consistent estimator of the dependence parameter (Genest & Rémillard, 2008). Inversion of Spearman's rho would provide a different estimator that is just as good but as it turns out, Kendall's tau is easier to compute analytically than Spearman's rho for most copula families currently used in practice. For this reason, inversion of Kendall's tau is often favoured and we follow this trend.

Selection of the copula: goodness-of-fit testing

When $\theta=\theta_n$ as per table 2, each of the copulas C_θ listed in table 1 meets the minimal requirement $\tau(C_{\theta n})=\tau_n=0.349$. However, this is not sufficient to guarantee that C_θ is a good model for the data. For example, it could be that $\rho(C_{\theta n})$ is very different from $\rho_n=0.508$ or that other important data characteristics have been missed.

To help us select the most appropriate copula from table 1, we must rely on testing. A review and comparison of goodness-of-fit procedures is given by Genest *et al* (2008). These authors favour 'blanket' tests, that is, rank-based procedures requiring no strategic choice such as kernel, bandwidth, etc. From their simulations, a good combination of power and conceptual simplicity is provided by the Cramér-von Mises statistic:

$$S_n = n\sum_{t=1}^{n}\left\{C_{\theta n}(u_t,v_t) - C_n(u_t,v_t)\right\}^2.$$

This statistic measures how close the fitted copula $C_{\theta n}$ is from the empirical copula C_n whose support appears in figure 2 for the data at hand. Because the definition of S_n involves θ_n, the distribution of this statistic depends on the unknown value of θ under the null hypothesis that C is from the class C_θ. Thus the P-value of the test must be computed using a parametric bootstrap procedure described by Genest *et al* (2008).

T1. Six families of copulas

Family	$C(u,v)$	Range of θ
Clayton	$\max[0,1-\{(1-u)^\theta+(1-v)^\theta\}^{1/\theta}]$	$[-1,\infty)-\{0\}$
Frank	$\theta^{-1}\ln\{1+(e^{\theta u}-1)(e^{\theta v}-1)/(e^\theta-1)\}$	$(-\infty,\infty)-\{0\}$
Gumbel	$\exp\{-(\lvert\ln u\rvert^\theta+\lvert\ln v\rvert^\theta)^{1/\theta}\}$	$[1,\infty)$
Plackett	$[1+(\theta-1)(u+v)-$ $[\{1+(\theta-1)(u+v)\}^2-4uv\theta(\theta-1)]^{1/2}]/\{2(\theta-1)\}$	$(0,\infty)-\{1\}$
Normal	$N_\theta(\Phi^{-1}(u),\Phi^{-1}(v))$	$[-1,1]$
Student	$T_{\theta,\gamma}(T_\gamma^{-1}(u),T_\gamma^{-1}(v))$	$[-1,1]$

Φ: cumulative distribution function (cdf) of a $N(0,1)$
N_θ: cdf of a standard bivariate Normal distribution with Pearson correlation θ
T_γ: cdf of a Student with γ degrees of freedom
$T_{\theta,\gamma}$: cdf of the bivariate Student distribution with γ degrees of freedom (Fang *et al*, 2002)

T2. Estimates of the dependence parameter θ for different copulas

Family	Formula	θ_n
Clayton	$\theta = 2\tau/(1-\tau)$ for $\tau > 0$	1.076
Frank	Obtained numerically	3.474
Gumbel	$\theta = 1/(1-\tau)$ for $\tau > 0$	1.538
Plackett	Obtained numerically	5.107
Normal	$\theta = \sin(\pi\tau/2)$	0.522
Student	$\theta = \sin(\pi\tau/2)$	0.522

T3. P-values for the goodness-of-fit of different copula models

Assumed copula	P-value	Assumed copula	P-value
Clayton	0.0000	Plackett	0.0080
Frank	0.0095	Normal	0.0140
Gumbel	0.0075	Student, γ=22	0.0300

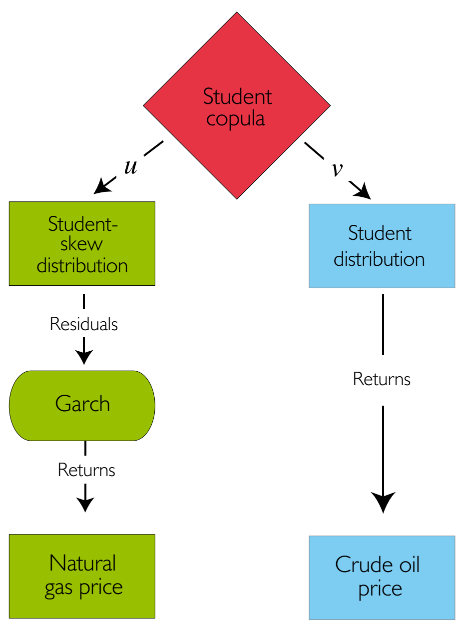

F3. A schematic description of the forecasting process

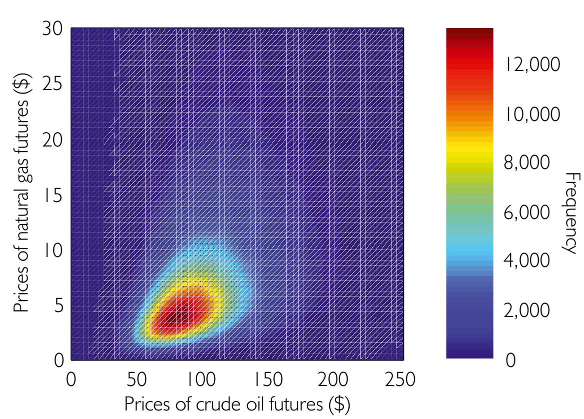

F4. Looking ahead
Predictive distribution for the prices of crude oil and natural gas futures in July 2007

In order to distinguish between the copula models listed in table 1, we applied the parametric bootstrap procedure of Genest *et al* (2008) to both the crude oil and natural gas data. For each model considered, we generated bootstrap values $S^*_1,\ldots,S^*_{2000}$ of the Cramér-von Mises test statistic and we determined the proportion of these values that are larger than S_n.

Table 3 lists the (one-sided) P-values that we obtained. In particular, we found $P\approx 0.03$ for the Student copula with $\theta = 0.522$ and $\gamma = 22$ degrees of freedom. All other models were rejected at this level, including the Student copulas with $\gamma \neq 22$ degrees of freedom (not displayed in table 3). With a P-value of 1.4%, the Normal copula would be a viable alternative but this is still a far cry from the assumption that the distribution of the pair (X_t, Y_t) is bivariate Normal. For, this assumption amounts to claiming that C is the Normal copula and that both F and G are also Normal. However, Normal margins are inappropriate in the present context.

In summary, an adequate description of the dependence in the pairs (X_t, Y_t) portrayed in figure 2 is given by

$$C(u,v) = T_{0.522,22}\left(T_{22}^{-1}(u), T_{22}^{-1}(v)\right). \quad (6)$$

This finding is consistent with recent studies showing that the Student copula often provides a much better fit of multivariate financial return data than the Normal copula; for example, see Breymann *et al* (2003). As mentioned by Demarta & McNeil (2005), a distinct advantage of Student copulas over Normal copulas is their ability to capture heavy-tail behaviour – that is, the phenomenon of dependent extreme values so often observed in financial returns. Another important feature of the final model is the non-linear dependence of Y_t on X_t. If we had taken for granted the traditional assumption that the joint distribution of (X_t,Y_t) is Normal, we would not only have gotten the marginal distributions wrong: we would have forced $E(Y_t|X_t=x)$ to be linear in x, which it isn't.

Forecasting

Forecasting with a copula model is simple. The basic steps are summarised in figure 3 for the specific model at hand. At time t, we generate a pair (u_t,v_t) from the Student copula C defined in (6). Then we set $X_t=F^{-1}(u_t)$ and $Y_t=G^{-1}(v_t)$, where F is the Student distribution with ψ=13.745 degrees of freedom, and G is the skewed t-distribution given in (4).

Finally, we compute $O_t=\sigma X_t+\mu$ and we obtain N_t via the Garch model described earlier. By using this procedure, we can construct price series for both crude oil and natural gas one day at a time in successive steps. To assess the uncertainty associated with this prediction, we must repeat the process a large number of times. This leads to an empirical predictive distribution.

Simulation results

We simulated 4,000,000 random paths using the above copula model for a one-year time horizon using July 19, 2006 as the initial time step.

The bivariate density of prices is shown in figure 4. It is obvious from this graph that the joint distribution is asymmetric and heavy-tailed. The same observation holds for the predictive distribution of natural gas, conditional on any price of crude oil. For example, figure 5 shows the distribution of forecasts for the price of natural gas in July 2007, assuming that the price of crude oil is around \$70. The median of this distribution is \$4.77, while the price of natural gas observed in that period was around \$6.50.

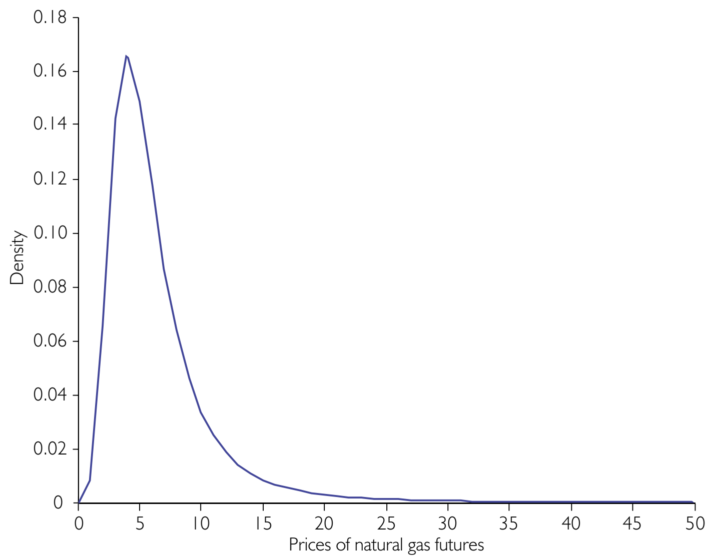

F5. Price prediction for natural gas

Predictive distribution for the price of natural gas in July 2007, assuming a price of \$70±\$5 for crude oil

Extension

An implicit assumption of the preceding analysis is that the dependence between the prices of crude oil and natural gas is (roughly) constant with time. Otherwise, we couldn't represent this dependence by a time-invariant copula C. To check this assumption, the value $\tau_{250}(t)$ of Kendall's tau between the log-returns for the two commodities was computed for successive periods of 250 working days ending at t. The results are plotted in figure 6. The graph shows an upward trend in $\tau_{250}(t)$ between April 1990 and July 2006. Starting in April 2004, however, $\tau_{250}(t)$ appears to have stabilised and is oscillating around 0.35. In other words, the dependence between X_t and Y_t is roughly constant since July 2003. This is consistent with the hypothesis that the dependence between pairs (X_t, Y_t) of error terms can be represented by a single copula in this period.

What if the dependence embodied in the copula C changed over time? To deal with this issue, we might assume that the joint distribution of the pair (X_t, Y_t) is of the form (1) with $C=C_\theta$ with $\theta=\theta_t$ a function of time t. To ensure that the dependence structure is identifiable, we would need to assume a pattern for the variation – for example, $\theta_t = at+b$. Van den Goorbergh *et al* (2005) illustrate this approach in option pricing based on the S&P 500 and the Nasdaq indexes. While a dynamic-copula approach could be adapted to the present context, much remains to be done to develop efficient estimation procedures and goodness-of-fit tests for cases where dependence varies over time. See Patton (in press) for further discussion. ER

Vincent Grégoire and Michel Gendron: Département de finance et assurance, Université Laval, Québec, Canada
Christian Genest: Département de mathématiques et de statistique, Université Laval, Québec, Canada
Emails: Vincent.Gregoire.1@ulaval.ca, Christian.Genest@mat.ulaval.ca, Michel.Gendron@fsa.ulaval.ca
Funding in support of this work was provided by the Natural Sciences and Engineering Research Council of Canada, the Fonds Québécois de la Recherche sur la Nature et les Technologies, and the Institut de Finance Mathématique de Montréal

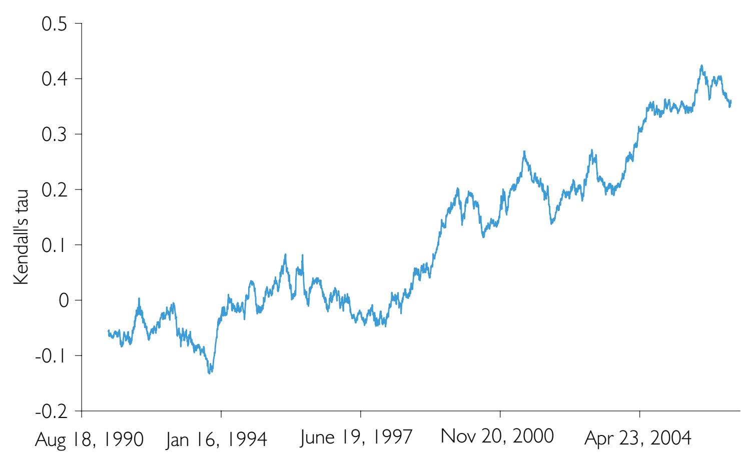

F6. Rolling window

Plot of $\tau_{250}(t)$ for the two series using a 250-day rolling window ending at date t

References

Alexander C, 2005
Correlation in crude oil and natural gas markets
In Managing Energy Price Risk: The New Challenges and Solutions, 3rd edition, Chapter 17, pages 573–604. Risk Books, London

Azzalini A and A Capitanio, 2003
Distributions generated by perturbation of symmetry with emphasis on a multivariate skew t-distribution
Journal of the Royal Statistical Society, Series B, 65, 367–389

Bollerslev T, 1996
Generalized autoregressive conditional heteroskedasticity
Journal of Econometrics, 31, 307–327

Breymann W, A Dias and P Embrechts, 2003
Dependence structures for multivariate high-frequency data in finance
Quantitative Finance, 3, 1–14

Capéraà P and C Genest, 1993
Spearman's rho is larger than Kendall's tau for positively dependent random variables.
Journal of Nonparametric Statistics, 2, 183–194

Cherubini U, E Luciano and W Vecchiato, 2004
Copula Methods in Finance – Wiley, New York

Deheuvels P, 1979
La fonction de dépendance empirique et ses propriétés: Un test non paramétrique d'indépendance
Académie Royale de Belgique, Bulletin de la Classe des Sciences (5), 65, 274–292

Demarta S and A McNeil, 2005
The t copula and related copulas
International Statistical Review, 73, 111–129

Fang H-B, K-T Fang and S Kotz, 2002
The meta-elliptical distributions with given marginals
Journal of Multivariate Analysis, 82, 1–16

Genest C and A-C Favre 2007
Everything you always wanted to know about copula modeling but were afraid to ask
Journal of Hydrologic Engineering, 12, 347–368

Genest C and B Rémillard, 2008
Validity of the parametric bootstrap for goodness-of-fit testing in semiparametric models
Annales de l'Institut Henri Poincaré, Probabilités et Statistiques, 44, in press

Genest C, B Rémillard and D Beaudoin, 2008
Goodness-of-fit tests for copulas: A review and a power study
Insurance: Mathematics and Economics, 42, in press

Joe H, 1997
Multivariate Models and Dependence Concepts
Chapman and Hall, London

McNeil A, R Frey and P Embrechts, 2005
Quantitative Risk Management: Concepts, Techniques and Tools
Princeton University Press, Princeton

Nelsen R, 1999
An Introduction to Copulas
Springer, New York

Patton A (in press)
Copula-based models for financial time series
In Handbook of Financial Time Series, Andersen T, R Davis, J-P Kreiss and T Mikosch (Eds), Springer, New York

Sklar A, 1959
Fonctions de répartition à n dimensions et leurs marges
Publications de l'Institut de statistique de l'Université de Paris, 8, 229–231

van den Goorbergh R, C Genest and B Werker, 2005
Bivariate option pricing using dynamic copula models
Insurance: Mathematics and Economics, 37, 101–114