DYNAMIC COPULA MODELS FOR THE SPARK SPREAD

FRED ESPEN BENTH AND PAUL C. KETTLER

ABSTRACT. We propose a non-symmetric copula to model the evolution of electricity and gas prices by a bivariate non-Gaussian autoregressive process. We identify the marginal dynamics as driven by normal inverse Gaussian processes, estimating them to a series of observed UK electricity and gas spot data. We estimate the copula by modeling the difference of the empirical copula to the independent copula. Following we simulate the joint process and price options written on the spark spread. We find that options prices are significantly influenced by the copula and the marginal distributions, along with the seasonality of the underlying prices.

1. INTRODUCTION

The spark spread measures the price spread between electricity and gas, and is of importance in the energy markets. Specifically, the spark spread as a function of time is defined as

$$S(t) = E(t) - H_r G(t), \tag{1.1}$$

where $E(t)$ and $G(t)$, respectively, are electricity and gas spot prices quoted in customary units. The constant H_r is called the *heat rate*, chosen to make approximate equivalence between the energy content of the two sources, adjusted by a factor consequential to the lesser efficiency of gas in typical applications. From a risk management point of view producers may use options written on the spark spread to hedge their price risks. Also, the spark spread option can be used to find the value of a gas-fired power plant project. The price of a spark spread option will inevitably depend on the joint distribution of electricity and gas prices. The aim of this paper is to propose a model for the joint probabilistic behavior of the prices, and apply this to investigate the consequences in pricing of spark spread options.

In particular, we propose a dynamic model for the prices of the two energies, including seasonality, mean reversion, non-Gaussianity, and a dependency structure. Marginally, each energy will be modeled by the exponential of an AR(1) process with a seasonally varying level and non-Gaussian innovations. The innovations of electricity and gas are next made dependent by a copula. Suggested from empirical studies of price data collected at the UK market, we model the innovations marginally by a normal inverse Gaussian distribution. The data show a dependency pattern which may be modeled by a certain one-parameter

Date: 19 October 2009.
2000 *Mathematics Subject Classification.* 62M10, 91B32, 91B70.
1991 *Journal of Economic Literature Subject Classification*. C51, G13.
Key words and phrases. copulas; mean reversion; seasonality; spark spread options.
The authors wish to thank Jūratė Šaltytė-Benth for providing us with estimates, Jan Kallsen for valuable discussions, and Heren Energy, Ltd., for the provision of the data used in this study. Two anonymous referees are thanked for their constructive criticism which led to a significant improvement of the presentation in this paper.

non-symmetric copula. The proposed copula, subtracted the independent copula, has a tent-shaped form in one coordinate and parabolic in the other. Our suggested copula model fits the data very well, and seems to be new in the literature.

The normal inverse Gaussian [NIG] distribution has been used successfully to model the returns of financial asset prices in different stock markets. See, *e.g.*, (Barndorff-Nielsen 1998; Raible 2000; Rydberg 1997). See (Benth and Šaltytė-Benth 2004) for a discussion of modeling gas spot prices in conjunction with a mean-reverting model, where it is demonstrated that the innovations in the gas price dynamics are very well explained by the NIG. Further, see (Börger, Cartea, Kiesel, and Schindlmayer 2007) for an analysis of the multivariate NIG distribution applied to forward price returns in a multi-commodity market, including electricity and gas.

Using copulas for modeling dependency between financial asset prices is gaining much attention. See, *e.g.*, (Cherubini, Luciano, and Vecchiato 2004) and the references therein. However, not much attention has been given to the energy markets. Copulas provide a very flexible tool for modeling the dependency between asset prices, in particular the returns. Due to the complexity of energy markets it is to be expected that the dependency structures are non-trivial, and far from being explainable by simple linear correlation. This view is confirmed by our data analysis.

Our marginal dynamics are a discrete-time version of the one-factor Schwartz model. In (Lucía and Schwartz 2002), the Schwartz model was applied to model electricity spot prices at the Nordic power exchange Nord Pool, assuming normal innovations. A generalization of this model including jumps is given in (Cartea and Figueroa 2005), where they study the UK electricity market. See also (Benth, Ekeland, Hauge, and Nielsen 2003) for a similar model applied to the Nord Pool market. Based on the studies by (Benth and Šaltytė-Benth 2004) for gas spot prices, we use a bivariate Schwartz model as the starting point for a joint modeling of the electricity and gas prices. Admittedly, this model has several shortcomings, one being the inability to model the spikes appropriately. Spikes are frequently observed in electricity spot prices, and sometimes also in gas prices. A more sophisticated mean-reverting model, with state-dependent jumps, was proposed in (Geman and Roncoroni 2006) and successfully fitted to many electricity price series. This model is able to capture the spikes, however, at the cost of being more complex for analytical purposes. A model for the gas spot prices with stochastic market price of risk is proposed in (Cartea and Williams). We believe that our model, although very simple, still captures many of the most important *distributional* properties of the prices, being essential for the valuation of spark spread option prices.

Our studies show that spark spread option prices are sensitive to the introduction of a copula. We compare the prices yielded by our model with those from a binormal model with linear correlation, and we observe generally an increase in price for out-of-the-money options, and the opposite for in-the-money options. The differences are close to negligible for far in-the-money options, but may become significant for far out-of-the money options. There is a clear sign of a combined effect of seasonality and heavy tails implied by the NIG and the copula chosen. Comparing with the independent copula, the model with NIG innovations will give slightly higher prices than when a copula is introduced. Thus, the variation is less when having a copula modeling the dependency structure. One may claim that supposing NIG innovations where the tails are heavy the option prices necessarily must increase over the binormal case. This is not so obvious, since the spark spread is a difference in electricity and gas prices, and with a dependency structure one could picture that the spark spread distribution becomes more concentrated. In fact, the NIG is more concentrated in the center than the normal for our parameter estimates. However, our findings tell that the prices indeed increase compared to the binormal case. Several authors have treated the problem of modeling and pricing spark spread options, where we here only list a few. For a broad presentation on the spark and other spreads, with stochastic analysis including mean reversion, jump terms, and attention to seasonality see (Carmona and Durrelman 2003), wherein the authors base their considerations on a multivariate Schwartz model, possibly including jumps. In (Benth and Šaltytė-Benth 2006) it is suggested to model the spark spread dynamics explicitly by a mean-reverting jump process, for which explicit option prices are obtained through Fourier transforms. This study generalizes some of the approximation results in (Carmona and Durrelman 2003) on the difference of two lognormal variables. In (Dempster and Hong 2000) the fast Fourier transform method is considered for pricing of different spread options.

This paper is structured as follows. In the next section we introduce our dynamic model for the joint electricity and gas price evolution, including the theoretical copula we propose as the dependency structure between the two energies. Based on UK spot price data of electricity and gas, we fit and validate our model in Section 3, following up with a simulation study pricing different call and put options written on the spark spread in the following section. Finally, we conclude.

2. A COPULA MODEL FOR ELECTRICITY AND GAS SPOT PRICE DYNAMICS

Let $(\Omega, \mathcal{F}, P)$ be a probability space equipped with the filtration $\{\mathcal{F}_t\}_{0\leqslant t\leqslant T}$ which is assumed to satisfy the *usual conditions*. In (Benth and Šaltytė-Benth 2004) it was argued that a reasonable model for the spot price dynamics $G(t)$ of gas is given as the exponential of a non-Gaussian Ornstein-Uhlenbeck process. More specifically, it was assumed that $G(t) = \Lambda(t)\exp\big(X(t)\big)$ where $\Lambda(t)$ is a deterministic seasonal function and $X(t)$ has the dynamics

$$dX(t) = -\alpha\,(X(t) - \mu)\,dt + dL(t) \tag{2.1}$$

In particular, an empirical study revealed that the $L(t)$ should be modeled as a Lévy process with normal inverse Gaussian distributed marginals.

In this paper we work with stochastic processes in discrete time, and we propose and analyze a two-dimensional stochastic process $\big(E(t), G(t)\big)$ as a joint model of electricity and gas spot prices. Time t is measured on a discrete set $\{t_0, t_1, \dots,\}$ with the convention $t_{i+1}-t_i = 1$, being typically hours or days. Motivated by Equation (2.1), we suppose that electricity and gas spot prices have the dynamics

$$E(t) = \Lambda_E(t)\exp\big(X(t)\big) \tag{2.2}$$

and

$$G(t) = \Lambda_G(t)\exp\big(Y(t)\big), \tag{2.3}$$

respectively, for two deterministic seasonal functions Λ_E, Λ_G. Furthermore, $X(t)$ and $Y(t)$ are two AR(1)-processes with non-Gaussian innovations $\varepsilon_E, \varepsilon_G$ defined by

$$X(t_{i+1}) = \mu_E + \alpha_E X(t_i) + \varepsilon_E(t) \tag{2.4}$$

and

$$Y(t_{i+1}) = \mu_G + \alpha_G Y(t_i) + \varepsilon_G(t) \tag{2.5}$$

Here, μ_E, μ_G, α_E, and α_G are constants, with $\alpha_{\{E,G\}}$ being positive and between 0 and 1 in value. Moreover, the innovations are independent and identically distributed in time, however, possibly being dependent. The standard choice of distributions is the normal one. Motivated from the data analysis in Section 3 we shall specify the distribution of $\varepsilon_{\{E,G\}}$ to be NIG.

The NIG distribution is a four-parameter family of distributions which has been used frequently to model financial data. See, *e.g.*, (Barndorff-Nielsen 1998; Raible 2000; Rydberg 1997). The density function of the NIG distribution is given as

$$f(x; \alpha, \beta, \mu, \delta) = a(\alpha, \beta, \mu, \delta) q^{-1}\left(\frac{x-\mu}{\delta}\right) K_1\left[\delta\alpha q\left(\frac{x-\mu}{\delta}\right)\right] e^{\beta x},$$

where

$$q(x) = \sqrt{1+x^2}, \qquad a(\alpha, \beta, \mu, \delta) = \frac{\alpha}{\pi}\exp\left[\delta\sqrt{\alpha^2-\beta^2} - \beta\mu\right],$$

and $K_1(x)$ is the modified Bessel function of the second kind of index one. The parameters μ, α, β and δ satisfy the following relations

$$\alpha > 0, \quad |\beta| \leqslant \alpha, \quad \delta > 0$$

The first parameter, μ, is the location of the NIG distribution, whereas β is the skewness and δ is the scale parameter. Finally, α measures the tail heaviness of the distribution. The formulas for the mean, variance, skewness, and kurtosis (denoted as (m, ν, s, k), respectively) for a NIG distributed random variable with parameters $(\alpha, \beta, \mu, \delta)$ are given as in (Rydberg 1997)).

$$m = \mu + \delta\frac{\beta}{\gamma} \qquad \nu = \delta\frac{\alpha^2}{\gamma^3}$$

$$s = 3\frac{\beta}{\alpha}\frac{1}{(\delta\gamma)^{\frac{1}{2}}} \qquad k = 3\left[1+4\left(\frac{\beta}{\alpha}\right)^2\right]\frac{1}{\delta\gamma}$$

For convenience, we have used the notation $\gamma = \sqrt{\alpha^2-\beta^2}$. These formulas are invertible, enabling us to express the NIG parameters in terms of the first four moments: In the following formulas, we have denoted $\zeta = 3k - 4s^2$ and $\eta = k - \frac{5}{3}s^2$.

$$\alpha = \frac{\sqrt{\zeta}}{\eta\sqrt{\nu}} \qquad \beta = \frac{s}{\eta\sqrt{\nu}}$$

$$\mu = m - \frac{3s\sqrt{\nu}}{\zeta} \qquad \delta = \frac{3^{\frac{3}{2}}\sqrt{\eta\nu}}{\zeta}$$

In the empirical analysis we apply these formulas for a method-of-moment estimation approach to fit the NIG distribution to the observed innovations ε_E and ε_G. For background on the NIG distribution see (Rydberg 1997; Barndorff-Nielsen 1998). For a comprehensive treatment of NIG distributions as they relate to Lévy processes see (Raible 2000).

Next, we introduce a copula for modeling the dependency structure between the innovations of electricity and gas, $\varepsilon_E(t)$ and $\varepsilon_G(t)$. In general, a function $C\colon [0,1]^2 \mapsto [0,1]$ is called a *copula* if it is *2-increasing* and *grounded*. A function $\widetilde{C}\colon [0,1]^2 \mapsto [0,1]$ is called 2-increasing when for every rectangle $[u_1,u_2] \times [v_1,v_2]$ with vertices in $[0,1]^2$, and such that $u_1 \leqslant u_2, v_1 \leqslant v_2$,

$$\widetilde{C}(u_2,v_2) - \widetilde{C}(u_2,v_1) - \widetilde{C}(u_1,v_2) + \widetilde{C}(u_1,v_1) \geqslant 0$$

Hence, a 2-increasing function assigns positive mass to every rectangle in its domain. Further, $\widetilde{C}$ is called grounded on $[0,1]^2$ if

$$\widetilde{C}(0,v) = \widetilde{C}(u,0) = 0$$

A copula C is a bivariate uniform distribution function. For a good foundation of copula theory, we refer to (Nelsen 1998). In this paper we propose a non-symmetric bivariate copula to model the dependency between electricity and gas. The copula has component being tent-shaped in one dimension, whereas being a parabola in the other, and having one parameter only. It is proposed on the background of empirical analysis of data to come, but we include the definition of it here along with some analysis of its properties.

Define the function $C_h\colon [0,1]^2 \longmapsto [0,1]$ for a parameter h (indicating the 'height' of C_h) as

$$C_h(u,v) = uv + h(1-|2u-1|)\left(1-(2v-1)^2\right) \tag{2.6}$$

We observe that the function C_h is a sum of the *independent copula* $C_\perp(u,v) = uv$ and a function

$$D(u,v) = (1-|2u-1|)\left(1-(2v-1)^2\right), \tag{2.7}$$

that is,

$$C_h(u,v) = C_\perp(u,v) + hD(u,v)$$

It turns out that C_h is only a copula for an interval of admissible parameter values of h. We now discuss this in more detail. We investigate the range of admissible values of h yielding a copula C_h. First, note that

$$c_h(u,v) = \frac{\partial^2 C_h}{\partial u \partial v} = 1 + hd(u,v),$$

where

$$d(u,v) = \frac{\partial^2 D}{\partial u \partial v} = 8\,\mathrm{sgn}(2u-1)(2v-1)$$

Hence, $c_h(u,v)$ is a density function on $[0,1]^2$ if it is non-negative and has mass equal to one over the unit square. It is easily seen that for

$$h \in [-1/8, 1/8], \tag{2.8}$$

we find that $c_h(u,v)$ is non-negative on $[0,1]^2$. Furthermore, since the mass of $d(u,v)$ is zero over the unit square, we conclude that $c_h(u,v)$ is a density function. Hence, $C_h(u,v)$ is 2-increasing. Finally, note that $D(u,v) = 0$ on the boundaries of $[0,1]^2$, which implies that $C_h(u,v)$ is grounded. Therefore, we have that $C_h(u,v)$ is a copula for $-1/8 \leqslant h \leqslant 1/8$. Data analysis will suggest that the innovation process $\left(\varepsilon_E(t), \varepsilon_G(t)\right)$ is well represented as *i.i.d*, with NIG distributed margins and a dependency structure described by the copula $C_h(u,v)$.

Much of the literature on derivatives pricing is based on continuous-time stochastic processes, and in our context it would have been natural to start with a bivariate non-Gaussian Ornstein-Uhlenbeck process, where the marginal deseasonalized log-spot prices are of the form of Equation (2.1). Electricity and gas would be driven by two Lévy processes $L_E(t)$

	a	b	c	ϕ
Electricity	2.79	0.000361	0.1603	−49.182
Gas	2.86	0.000116	0.2748	−21.4688

TABLE 1. Estimates of the seasonal function defined in (3.1) for electricity and gas log-spot price data

and $L_G(t)$, each having NIG-distributed increments, say. The next step would then be to model the dependency of the two Lévy processes using a copula C. However, as shown by (Kallsen and Tankov 2004), one can not use a standard copula to link two Lévy processes, and preserve the Lévy property for the two-dimensional process. First, the copula needs to be introduced for a prescribed time step, and it is not preserved for all other increments in time of the Lévy processes. Secondly, the stochastic properties of the two-dimensional process does not seem to be clear, and the analytical tractability of the model is not straightforward. This question is resolved in (Kallsen and Tankov 2004) by introducing a so-called *Lévy copula*, which means that the Lévy measures (jump measures) are coupled by a function sharing similar properties as a copula. In this way one can define a two-dimensional Lévy process. However, the connection to the statistical methods and intuition on copulas is lost, and one must analyze data in frequency space instead. In order to keep the connection to "statistical" copulas as introduced above, we investigate time series models in this paper. The continuous-time process related to our suggested model is an open, yet very interesting question.

We now move on to the empirical analysis of electricity and gas data, and fitting of our proposed model.

3. THE DATA AND MODEL FITTING

The data of this study are 805 parallel observations of daily spot prices for electricity and gas from the United Kingdom, 6 February 2001 through 26 April 2004. On average we had 252 price data a year. Of these observations seven were considered as outliers and truncated. The estimation of the spot price model $E(t)$ and $G(t)$ defined in Equations (2.3)-(2.2) goes as follows. First, the price data are transformed to a logarithmic scale, where a seasonal function is estimated for both electricity and gas. Next, we fit an AR(1) process using a linear regression procedure. The residuals from this regression are used for estimating the parameters in the NIG distribution supposed for the innovations marginally. Finally, we fit our proposed copula to these residuals.

3.1. **Estimation of the marginal models.** We choose the seasonal functions $\Lambda_E(t)$ and $\Lambda_G(t)$ to be of the form

$$\log \Lambda(t) = a + bt + c\cos(2\pi(t - \phi)/252), \tag{3.1}$$

where the four parameters a, b, c and ϕ are estimated for logarithmic electricity and gas spot prices, respectively. Using a least squares approach, we reach the estimates in Table 1.

In Figure 1 we plot the data series along with the fitted seasonal functions on log-scale. Although the plot shows a reasonable fit to the deterministic variations of the data, the estimation of a seasonal function to power prices is a delicate issue. See (Geman and Roncoroni

2006) for applying a threshold technique where the spikes are filtered out. The spikes frequently encountered in power prices may lead to a bias in the estimation of the seasonal function. This is true for data series with many spikes, but is not obvious in data series where spikes are less frequently occuring. In (Benth, Kiesel, and Nazarova 2009) the threshold method of (Geman and Roncoroni 2006) was applied to spot prices from the German power exchange EEX where only a small effect was detected from the spike influence. We believe that the seasonal function estimated for the gas and electricity data considered in this paper is sufficiently reflecting the variations observed. We also refer to (Cartea and Figueroa 2005) for more on seasonality estimation.

After subtracting the estimated seasonal function from the log-spot data, we are left with two data series which we recall from Equations (2.5) and (2.4) are assumed to follow a non-Gaussian AR(1) model each. We regressed next-day data against today's data, and found autoregression coefficients equal to $\widehat{\alpha}_E = 0.7278$ and $\widehat{\alpha}_G = 0.8849$. The intercepts μ_E and μ_G were found to be insignificant and therefore put to zero. All these estimates were kindly provided to us by Jūratė Šaltytė-Benth.

We now analyze empirically the residuals of electricity and gas resulting from the regression. In Figure 2 we show the Q-Q plot of electricity residuals. A qualitative interpretation reveals leptokurtosis insofar as the median frequency of the electricity variate is greater than the median frequency of the normal variate, *i.e.*, the slope in the center is greater than one. The same phenomenon obtains in the tails, giving a fat tail interpretation to the electricity data. In contrast, the Q-Q gas function depicted in Figure 3, while also exhibiting leptokurtosis, shows distinctly thin tails.

This comparative distinction appears directly in the cross reference electricity-gas Q-Q function plotted in Figure 4, which in the center reveals the electricity variate as slightly more leptokurtic than the gas variate. As well, the dramatic slopes in the tails display the combined effect of fat electricity tails and thin gas tails.

Considering the three P-P functions — electricity, gas, and cross reference — the increase in ordinate for each successive point is the same, namely $1/n$. Thus a projection of these points on the vertical axis is an instance of the uniform distribution. See Figures 5, 6, and 7.

A qualitative interpretation of the P-P electricity function reveals two evident features. First is an early and steep traverse of the equivariate line approaching the center of the plot from the left. In consequence, the median of the electricity residuals corresponds to a significantly higher fractile in the comparing normal distribution. This sighting is consistent with right skewness in the electricity variate, along with excess kurtosis corresponding to the steep slope. Second is the late traverse in the opposite direction at the right edge of the plot. This pattern shows that the fat right tail, to accompany the excess kurtosis, diminishes to a very thin tail in the extreme; it is as if there were some institutional proscription militating against large upward jumps in the process. The bidding system possibly plays a role here, as could explicit or implicit intraday price move limits.

In contrast, the P-P gas function is much more regular, though decidedly non-Gaussian. The medians of the variables are insignificantly different and skewness is not apparent, though kurtosis and fat tails clearly obtain.

With the cross reference P-P electricity-gas function we see distributions which are not as far from each other as either is to the normal distribution; however, the early traverse of the equivariate line still shows the comparative right skewness of the electricity variate, whereas

	μ	α	β	δ
Electricity	−0.0694	6.93	2.80	0.151
Gas	−0.0122	7.77	−0.998	0.0831

TABLE 2. Estimated parameters of the NIG distribution for electricity and gas innovations

the relatively thin right electricity tail is revealed to an even greater extent when compared to the relatively fat tail in the gas variate.

These empirical findings point toward a non-Gaussian distribution for modeling the innovations. The NIG has shown to be a flexible family of distribution, which by now is frequently used for modeling financial returns. We apply the method of moments, in particular the formulas in the previous section linking the first four empirical moments to the four NIG parameters. The estimates are presented in Table 2. We have plotted the empirical density against the fitted NIG distribution in Figures 8 and 9 for electricity and gas, respectively. The NIG distribution seem to fit both data set reasonably well.

3.2. **Estimation of the copula.** Before estimating our proposed copula model C_h in Equation (2.6), we investigate the empirical copula for the residuals. Moreover, in order to reveal the dependence structure, we consider the empirical copula after subtracting the independent copula, that is, the *empirical difference copula*. Recall that our model for this is $D(u,v) = C_h(u,v) - C_\perp(u,v)$ in Equation (2.7).

In practice we calculated $D(u,v)$ by first constructing $C_h(u,v)$. To this end we assigned a pair of ranks, integers $(\rho_1, \rho_2) \in \{1,\dots,n\}^2 = \{1,\dots,805\}^2$, to each data pair, and then figuratively placed the binary digit 1 on the unit square at each reduced rank pair $(u,v) = (\rho_1/n, \rho_2/n)$, all other grid positions having been initialized to 0. The empirical copula was generated simply by counting the bits for each grid position not greater in both variables, expressing the total as a fraction of n. Note that for the n observations there is exactly a single digit 1 in each row and in each column of the grid, providing the characteristic uniform marginal distributions of a copula. It was a simple matter then, finally, to subtract the value $C_\perp(u,v) = uv$ from $C_h(u,v)$ to arrive at $D(u,v)$.

Figures 10 and 11 show the empirical copula difference function, also from the electricity axis into the domain, and toward the gas axis out of the domain. What we see is an asymmetrical function, being close to tent-shaped in one direction and parabolic in the other. The qualitatively different view must owe to the different character of the separate markets for electricity and gas. Gas is storable, whereas electricity is not. The development of our model captures this distinction. Figure 12 shows the density of this empirical copula, a "top view," of which Figures 10 and 11 are the "side views."

We estimate the parameter h in C_h (recall Equation (2.8) by using least squares, and find $\widehat{h} = 0.0848$. This estimate lies well within the interval $[-1/8, 1/8]$ of admissible parameter values for C_h, and thus the estimated function $C_h(u,v)$ is a copula. The coefficient of determination is rather high, being $r^2 = 0.8949$, implying a reasonably good fit of the model. Figures 13 and 14 show an estimated theoretical copula difference function as in Equation (2.7), evaluated on the domain of our sample. The views are first from the electricity axis into the domain, and then toward the gas axis out of the domain.[1]

[1]We take these views to keep the origin on the left in each case.

In connection with copula modeling, the question of tail dependence is considered. We recall the definition of tail dependency. A bivariate distribution $F(x, y)$ of the random variable (X, Y) with margins F_X and F_Y, is *lower tail dependent* with coefficient λ_L if

$$\lambda_L := \lim_{a \to -\infty} \Pr\left(Y \leqslant a \mid X \leqslant a\right) = \lim_{a \to -\infty} \frac{F(a, a)}{F_X(a)}$$

Letting the copula of the bivariate distribution be $C(u, v)$, we can state the lower tail dependency in terms of this:

$$\lambda_L = \lim_{u \downarrow 0} \frac{C(u, u)}{u}$$

We say that (X, Y) is *upper tail dependent* with coefficient λ_U if $(-X, -Y)$ is lower tail dependent with that coefficient. Note that in this case we have

$$\lambda_U = \lim_{v \uparrow 1} \frac{1 - 2v + C(v, v)}{1 - v}$$

The numerator on the right-hand side is often called the *survival function* of the copula C.

When the data of this study are examined for tail dependence an interesting pattern emerges. First, bear in mind that with a finite data set it is impossible to examine limit behavior in any rigorous context. This is especially so with the relatively small bivariate sample of 798 points as herein. Nonetheless, it is possible to say something.

As an *ad hoc* test we look to the first and last deciles in each variable, and specifically seek the points which are in both lower or in both upper deciles. We test then on the null hypothesis that the distributions from which these data emerge are independent. With 798 points we would expect 7.98 in both lower and both upper deciles. In fact, the lower joint tail has 23 points, and the upper joint tail has 15 points, both significantly higher at the 1% level of confidence to reject the null hypothesis on a binomial test with continuity correction. In fact, the lower joint tail shows significant departure from the independent assumption at a much smaller fraction than 1%.

We should note here that the copula C_h we employ has zero tail dependence. We made a conscious decision not to include tail dependence in the theory, for as noted the number of points modeled is small and the fit overall to our choice of copula is excellent. However, since there are signs of tail dependency, one may wish to refine our copula model to include such a possibility. We believe that with an increasing number of data, it could be possible to reveal the true structure of the tail dependency, being very difficult in our situation.

4. Simulation of spark spread option prices

In this section we describe how to simulate paths for our proposed electricity and gas price dynamics based on the copula C_h, and apply the methodology to analyze prices of spark spread options.

Drawing samples from the copula C_h is a straightforward process using the standard approach of inverting the conditional marginals of the copula. See, *e.g.*, (Nelsen 1998). We recall this procedure here in the special case of C_h. Suppose that (U, V) are bivariate uniform with distribution given by the copula $C_h(u, v)$. Then we find that

$$C_h^v(u) := \Pr\left\{U \leqslant u \mid V = v\right\} = \frac{\partial}{\partial v} C_h(u, v) = u - 4h(1 - |2u - 1|)(2\bar{v} - 1)$$

Next, start with two independent draws (u_1, u_2) from the uniform distribution, and let $v = u_2$. To find the draw u, we solve the equation $u_1 = C_h^v(u)$, which results in

$$u = \begin{cases} z, & \text{if } z \leqslant \frac{1}{2} \\ \dfrac{u_1 + \widetilde{v}}{1 + \widetilde{v}}, & \text{if } z > \frac{1}{2} \end{cases}$$

Here, we have introduced the notation $\widetilde{v} = 8h(2v - 1)$ and $z = u_1/(1 - \widetilde{u})$. This gives us a pair (u, v) drawn from C_h. Figure 15 shows a simulation of 2000 of such draws.

By inversion of the NIG distribution, we can feed in a draw (u, v) from the copula C_h to find a draw from a distribution with NIG marginals and dependence structure given by the copula. For the parameters given from electricity and gas, we performed a simulation of 1000 pairs. The scatter plot of this exercise is given in Figure 16. For comparison we simulated the same number of pairs from a bivariate normal distribution, using the means and covariances of the samples from the NIG-copula simulation. The resulting scatterplot is found in Figure 17. As we can see, there is a significant difference between the two simulations. The NIG-copula simulation shows the effect of both the NIG distribution marginally and the copula jointly over the bivariate normal. The latter simulation corresponds to the assumption that the innovation processes $\big(\varepsilon_E(t), \varepsilon_G(t)\big)$ are supposed to be *i.i.d* bivariate normally distributed.[2]

Based on the simulation of the innovation process $\big(\varepsilon_E(t), \varepsilon_G t)\big)$ just described, we use the dynamical model of electricity and gas spot prices in Equations (2.3) and (2.2) to generate sample paths. We also generated paths based on the bivariate normal. As an illustration, we compare in Figures 18 and 19, the distribution of electricity and gas using the NIG assumption with the normal after 20 days.

The comparisons are based on 1000 simulations of the terminal distributions. Since our main application is the study of the spark spread, we also plot in Figure 20 the distribution of the spark spread after 60 days corresponding to a relative efficiency of gas to electricity of 40%, that is, $H_r = 0.853$. See the Appendix for details. The outcomes of the spark spread are in terms of £/MWh.

Let us explain these charts in more detail. The charts show the difference of the two distribution samples, point by point, on the vertical axis, compared to sample point of the second variable on the horizontal axis. Thus to recover the point of the first variable one adds the abscissa and ordinate. This mode of presentation is possible owing to the coordinated draws from the uniform distributions generating the price paths. This explicit choice, made to avoid all manner of induced Monte Carlo errors, now bears fruit in the form of permitting this visualization without having to introduce other errors, for instance, from interpolation of the plots.

The insight to gain is that these three plots are roughly serpentine shaped, informing us that the pattern is for the extreme copula-driven NIG-based paths to terminate farther from the mean than the corresponding binormal-based paths, whereas the central NIG paths terminate more closely to the mean than their normal counterparts. This is the pattern of the leptokurtic NIG distribution compared with the normal. Observe that the gas data, being

[2]From the observed innovations, we estimated the mean and standard deviation for the electricity data to be -0.0025 and 0.17, whereas the corresponding estimates for gas were 0.0014 and 0.10. The correlation coefficient is 0.27.

more regular than the electricity data, produce a clearer pattern, whereas the spread data show regularity between that of the electricity and gas.

We move on to the problem of pricing options written on the spark spread. The arbitrage-free price of a spark spread call option with maturity T and strike K is

$$C = e^{-rT} E_Q \left[\max \left(S(T) - K, 0 \right) \right]$$

where Q is a risk-neutral probability and r is the constant risk-free interest rate. Since electricity is non-storable and it is costly to store gas, spark spread options can not be hedged, implying that the market is so-called incomplete.

There exist in general many risk-neutral probabilities, meaning that there is no unique theoretical arbitrage-free price. As is standard in such situations, one introduces a parameter (possibly time-dependent) adjusting the level to which the processes $E(t)$ and $G(t)$ mean revert. This parameter, being two-dimensional, is often interpreted as a market price of risk, and can be estimated from observed option prices quoted in the market.

If we have a model based on normally distributed innovations, the introduction of the market price of risk can be viewed (from the continuous-time perspective) as resulting from a standard Girsanov transformation. For other choices of the innovation processes, one may resort to the Esscher transform. This transform will preserve the distributional properties of the innovation process, but possibly change the parameters. In addition, the seasonal level will be adjusted with the interpretation of being a market price of risk. Since we have no access to market quotes on spark spread options, it is not straightforward to specify any market price of risk. One possibility would be to estimate it from traded forward/futures prices. However, this approach would introduce additional complications like establishing the link between forward and spot prices in the electricity market, and to reveal the correct convenience yield or storage cost structure in the gas market. See (Benth, Cartea, and Kiesel 2008) and (Benth and Meyer-Brandis 2009) for analyses of the spot-forward relation in electricity markets. See (Geman 2005) for a discussion on forward pricing in gas markets.

There are also clear liquidity issues in these markets, complicating such analysis. Furthermore, it is not clear how to parametrize the market price of dependency risk. We leave these issues for future studies, and focus here on the possible effect of the copula model on prices where the risk premium is ignored. Obviously this will not reflect the market situation completely, but will reveal the significance of dependency risk on prices compared to other sources of risk.

In view of the above discussions on risk premium issues in these markets we choose $Q = P$ in order to pursue our study of the influence of our model on prices. Note that from the non-tradeability of the underlying processes in the spark spread, one does not need to require that the discounted spot prices to be martingales under Q, validating that the objective measure can be used as a pricing measure. The choice $Q = P$ can also be regarded as deriving the predicted payoff from the option, rather than the price, which sometimes is an interesting figure by itself. For example, if option prices are available, one can use this method to analyse the risk premium in the market. Based on these arguments, we make the choice $Q = P$ in the further analysis of spark spread prices. This choice is also advantageous, because it is neutral with respect to the specification of the marginal distributions and copula, and it is easier to reveal possible effects coming from these. We suspect, however, that choosing Q by Esscher or Girsanov transforms will not alter our results and conclusions in any significant way, except for giving different values, of course.

We make the further simplification that $r = 0$, which is not essential for the discussion of the results to come. The prices of call options are now straightforward to calculate based on Monte Carlo simulations of the price paths up to exercise time for the electricity and gas spot prices. The simulation procedure described above is used. We also calculate put option prices, which is based on the put-call parity

$$C - P = E\left[\max\left(S(T) - K, 0\right)\right] - E\left[\max\left(K - S(T), 0\right)\right] = E\left[S(T)\right] - K$$

We simulated call and put prices for different strike levels K and exercise times T for our proposed NIG-copula model based on the parameter estimates found in Section 3. To have an idea of the influence of the copula on the option price, we included simulations assuming independence between the electricity and gas innovation processes, but still having NIG distributions. The bivariate case was also considered, to compare the mixed effect of NIG and copula assumptions in our proposed dynamics with the standard model based on normal distributions and correlation. The strikes are chosen to be $\{-10, -5, 0, 5, 10\}$, where we note that $K = 0$ with a normal assumption on the innovations can be calculated explicitly by using a modification of Margrabe's formula. See (Carmona and Durrelman 2003). Further, we consider 4 maturity times, being the first 4 quarters, that is, $T \in \{60, 120, 180, 240\}$ days. We refer to these as periods 1, 2, 3 and 4. The results for call options are reported in Table 3, while the put options prices for the three different model assumptions are listed in Table 4.

From the tables we observe that differences in prices between the NIG and binormal models are relatively small, except for out-of-the money options where the binormal assumption are nearly worthless, whereas the same options under the NIG assumptions have some value. We also see differences in prices between the NIG with fitted and independent copula, showing that introducing a dependency structure has an impact on pricing. Overall, the independent case gives higher prices, which is not surprising since the fitted copula introduces a dependency which makes the variation of the spark spread smaller than for the independent case. For in-the-money options, the difference between the two NIG assumptions is negligible; however, for options moving toward out-of-the money it becomes gradually more significant. Considering for instance a call option with strike $K = 0$ in period 1, we see that the independent case yields a price increase of about 20%. Tracing through the different maturities, we see that the price increase from copula to independent case is varying significantly for $K = 0$; the second period has only 2.2% increase, with an even lower figure 0.4% for period 3. The last period is bigger again, with 9%. Thus, the seasonality of prices plays a strong role, as well.

The spark spread option study tells us that the seasonality of the price dynamics plays a significant role in the formation of option prices, along with the introduction of a copula and heavy tailed NIG distributions. We emphasize that the introduction of our proposed copula yields a price *reduction* compared to the independent case. Furthermore, the introduction of a bivariate heavy-tailed distribution (NIG) in the price dynamics yields higher prices than the light-tailed binormal assumption. At first sight this may seem reasonable, but on the other hand we are looking at the price difference between electricity and gas when pricing a spark spread option, which could potentially cancel out heavy tailed effects and therefore make variation — and thus the option price — smaller in conjunction with a dependency structure. This is not the case, as our results show.

We clearly see how seasonal variation in spot prices influences the option prices by looking at the variation in Tables 3–4 over the quarters of the year (corresponding to the four periods we investigate.) To quantify this effect more precisely, one could perform a study

Strike	Period 1	Period 2	Period 3	Period 4	Process
-10	8.7248	12.4448	14.5029	11.3708	NIG, copula
-10	8.7216	12.4408	14.4981	11.3658	NIG, indep.
-10	8.7681	12.4869	14.5545	11.4260	binormal
-5	3.7967	7.4499	9.5079	6.3978	NIG, copula
-5	3.8470	7.4465	9.5033	6.4135	NIG, indep.
-5	3.8981	7.4880	9.5546	6.4728	binormal
0	0.6360	2.5609	4.5424	2.0509	NIG, copula
0	0.7640	2.6192	4.5603	2.2364	NIG, indep.
0	0.7152	2.6820	4.6235	2.3058	binormal
5	0.1013	0.3004	0.9295	0.4443	NIG, copula
5	0.1340	0.3585	1.0457	0.5415	NIG, indep.
5	0.0429	0.2720	1.0604	0.4170	binormal
10	0.0233	0.0442	0.1746	0.1134	NIG, copula
10	0.0371	0.0625	0.2058	0.1487	NIG, indep.
10	0.0008	0.0075	0.0921	0.0384	binormal

TABLE 3. Simulated spark spread call option prices for various periods and strikes, with NIG and binormal process assumptions. The NIG with fitted and independent copula, respectively.

Strike	Period 1	Period 2	Period 3	Period 4	Process
-10	0.0107	0.0013	0.0011	0.0105	NIG, copula
-10	0.0127	0.0021	0.0020	0.0122	NIG, indep.
-10	0.0033	0.0000	0.0000	0.0017	binormal
-5	0.0827	0.0065	0.0061	0.0376	NIG, copula
-5	0.1381	0.0077	0.0072	0.0599	NIG, indep.
-5	0.1333	0.0011	0.0001	0.0484	binormal
0	1.9220	0.1175	0.0406	0.6906	NIG, copula
0	2.0552	0.1804	0.0642	0.8828	NIG, indep.
0	1.9504	0.1950	0.0690	0.8815	binormal
5	6.3873	2.8569	1.4277	4.0840	NIG, copula
5	6.4251	2.9197	1.5496	4.1879	NIG, indep.
5	6.2781	2.7851	1.5058	3.9927	binormal
10	11.3092	7.6008	5.6728	8.7531	NIG, copula
10	11.3283	7.6237	5.7097	8.7951	NIG, indep.
10	11.2359	7.5206	5.5376	8.6140	binormal

TABLE 4. Simulated spark spread put option prices for various periods and strikes, with NIG and binormal process assumptions. The NIG with fitted and independent copula, respectively.

where the seasonal functions of the spot prices were artifically omitted. We refer to (Fusai, Marena, and Roncoroni 2008), where a complete study of the seasonality effects in the volatility structure for discretely-monitored Asian options are presented. Their findings seem to support our claim.

5. Conclusions

Joint electricity and gas price discovery is better modeled by NIG marginal processes with a specialized copula joining them than by a binormal model. The latter, more commonly used to model joint price processes, falls short on both the proper description of the marginal distributions, and on their dependence relationship. With this work we free ourselves of the Gaussian constraint to face more realistically the dynamics of the marketplace.

For option valuation the most significant finding is that out-of-the-money puts and calls are valued higher under the NIG than the binormal assumption, which we believe is an attribution of the fatter tails of the NIG. Introducing a copula yields lower prices than for the independent case, demonstrating the importance of having a reasonable model for the dependency structure. The seasonality also plays an important part in price formation. Otherwise option evaluation is robust across the distributional assumptions.

The model proposed in this paper may be extended in several different ways. Firstly, one may argue that spikes are not well-captured by our simple one-factor model, and more realistic jump terms are called for. Here one could think of multi-factor models, where we have several AR(1) processes driving the dynamics of electricity and gas. One or more of these factors could consist of innovations with rare, but big jumps, along with a strong mean reversion, mimicking the spikes. Another possibility could be to make the jumps state dependent, as is done in (Geman and Roncoroni 2006) for electricity prices.

Further, we observed some signs of tail dependency in the data, which calls for more sophisticated copula models than the one we propose here. A longer data series would enable stronger conclusions in this respect, and the question of tail dependency, could be settled, and possible extensions of our proposed copula could be suggested.

Finally, it would be interesting to have a continuous-time dynamics for the electricity and gas prices, with a copula reflecting the data. This calls for a closer study of the connection between Lévy copulas as introduced by (Kallsen and Tankov 2004) and the statistical copulas which we have discussed here. These possible extensions are left for future research.

APPENDIX: THE HEAT RATE

Customarily, the unit of electrical energy is the Mega Watt-hour (MWh), whereas the unit of gas energy is the Giga joule (GJ). The heat rate $\mathrm{H_r}$ is the dimensionless conversion rate of 3.6 GJ/MWh divided by the relative efficiency of gas compared to electricity, here assumed to be 40%. Thus, $\mathrm{H_r} = 9.0$ in our model. This conversion presumes that both kinds of energy are quoted in the same monetary units. In the marketplace, gas energy is commonly measured in therms, a therm being 105.5 MJ. Gas is quoted in pence/therm, whereas electricity is in £/MWh. The conversion for equal amounts of energy, not including a correction for relative efficiency, goes as follows.

$$1\,\frac{\text{penny}}{\text{therm}} \cdot \left[\frac{1\text{ therm}}{105.5\text{ MJ}} \cdot \frac{1000\text{ MJ}}{\text{GJ}} \cdot \frac{3.6\text{ GJ}}{\text{MWh}} \cdot \frac{\$1}{100\text{ pence}}\right] = 0.34123223\,\frac{\$}{\text{MWh}}$$

In this system of quotation the heat rate, now allowing for relative efficiency of 40%, is 0.85308057. For practical calculations these coefficients have at most two significant digits.

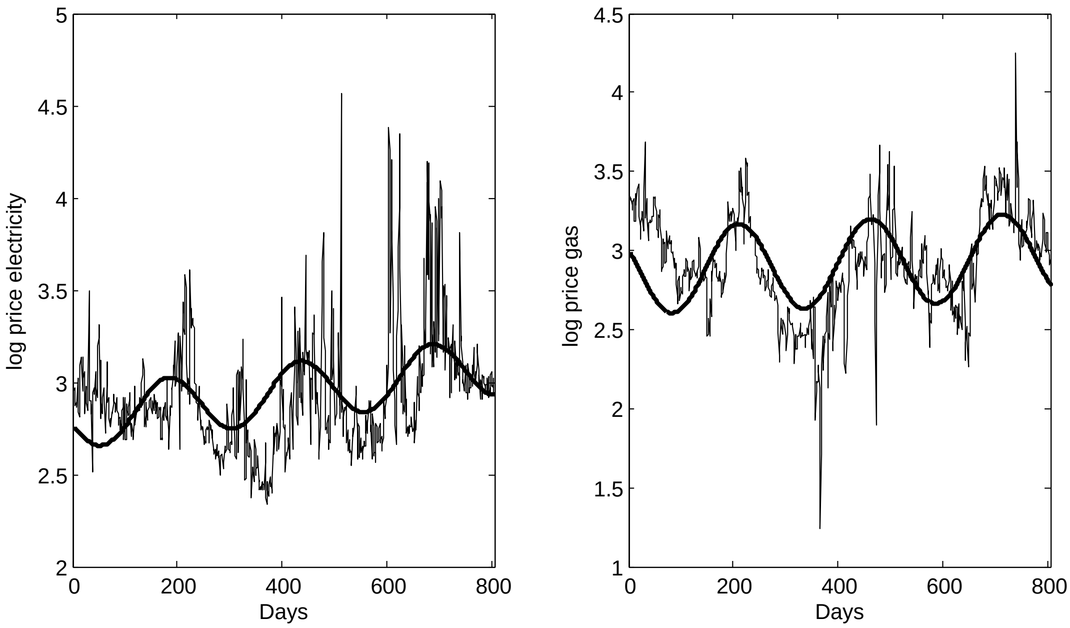

FIGURE 1. Data Series with Fitted Seasonal Functions

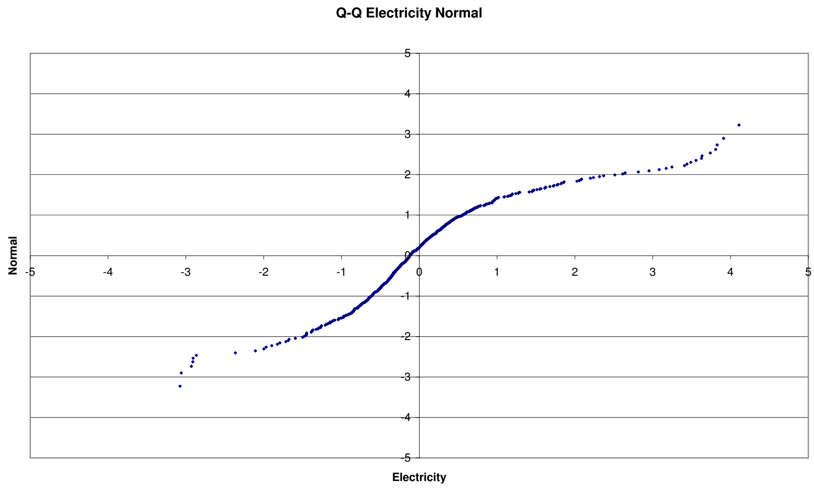

FIGURE 2. Q-Q Electricity Normal

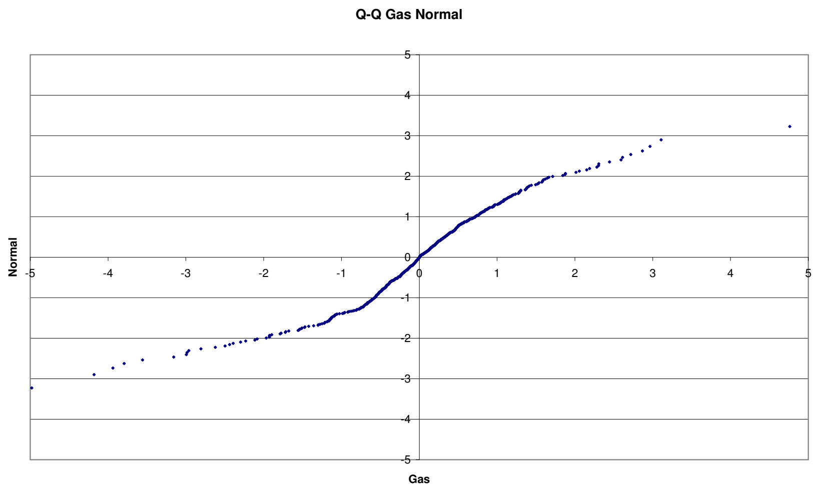

FIGURE 3. Q-Q Gas Normal

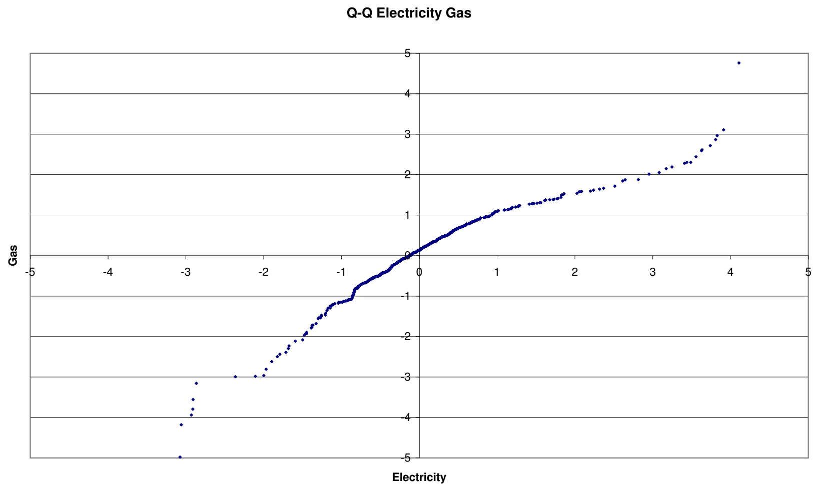

FIGURE 4. Q-Q Electricity Gas

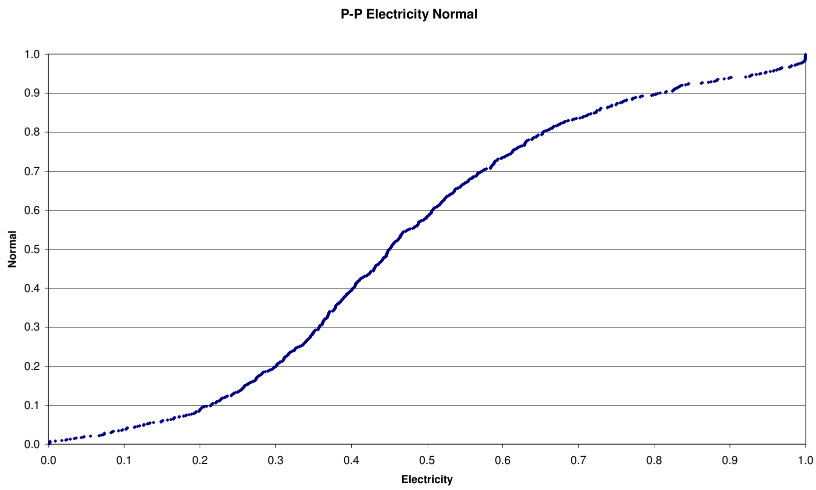

FIGURE 5. P-P Electricity Normal

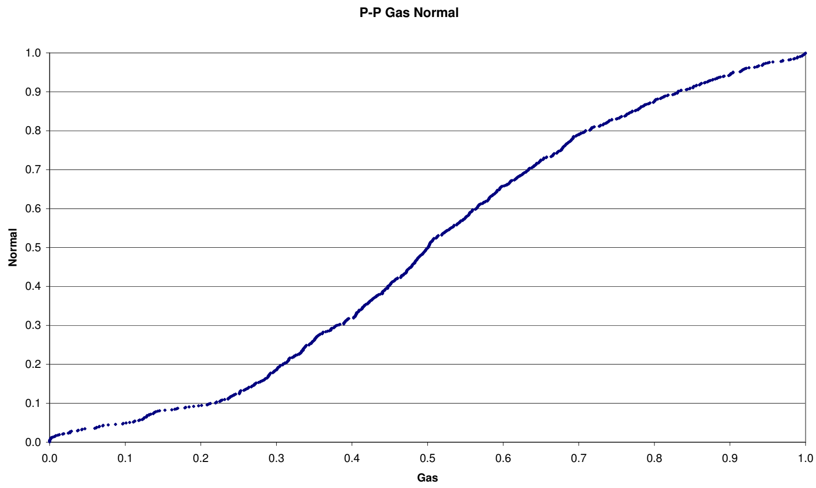

FIGURE 6. P-P Gas Normal

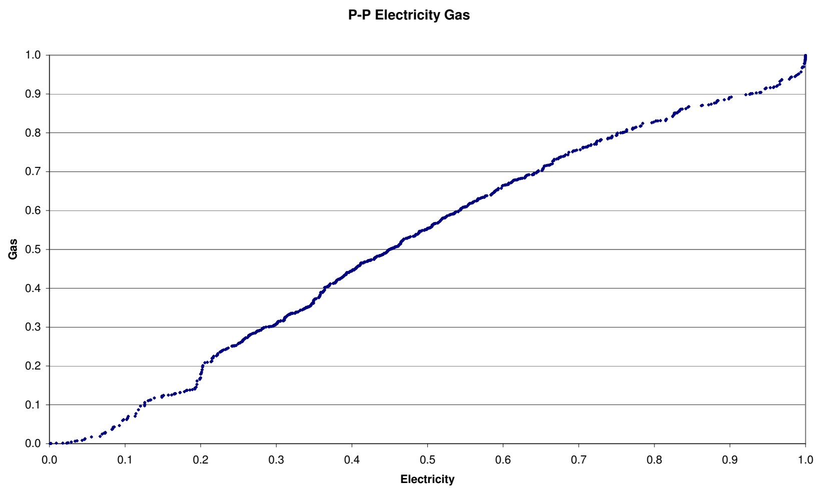

FIGURE 7. P-P Electricity Gas

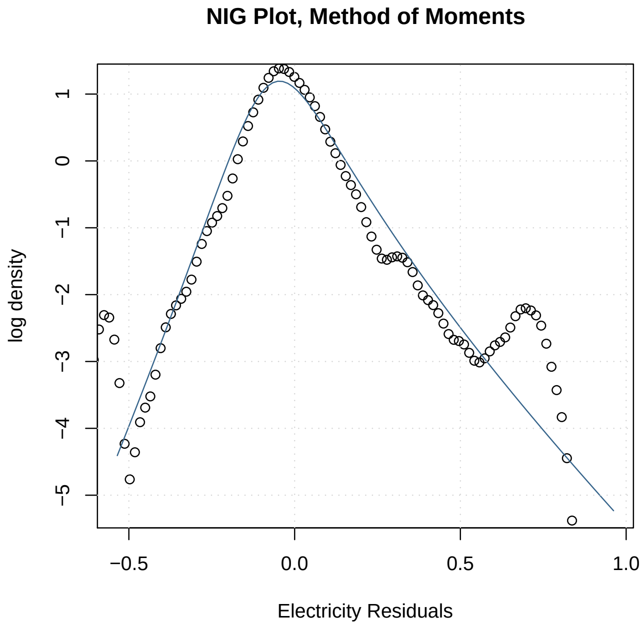

FIGURE 8. Electricity residuals, NIG fit, $(\alpha, \beta, \mu, \delta)$ =(6.9342, +2.8003, -0.0694, 0.1514)

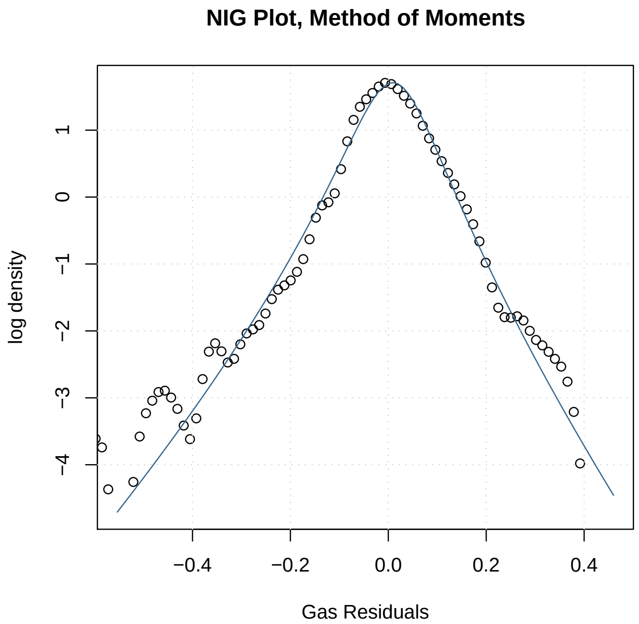

FIGURE 9. Gas residuals, NIG fit, $(\alpha, \beta, \mu, \delta)$ =(7.7740, -0.9982, +0.0122, 0.0831)

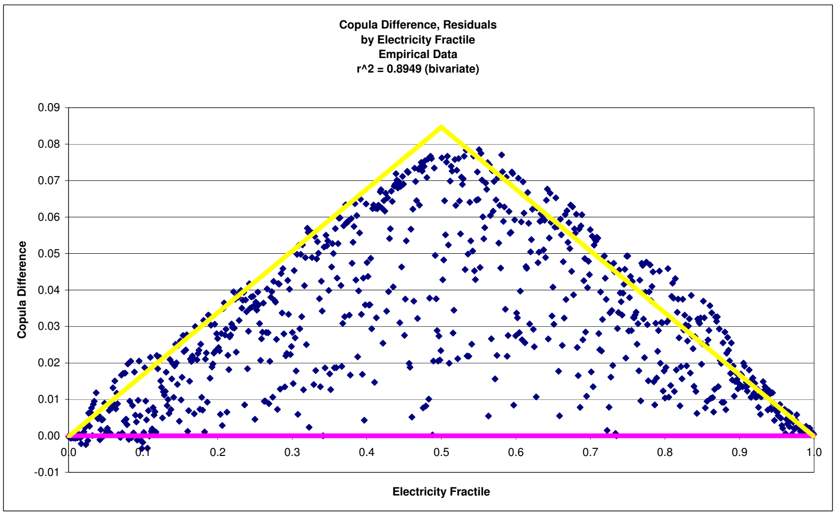

FIGURE 10. Copula Difference Electricity

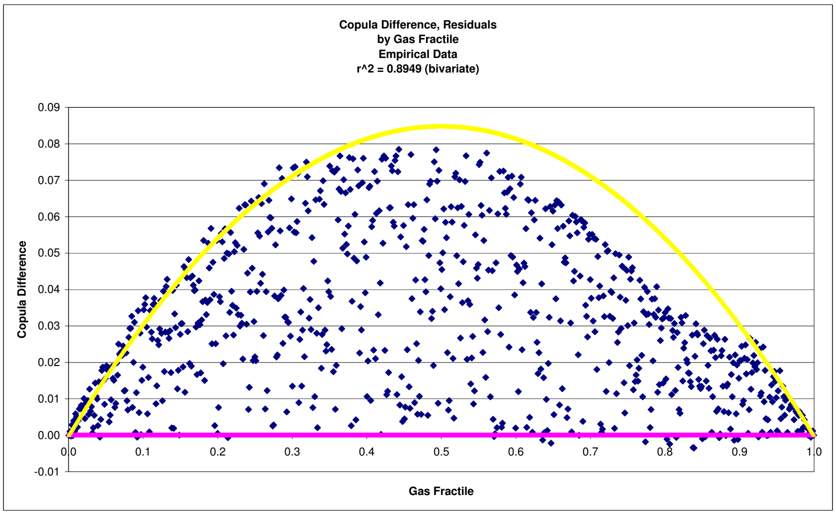

FIGURE 11. Copula Difference Gas

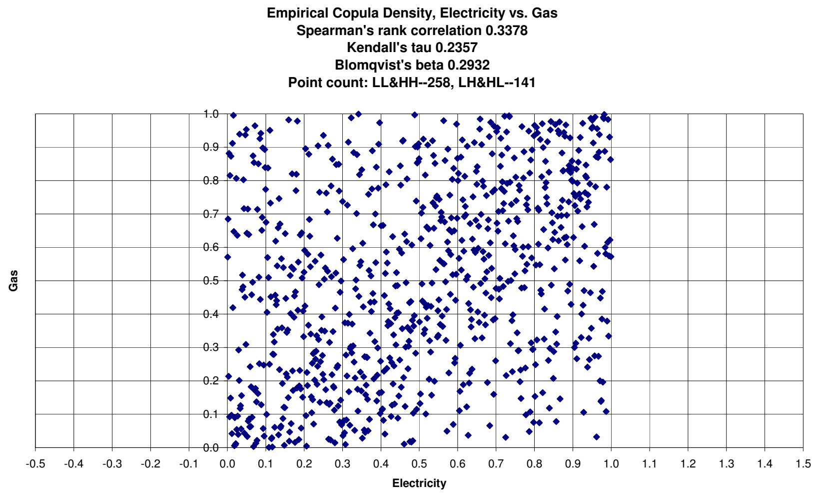

FIGURE 12. Empirical Copula Density

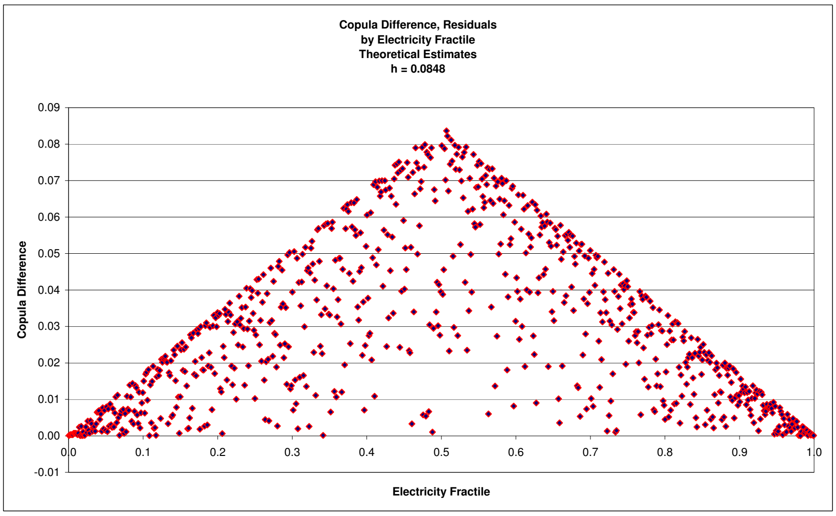

FIGURE 13. Copula Difference Electricity, Estimate

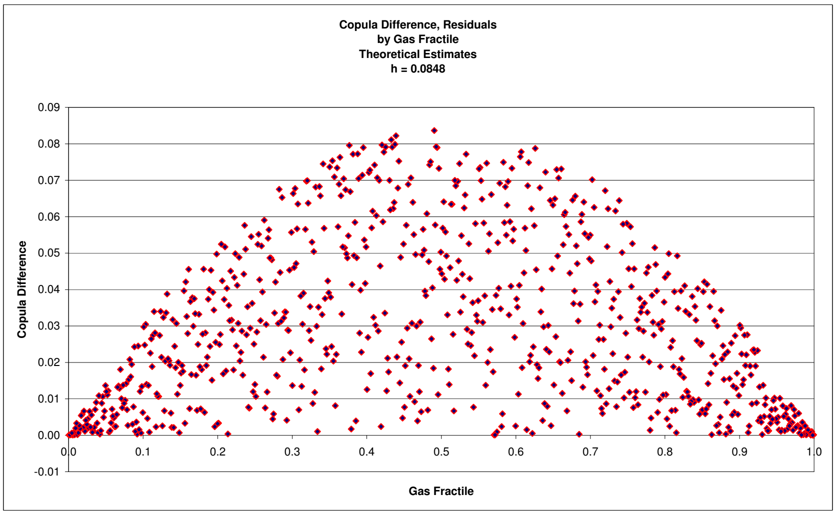

FIGURE 14. Copula Difference Gas, Estimate

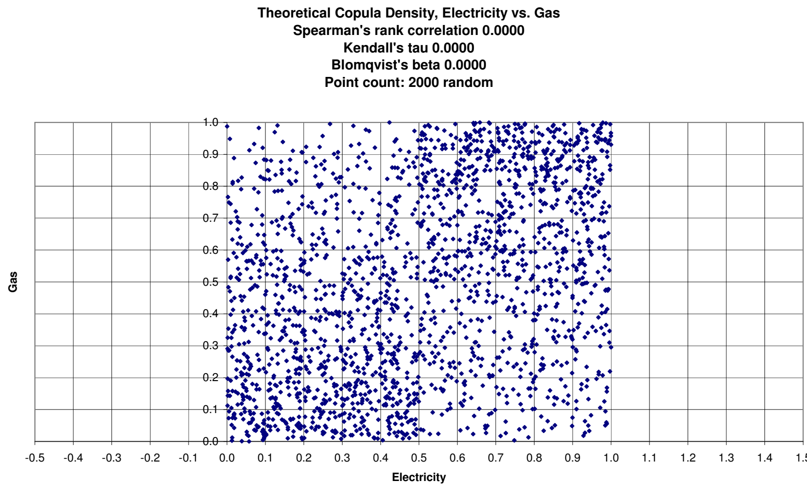

FIGURE 15. Theoretical Copula Density

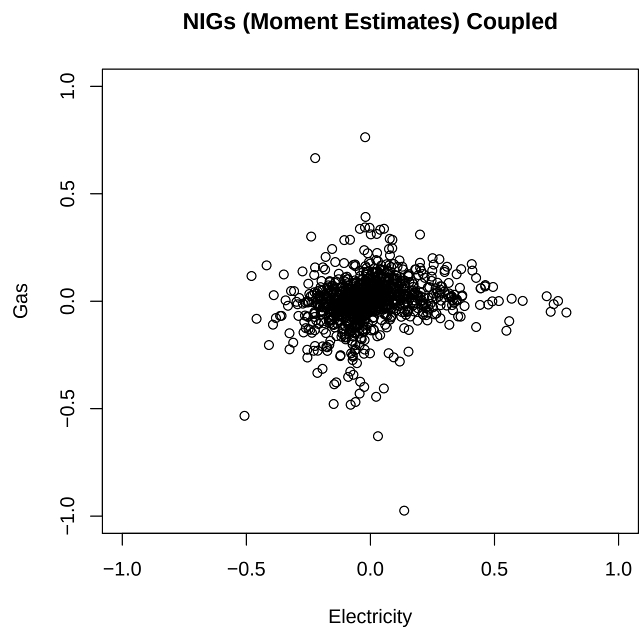

FIGURE 16. Pair estimates, 1000 points, NIG fit

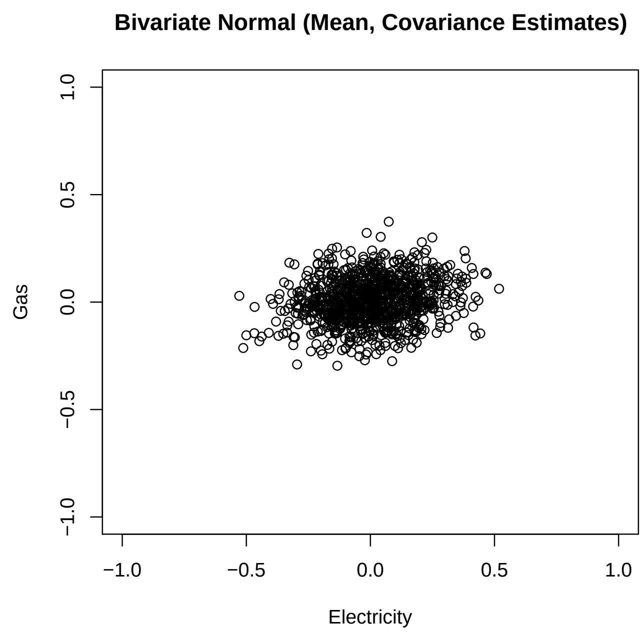

FIGURE 17. Pair estimates, 1000 points, Binormal fit

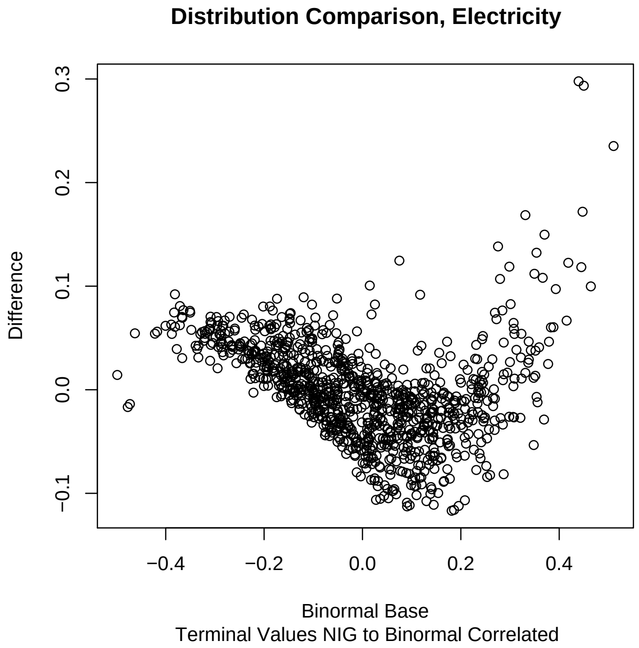

FIGURE 18. NIG minus binormal terminal distributions, compared to binormal, electricity axis

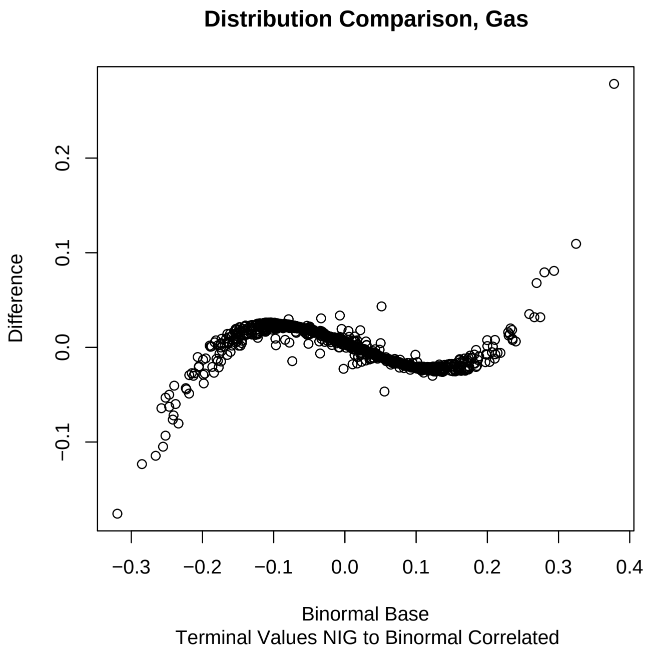

FIGURE 19. NIG minus binormal terminal distributions, compared to binormal, gas axis

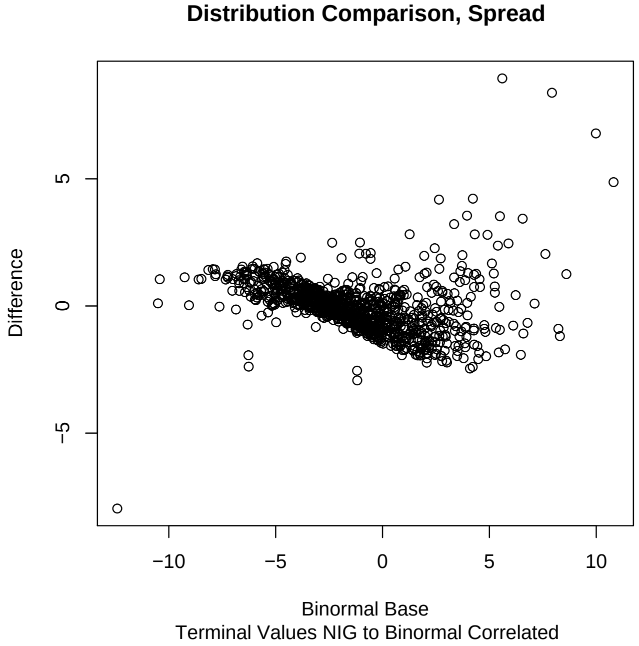

FIGURE 20. NIG minus binormal terminal spreads, compared to binormal spread, Q1

REFERENCES

Barndorff-Nielsen, O. E. (1998). Processes of normal inverse Gaussian type. *Finance Stochastics 2*, 41–68.

Benth, F. E., Á. Cartea, and R. Kiesel (2008). Pricing forward contracts in power markets by the uncertainty equivalence principle: Explaining the sign of the market risk premium. *J. Banking Finance 32*(10), 2006–2021.

Benth, F. E., L. Ekeland, R. Hauge, and B. F. Nielsen (2003). A note on arbitrage-free pricing of forward contracts in energy markets. *Appl. Math. Finance 10*(4), 325–336.

Benth, F. E., R. Kiesel, and A. Nazarova (2009, Oct.). A critical empirical study of two electricity spot price models. Preprint, Pure Mathematics, Department of Mathematics, University of Oslo.

Benth, F. E. and T. Meyer-Brandis (2009). The information premium for non-tradable commodities. *J. Energy Markets 2*(3). Special issue for the Conference on Financial Modelling in Energy Markets, Karlsruhe, Germany, October 9–10, 2008).

Benth, F. E. and J. Šaltytė-Benth (2004). The normal inverse Gaussian distribution and spot price modelling in energy markets. *Int. J. Theoretical Appl. Finance 7*(2), 177–192.

Benth, F. E. and J. Šaltytė-Benth (2006). Analytical approximation for the price dynamics of spark spread options. *Stud. Nonlinear Dynam. Econometrics 10*(3). Article 8 (electronic).

Börger, R., Á. Cartea, R. Kiesel, and G. Schindlmayer (2007). A multivariate commodity analysis and applications to risk management. Birkbeck College, London.

Carmona, R. and V. Durrelman (2003). Pricing and hedging spread options. *SIAM Rev. 45*(4), 627–685.

Cartea, Á. and M. G. Figueroa (2005, Dec.). Pricing in electricity markets: a mean reverting jump diffusion model with seasonality. *Appl. Math. Finance 12*(4), 313–335.

Cartea, Á. and T. Williams. UK gas markets: the market price of risk and applications to multiple interruptible supply contracts. To appear in Energy Economics.

Cherubini, U., E. Luciano, and W. Vecchiato (2004). *Copula Methods in Finance*. Chichester: Wiley.

Dempster, M. A. H. and S. S. G. Hong (2000). Spread option valuation and the fast Fourier transform. The Judge Institute of Management Studies, University of Cambridge, WP 26/2000.

Fusai, G., M. Marena, and A. Roncoroni (2008, Oct.). Analytical pricing of discretely monitored Asian-style options: Theory and application to commodity markets. *J. Banking Finance 30*(10), 2033–2045.

Geman, H. (2005). *Commodities and Commodity Derivatives: Modeling and Pricing for Agriculturals, Metals, and Energy*. West Sussex: Wiley Finance.

Geman, H. and A. Roncoroni (2006). Understanding the fine structure of electricity prices. *J. Bus. 79*(3), 1225–1261.

Kallsen, J. and P. Tankov (2004). Characterization of dependence of multidimensional Lévy processes using Lévy copulas. Preprint, Technische Universität München.

Lucía, J. J. and E. S. Schwartz (2002). Electricity prices and power derivatives: evidence from the Nordic Power Exchange. *Rev. Derivatives Res. 5*(1), 5–50.

Nelsen, R. B. (1998). *An Introduction to Copulas*. New York: Springer-Verlag.

Raible, S. (2000, Apr.). *Lévy Processes in Finance: Theory, Numerics, and Empirical Facts*. Ph. D. thesis, Albert-Ludwigs-Universität, Frieburg.

Rydberg, T. H. (1997). The normal inverse Gaussian Lévy process: simulation and approximation. *Comm. Statist. Stochastic Models 13*(4), 887–910.

(Fred Espen Benth)
CENTRE OF MATHEMATICS FOR APPLICATIONS
DEPARTMENT OF MATHEMATICS
UNIVERSITY OF OSLO
P.O. BOX 1053, BLINDERN
N–0316 OSLO
NORWAY
AND
UNIVERSITY OF AGDER
SCHOOL OF MANAGEMENT
SERVICEBOKS 422
N-4604 KRISTIANSAND
NORWAY
E-mail address: fredb@math.uio.no
URL: http://www.math.uio.no/~fredb/

(Paul C. Kettler)
CENTRE OF MATHEMATICS FOR APPLICATIONS
DEPARTMENT OF MATHEMATICS
UNIVERSITY OF OSLO
P.O. BOX 1053, BLINDERN
N–0316 OSLO
NORWAY
E-mail address: paulck@math.uio.no
URL: http://www.math.uio.no/~paulck/