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 Ornstein-Uhlenbeck pure jump process. We identify the marginal processes as driven by normal inverse Gaussian Lévy processes, estimating them by the method of moments. We estimate the copula by modeling the difference of the empirical copula to the independent copula. Following we simulate the joint process and conclude with valuation of options.

1. Introduction

We propose a new copula to model the dependence relationship in the joint evolution of electricity and gas prices. For the marginal processes we choose discrete-time Ornstein-Uhlenbeck Lévy processes having normal inverse Gaussian (NIG) jump components. With the model calibrated we simulate the joint price series and produce option values on the spark spread.

The principal findings of the study are that the proposed theoretical copula and marginal assumptions fit the data well, that the NIG assumption, compared to the more traditional binormal concept, produces fat tails consistent with observation, and that the model's predicted prices are nearly invariant to the choice of distribution inside those tails.

Specifically, the spark spread as a function of time is this.

$$(1.1) \qquad S(t) := E(t) - c\,G(t),$$

where $E(t)$ and $G(t)$, respectively, are electricity and gas prices quoted in customary units. The constant c is a *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. See Appendix A for a discussion in detail on this point.

For inspiration and motivation we observe and comment upon a series of graphs depicting the difference between our proposed theoretical copula and the independent copula. As well, we compare the empirical copula to the independent copula. By various tests the proposed theoretical copula provides a good fit to the data, and therefore serves as a basis for the simulation study to follow. We estimate by the method of moments parameters of the NIG distribution for each marginal series, and with these distributions and the copula connecting them the simulations follow.

For a presentation on the spark and other spreads, with stochastic analysis including mean reversion, jump terms, and attention to seasonality see (Carmona and Durrelman 2003). For further analysis of the economics and dynamics of spark spreads see (Benth and Šaltytė-Benth 2006). For a specific treatment of the Nordic power market employing one and two factor models with deterministic and stochastic components (but without jump terms) see (Lucía and Schwartz 2002). Providing a detailed microanalytic stochastic approach with a novel likelihood estimator applied to jump processes is this paper (Geman and Roncoroni 2006). Delving into a jump diffusion model of the English and Welsh spot and forward electricity markets is this paper (Cartea and Figueroa 2005). For a comprehensive exposition on an exponential mean reversion model with NIG marginals see (Benth and Šaltytė-Benth 2004).

We set forth our model for electricity and gas prices in several stages, reported in the subsequent sections. First is a discussion of the model with emphasis on the copula. Next comes a section on theory incorporating a description of the normal inverse Gaussian distribution, with formulation for estimating parameters by the method of moments. Within this section are items on a property of the copula difference function — important to our estimation procedure – and a definition for tail dependence, on which we make a later comment. We continue with a discussion of the data and inferences. This section also delves into the NIG estimations of the marginals, and the copula estimation using the empirical copula.

Next we offer a section on the simulations beginning with our method to make a random draw from the fitted copula, proceeding to the generation of the sample sequences, and concluding with inferences from the terminal distributions of the paths. Graphical material here assists in interpretation. Before conclusions comes the section on option pricing, with detail on those methods. Included here is a table illustrating theoretical prices for 20-day options scheduled to four seasons of the year.

2. A COPULA MODEL FOR ELECTRICITY AND GAS

Let $(\Omega, \mathcal{F}, P)$ be a standard filtered probability space, wherein $\{\mathcal{F}_t\}_{0 \leq t \leq T}$ is the filtration. The fundamental joint model we consider is this, incorporating pure jump Lévy processes. Thus we have

$$
\begin{aligned}
\mathrm{d}\widetilde{E}(t) &= -\alpha_E \left(\widetilde{E}(t) - \mu_E \right) \mathrm{d}t + \mathrm{d}L_E(t) \\
\mathrm{d}\widetilde{G}(t) &= -\alpha_G \left(\widetilde{G}(t) - \mu_G \right) \mathrm{d}t + \mathrm{d}L_G(t)
\end{aligned}
$$
(2.1)

In this model $\widetilde{E}(t)$ and $\widetilde{G}(t)$ are the *detrended and deseasonalized logarithmic electricity and gas processes*, and $L_E(t)$ and $L_G(t)$ are the corresponding pure jump Lévy processes. The parameters α_E and α_G describe rates of mean reversion, with μ_C and μ_G the implied means of those reversions.

Here is the discrete-time version of these formulas, which we employ in the rest of the paper; it replaces the differentials with differences. In the further analysis the $\{t_i\}$ constitute a natural number index set, so $t_{i+1} - t_i = 1$. However, here we state the formulas in generality (with slight abuse of notation in using the same symbols as for the continuous case.)

$$
\begin{aligned}
\widetilde{E}(t_{i+1}) - \widetilde{E}(t_i) &= -\alpha_E \left(\widetilde{E}(t_i) - \mu_E \right) (t_{i+1} - t_i) + \varepsilon_E(t_i) \\
\widetilde{G}(t_{i+1}) - \widetilde{G}(t_i) &= -\alpha_G \left(\widetilde{G}(t_i) - \mu_G \right) (t_{i+1} - t_i) + \varepsilon_G(t_i),
\end{aligned}
$$
(2.2)

where the independent and identically distributed Lévy increments $\varepsilon_E(t_i) := L_E(t_{i+1}) - L_E(t_i)$ and $\varepsilon_G(t_i) := L_G(t_{i+1}) - L_G(t_i)$. We join then the Lévy processes with a copula, estimated from data.

The emphasis of the model is on using a copula, a function joining two marginal distributions and containing all the dependency information between the random variables having those distributions. For a good foundation of copula theory see (Nelsen 1998). We propose a theoretical copula, which fits the data well and provides insight into the nature of electricity and gas price discovery. As well, we perform studies on the separate electricity and gas data resulting in additional series of residuals, necessary in modeling the marginal distributions.

Remark. This study develops a discrete-time model in order to take advantage of the discrete-time empirical copula provided by the source data, as reduced through the detrending and deseasonalizing procedures. We then are able to compute a theoretical copula fitted to this empirical copula. The plan of research herein, therefore, is to couple the marginal distributions, estimated directly as NIG variates by the method of moments, rather than to couple marginal Lévy measures with a Lévy copula. Marginal Lévy measures are significantly more challenging to estimate. For a good theoretical foundation on using the Lévy copulas for this and other purposes see (Kallsen and Tankov 2004).

The model we propose, having implications for the price processes generating the observed empirical distribution, is described by a copula difference function $\Delta(v, z)$ to the independent copula $C_\perp(vz) := vz$. This function is triangular in the electricity variable and quadratic in the gas variable with one parameter h, the height of the copula.

For the discrete formulation of the model, let $\{\varepsilon_E(t_i)\}$ and $\{\varepsilon_G(t_i)\}$, as above, be the data corresponding to $\widetilde{E}(t)$ and $\widetilde{G}(t)$, respectively. These discrete points are the residuals of the autoregression study to estimate the Ornstein-Uhlenbeck reversion coefficients α_E and α_G of Equations (2.1). Then let $F^E(\cdot)$ and $F^G(\cdot)$ be their empirical distributions and let $D(\cdot, \cdot)$ be the copula joining them.

Specifically, set

$$(2.3) \qquad \Delta(v, z) := (1 - |2v - 1|)\left(1 - (2z - 1)^2\right),$$

and then look to the model residuals

$$(2.4) \qquad \left\{\rho_j := D\left(F^E\left(r_j^E\right), F^G\left(r_j^G\right)\right) - h\Delta\left(F^E\left(r_j^E\right), F^G\left(r_j^G\right)\right)\right\}_{j=1}^n,$$

adjusting h so to minimize the sum of their squares. We then refer to

$$(2.5) \qquad C_h(v, z) := C_\perp(v, z) + h\Delta(v, z) = vz + h(1 - |2v - 1|)\left(1 - (2z - 1)^2\right)$$

as the copula we propose to link electricity and gas.[1]

3. Theoretical considerations

For background on the behavior of the NIG distribution under convolution, and on general principles for fitting it, see these (Rydberg 1997; Barndorff-Nielsen 1998). For a comprehensive treatment of NIG distributions as they relate to Lévy processes see this (Raible 2000).

---

[1] These model residuals are not to be confused with the basic data $\{r_j^E\}$ and $\{r_j^G\}$

The forward formulas for the statistics (m, v, s, k), given the parameters $(\alpha, \beta, \mu, \delta)$ are given here (Rydberg 1997) and now repeated. For convenience let $\gamma = \sqrt{\alpha^2 - \beta^2}$.

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

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

These formulas are invertible to the following backward formulas. For convenience let $\zeta = 3k - 4s^2$ and $\eta = k - \frac{5}{3}s^2$.

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

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

The NIG density $g(x; \alpha, \beta, \mu, \delta)$ is known in closed form.

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

where

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

and $K_1(x)$ is the modified Bessel function of the second kind of index one.

Consider next the estimated theoretical copula $C_h(v, z)$ along with the marginal distributions $\widehat{F}^E(a)$ and $\widehat{F}^G(b)$. One may recover the estimated joint distribution $\widehat{F}(a, b)$, as follows.

$$\widehat{F}(a, b) = C_h\left(\widehat{F}^E(a), \widehat{F}^G(b) \right)$$

Recall that $\Delta(v, z)$ is not arbitrarily, for it must conform to a conditions of Lemma 3.1 below. It does for $h \in [-0.1250, +0.1250]$.

Lemma 3.1. $\partial\Delta(v, z) \equiv 0$ and $\displaystyle\int_0^1 \int_0^1 \frac{\partial^2}{\partial v \partial z} \Delta(v, z)\, \mathrm{d}v\, \mathrm{d}z = 0.$

Proof. The first part follows readily as all copulas have common values on their boundaries. For the second part consider that the mixed second partial derivative of a differentiable distribution function is its density. Therefore, from Equation (2.5) we have

$$\int_0^1 \int_0^1 \frac{\partial^2}{\partial v \partial z} C_h(v, z)\, \mathrm{d}v\, \mathrm{d}z = 1 + h \int_0^1 \int_0^1 \frac{\partial^2}{\partial v \partial z} \Delta(v, z)\, \mathrm{d}v\, \mathrm{d}z = 1,$$

whence the conclusion follows.   □

We turn attention now to tail dependence with a definition. For a treatment of this concept see (Cherubini, Luciano, and Vecchiato 2004, Subsection 1.8.5, pp. 42–43).

Definition 3.2. A bivariate distribution is *lower tail dependent* with coefficient λ_L, $0 \leq \lambda_L \leq 1$ if

$$\lim_{a \to -\infty} \Pr\left\{ Y \leq a \mid X \leq a \right\} = \lim_{a \to -\infty} \frac{G(a,a)}{F(a)} = \lim_{\alpha \to 0} \frac{H(\alpha, \alpha)}{\alpha} = \lambda_L$$

A bivariate distribution is *upper tail dependent* with coefficient λ_U, $0 \leq \lambda_U \leq 1$ if the distribution of $(-X, -Y)$ is lower tail dependent with coefficient λ_U. A distribution is either *lower tail independent* or *upper tail independent*, respectively, as $\lambda_L = 0$ or $\lambda_U = 0$. A distribution is either *lower tail completely dependent* or *upper tail completely dependent*, respectively, as $\lambda_L = 1$ or $\lambda_U = 1$.

4. 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. Of these observations we truncate seven as extreme. This truncation changes the statistics of the residuals slightly, but insignificantly, and has the effect of eliminating those few observations as jumps. The intent, however, is not to eliminate jumps — as the decision has been made to combine the residuals into a jump term — but rather to leave the sample more manageable in other ways, as in performing a least squares fit of the difference copula.

4.1. **Data preparation.** Detrending of the logarithmic data follows this discrete model, with these estimates. Herein the original logarithmic data are $\widetilde{E}_1(t_i) = \log E(t_i)$ and $\widetilde{G}_1(t_i) = \log G(t_i)$, respectively for electricity and gas, and the detrended data are the residuals $\{\varepsilon_{E,1}(t_i)\}$ and $\{\varepsilon_{G,1}(t_i)\}$.

$$(4.1) \qquad \begin{aligned} \widetilde{E}_2(t_i) &= a_E^{(1)} + a_E^{(2)} \widetilde{E}_1(t_i) + \varepsilon_{E,1}(t_i) \\ \widetilde{G}_2(t_i) &= a_G^{(1)} + a_G^{(2)} \widetilde{G}_1(t_i) + \varepsilon_{G,1}(t_i) \end{aligned}$$

The estimated coefficients are $a_E^{(1)} = 2.7851$, $a_E^{(2)} = 0.0004$, $a_G^{(1)} = 2.8737$, $a_G^{(2)} = 0.0001$. The coefficients $a_E^{(2)}$ and $a_G^{(2)}$ are determined insignificantly different from zero.

Deseasonalizing the detrended logarithmic data follows this subsequent model, with these estimates. The final data are indicated without subscript, as in Equations (2.1). $D = 252$ is the assumed number of trading days in a year. Analogously, the reduced data are the residuals $\{\varepsilon_E(t_i)\}$ and $\{\varepsilon_G(t_i)\}$.

$$(4.2) \qquad \begin{aligned} \widetilde{E}_3(t_i) &= \varepsilon_{E,1}(t_i) + b_E^{(1)} + b_E^{(2)} \cos\left[2\pi \left(t + b_E^{(3)}\right)/D\right] + \varepsilon_E(t_i) \\ \widetilde{G}_3(t_i) &= \varepsilon_{G,1}(t_i) + b_G^{(1)} + b_G^{(2)} \cos\left[2\pi \left(t + b_G^{(3)}\right)/D\right] + \varepsilon_G(t_i) \end{aligned}$$

The estimated coefficients are $b_E^{(1)} = 0.0026$, $b_E^{(2)} = 0.1603$, $b_E^{(3)} = 49.1820$, and $b_G^{(1)} = -0.0062$, $b_G^{(2)} = 0.2748$, $b_G^{(3)} = 21.4688$. The coefficients $b_E^{(1)}$ and $b_G^{(1)}$ are determined insignificantly different from zero.

We perform statistical tests on these series $\{\varepsilon_E(t_i)\}$ and $\{\varepsilon_G(t_i)\}$ to satisfy ourselves that they are well behaved and essentially stationary. Among such tests is Fisher's variance ratio test for absence of heteroskedasticity. The F statistic performed on the first and last thirds of the data (points 1–266 and 533–798) in each series results in the acceptance of null hypotheses

of homoskedasticy. Specifically, F ratios of 0.3468 and 0.7020 are observed in the electricity and gas series, with corresponding confidence intervals of $[0.2725, 0.4415]$ and $[0.5515, 0.8936]$.

Now make the identifications

$$\widetilde{E}(t_i) := \varepsilon_E(t_i)$$
$$\widetilde{G}(t_i) := \varepsilon_G(t_i)$$

These are the variables of the Ornstein-Uhlenbeck Equations (2.2).

To review, this is the sequence transformations and generated variables to arrive at the basic residual series data for the study.

$$E(t_i) \xrightarrow[\text{transform}]{\log} \widetilde{E}_1(t_i) \xrightarrow[\text{regression}]{\text{trend}} \varepsilon_{E,1}(t_i) \xrightarrow[\text{regression}]{\text{season}} \varepsilon_E(t_i) =: \widetilde{E}(t_i)$$

$$G(t_i) \xrightarrow[\text{transform}]{\log} \widetilde{G}_1(t_i) \xrightarrow[\text{regression}]{\text{trend}} \varepsilon_{G,1}(t_i) \xrightarrow[\text{regression}]{\text{season}} \varepsilon_G(t_i) =: \widetilde{G}(t_i)$$

Consistent with the process descriptions of these equations we estimate by one period autoregression the coefficients $\alpha_E = 0.7278$, $\mu_E = -0.0002$, $\alpha_G = 0.8849$, $\mu_G = -0.0004$, concluding that μ_E and μ_G are insignificantly different from zero.[2] The prior detrending and deseasonalizing of the original logarithmic data follow linear and cyclic models presented here. These results we use in reverse order for the final recovery of the electricity and gas price sequences, and the consequent computation of the spread series on which we estimate spark spread option prices.

4.2. **Marginal and copular relationships.** We model the spark spread residuals as NIG variates on the separate electricity and gas series. This methodology makes possible the specific estimation of the joint distribution by copula methods, inspiring simulation of the joint process and subsequent pricing of derivative securities on it. The series are separately estimated, not jointly estimated.

The selection and fitting of a theoretical one-parameter copula to the joint process is central to the model. This copula joins $L_E(1)$ to $L_G(1)$, the value $t = 1$ signifying the passage of one trading day, or one observation in the data.

We also looked at the Q-Q and P-P functions of the variables and can report our inferences. Begin with the Q-Q functions. See Figures 8, 9, and 10. A qualitative interpretation of the electricity function 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, while also exhibiting leptokurtosis, shows distinctly thin tails.

This comparative distinction appears directly in the cross reference electricity-gas Q-Q function, 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.

In each of 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 11, 12, and 13.

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

---

[2]These estimates were provided by Jūratė Šaltytė-Benth.

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

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 we employ has zero tail dependence by Definition 3.2. 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.

4.3. Normal inverse Gaussian estimation of the marginal distributions. We estimate the marginal distributions for electricity and gas residuals from the empirical distributions F^E and F^G by the method of moments using the NIG distribution, as articulated in Section 3.

Implementing these estimates provides fits of theoretical densities $\widehat{f}^E(x)$ and $\widehat{f}^G(y)$ on histograms exhibited in Figures 14 and 15, with parameters $(\alpha, \beta, \mu, \delta) = (6.9342, +2.8003, -0.0694, 0.1514)$ for electricity, and $(\alpha, \beta, \mu, \delta) = (7.7740, -0.9982, +0.0122, 0.0831)$ for gas. Associated are distributions $\widehat{F}^E(a)$ and $\widehat{F}^G(b)$.

4.4. Estimation of the copula. The estimated value by least squares on the empirical copula of h is 0.0848, well within the Fréchet-Hoeffding interval [-0.1250, +0.1250], implying a coefficient of determination $r^2 = 0.8949$. Figure 1 exhibits $\Delta(v, z)$, a function on the unit square. Figure 2 shows the level curves of $\Delta(v, z)$. In these charts $h = 1$.

Look next to Figures 3 and 4. These figures show an estimated theoretical copula difference function as in Equation (2.3), 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.[3]
Look further to Figures 5 and 6. These figures show the empirical copula difference function,
also from the electricity axis into the domain, and toward the gas axis out of the domain.
What you see is an asymmetrical function in keeping with the model choice. The views are
qualitatively different, owing to the different character of the separate markets for gas and
electricity. Gas is storable, whereas electricity is not. The development of our model captures
this distinction. Figure 7 shows the density of this empirical copula, a "top view," of which
Figures 5 and 6 are the "side views."

5. Simulation of the joint electricity and gas processes

Simulation begins by taking random draws from the theoretical copula $C_h(v, z)$ of Equation (2.5). To gain insight into how a plot of such draws may look, we compute the copular density as the second mixed partial derivative of the copula, thus.

$$c_\theta(v, z) := \frac{\partial^2}{\partial v \partial z} C_h(v, z) = 1 + \text{sgn}(2v - 1) \cdot 8h(2z - 1)$$

This density $c_\theta(v, z)$ is constant on v and linear on z in the subsets of the unit square $v \leq 1/2$ and $v > 1/2$. Also, $c_\theta(v, z)$ is symmetric about $[1/2, 1/2]$ in that $c_\theta(v, z) = c_\theta(1 - v, 1 - z)$. Accordingly, $c_\theta(v, z)$ is discontinuous for $v = 1/2$, except for $v = z = 1/2$. Further, $c_\theta(v, 1/2) = 1$.

Drawing samples from this copula is a straightforward process. One may select either variable randomly by the uniform distribution, and then select the alternative variable by inverting on the conditional distribution for it.

The two conditional distributions $C_h^{(\bar{v})}(\bar{z})$ and $C_h^{(\bar{z})}(\bar{v})$ are as follows.

(5.1)

$$C_h^{(\bar{v})}(\bar{z}) = \Pr\left\{z \leq \bar{z} \mid v = \bar{v}\right\} = \frac{\partial}{\partial v} C_h(v, z)\big|_{(\bar{v}, \bar{z})} = \bar{z} - \text{sgn}(2\bar{v} - 1) \cdot 2h[1 - (2\bar{z} - 1)^2]$$

(5.2)

$$C_h^{(\bar{z})}(\bar{v}) = \Pr\left\{v \leq \bar{v} \mid z = \bar{z}\right\} = \frac{\partial}{\partial z} C_h(v, z)\big|_{(\bar{v}, \bar{z})} = \bar{v} - 4h(1 - |2\bar{v} - 1|)(2\bar{z} - 1)$$

The way is clear now to select a pair $(\bar{v}, \bar{z})$ by the distribution $C_h(v, z)$. Start with two independent draws, U_1 and U_2, from the uniform distribution on the unit interval. One may begin with either argument of $C_h(v, z)$. We choose the second (to avoid a quadratic equation.) Therefore, let

(5.3) $$\bar{z} = U_2$$

To continue, invert Equation (5.2) above. To that end let

$$C_h^{(\bar{z})}(\bar{v}) = U_1$$

---

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

and solve for $\bar{v}$. For convenience let $\widehat{h} = 8h(2\bar{z}-1)$ and then let $\widehat{v} = U_1/(1-\widehat{h})$. The inversion splits neatly into two cases, as follows.

$$(5.4) \qquad \bar{v} = \begin{cases} \widehat{v} & \text{if } \widehat{v} \leq \frac{1}{2} \\ \dfrac{U_1 + \widehat{h}}{1 + \widehat{h}} & \text{if } \widehat{v} > \frac{1}{2} \end{cases}$$

Taking $\bar{v}$ from Equation (5.4) and $\bar{z}$ from Equation (5.3), the pair $(\bar{v}, \bar{z})$ is a draw from $C_h(v, z)$. Figure 16 shows 2000 of these draws.

We proceed to construct a series of random paths based on draws from the theoretical copula to the estimated NIG marginal distributions. For control we generate paths with the binormal distribution assumption for the residuals, first by a method of moments using the respective means and covariance of the samples, and then assuming zero covariance.

The parallel studies employ coordinated draws from two independent uniform distributions, using those variates to choose respectively a copula pair and bivariate normal pairs. From the copula pair we invert the NIG marginal distributions to establish points in the time series, as we do directly with the bivariate normal pairs.

In order to avoid Monte Carlo error we use the same pair of independent uniform variates in the NIG study as in the control binormal study. This requirement necessitates separate scripting of the binormal draws, as the available software takes four uniform draws, not two, to create each binormal point when presented with a 2x2 covariance matrix.

The first step of the simulation is to make $40,000$ draws from the uniform distribution to construct 1000 sequences of pairs of points for the NIG and binormal studies. These data are necessary and sufficient to produce independent pairs for $p := 20$-day periods of trading days, the experimental design interval. The next step is to generate the sequences using the iterative formula implied by the autoregression study, and thereafter to correct the paths for seasonal and trend influence, finally exponentiating to arrive at prices.

Following the generation of these paths we correct the sets of endpoints for seasonal influences according to Equations (4.2), choosing four separate periods equally spaced around the yearly cycle. Specifically, we correct the distributions for the ending days of periods beginning on days 0, 63, 126, and 189. As we choose $p = 20$ for the cycle length, those ending days are numbers 20, 83, 146, and 209. Then we correct the adjusted endpoint distributions further for trend according to Equations (4.1). The resulting values we exponentiate to produce prices. Finally, we compute spread distributions according to Equation (1.1).

With the paths so generated we examine several data sets along the way with tests, charts, and findings. First, we test the series of detrended, deseasonalized logarithms of prices for normality. By the Anderson-Darling (A-D) test both series for electricity and gas have the null hypotheses of normality rejected at the 1% level. Respectively, the A values are 25.356 and 12.519, with P values less than 10^{-15} in each case. We then fit NIG distributions, hypothesized as better descriptions of the data. Refer to Figures 14 and 15.

Figures 17 and 18 show 1000 pairs of draws, respectively, from the theoretical copula with inversions on the parametrized NIG marginal distributions, and analogous pairs drawn from the binormal distribution with means and covariance estimated from the samples.

Figures 19 through 22 show histograms of the terminal distributions in the NIG and binormal study for each of the electricity and gas series, as captioned. A-D normality tests cause rejection of the null hypothesis in the NIG examples, and acceptance of the hypothesis in the

binormal example. In the order of the charts, A-D values of the A statistic, with corresponding P values parenthesized are $9.246(< 10^{-15})$, $8.185(< 10^{-15})$, $0.3136(0.5465)$, $0.1945(0.8920)$.

We examined similar histograms, now for prices, under the NIG and binormal electricity and gas assumptions for four equally spaced 20-day periods around the year. The two sets of four means of these distributions, one set for the NIG and binormal fit, are these, in chronological order. NIG: -1.2860, 2.4434, 4.5018, and 1.3603; binormal: -1.2352, 2.4869, 4.5546, and 1.4243.

As well, we examined the price spread distributions, NIG and binormal, under the same assumptions as above. These spread prices appear as £/MWh, the unit for electricity. (See Appendix A.) By the A-D test each of these is rejected for normality at the 1% level. Respective $A(P)$ values are $15.985(< 10^{-15})$ and $1.1241(0.006068)$.

To conclude the study we show charts depicting various aspects of the terminal distributions. Figures 23, 24, and 25, respectively, compare the terminal distributions for the electricity and gas prices, and the spread between them. The spread distribution in this context refers to the terminal distribution of the first of the four simulated quarters.

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

6. Pricing of spark spread options

The arbitrage-free price of a spark spread call and options with maturity T and strike K are

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

where Q is a risk-neutral probability.

We choose to disregard any price of risk incurred by the agent, and let $Q = P$. Thus, after generating the simulated paths pricing of derivative securities is straightforward. We utilize the simulations of Section 5, relying first on the proposed theoretical copula with NIG margins, and then for control on the binormal copula with normal margins. We choose five strike prices $[-10, -5, 0, 5, 10]$ for pricing the European calls and puts by the NIG and binormal processes, for the four time periods — those 20-day periods equally spaced around the year. The results appear in Table 1. As for spread quotations all numbers in this table are in units of £/MWh. The assumed discount rate of interest r for illustration is 0%. All data conform to the parity relationship, where C and P are call and put prices in any context, $S(T)$ is the spread price at terminal time T, $\overline{S}$ is the mean of the spread distribution, and K is the strike price,

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

The price differences, call to call and put to put, comparing the NIG and binormal processes are small. Observe that out-of-the-money options under the binormal assumption are nearly worthless, whereas the same options under the NIG assumption have some value. These results are consistent with the relatively fat tails of the NIG distribution. It is difficult to see any influence of the copula in the option prices.

Strike	Period 1	Period 2	Period 3	Period 4	Type	Process
-10	8.7248	12.4448	14.5029	11.3708	Call	NIG
-10	0.0107	0.0013	0.0011	0.0105	Put	NIG
-10	8.7681	12.4869	14.5545	11.4260	Call	binormal
-10	0.0033	0.0000	0.0000	0.0017	Put	binormal
-5	3.7967	7.4499	9.5079	6.3978	Call	NIG
-5	0.0827	0.0065	0.0061	0.0376	Put	NIG
-5	3.8981	7.4880	9.5546	6.4728	Call	binormal
-5	0.1333	0.0011	0.0001	0.0484	Put	binormal
0	0.6360	2.5609	4.5424	2.0509	Call	NIG
0	1.9220	0.1175	0.0406	0.6906	Put	NIG
0	0.7152	2.6820	4.6235	2.3058	Call	binormal
0	1.9504	0.1950	0.0690	0.8815	Put	binormal
5	0.1013	0.3004	0.9295	0.4443	Call	NIG
5	6.3873	2.8569	1.4277	4.0840	Put	NIG
5	0.0429	0.2720	1.0604	0.4170	Call	binormal
5	6.2781	2.7851	1.5058	3.9927	Put	binormal
10	0.0233	0.0442	0.1746	0.1134	Call	NIG
10	11.3092	7.6008	5.6728	8.7531	Put	NIG
10	0.0008	0.0075	0.0921	0.0384	Call	binormal
10	11.2359	7.5206	5.5376	8.6140	Put	binormal
mean	-1.2860	2.4434	4.5018	1.3603	—	NIG
dev.	3.0141	2.5945	3.1629	3.8034	—	NIG
mean	-1.2352	2.4869	4.5546	1.4243	—	binormal
dev.	3.0728	2.6160	3.1761	3.8535	—	binormal

TABLE 1. Simulated option prices for various periods and strikes, with NIG and binormal process assumptions, and means and standard deviations of the terminal spread distributions

7. 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, reflecting the fatter tails of the more realistic NIG. Otherwise option evaluation is robust across the distributional assumptions.

Continuation of this study would look to other markets beyond the United Kingdom, perhaps to the Nordic Power Exchange, for evidence of similarity in parametric estimation — or for explanation why the markets may differ.

Beyond such empirical studies could lie a comprehensive theory of joint price determination in markets demonstrating dependence, and the further development of continuous time modeling of these processes.

APPENDIX A. 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 c 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, $c = 9.0$ in our model. This conversion presumes that both kinds of energy are quoted in the same monetary units, *e.g.*, pence.

Prior to 1999, the United Kingdom measured gas energy in therms, which is not an SI unit as the others are, a therm being 105.5 MJ. In such period the quotations were in pence/therm and £/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.

LIST OF FIGURES

FIGURE 1.   Copula Difference, unscaled, All Quadrants
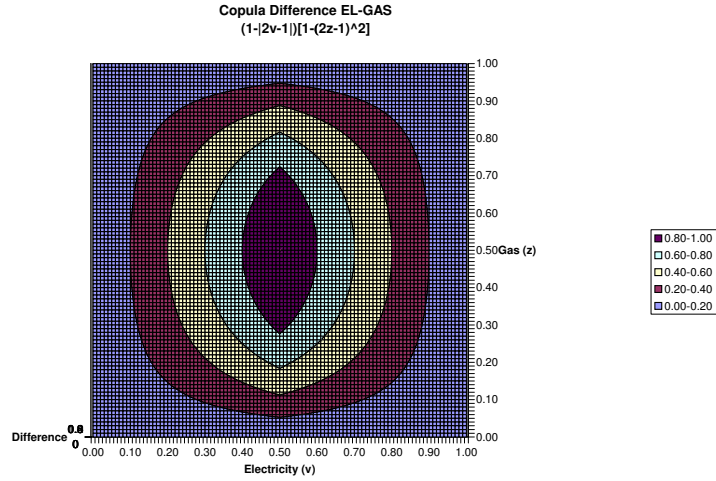
FIGURE 2.   Copula Difference, Level Curves
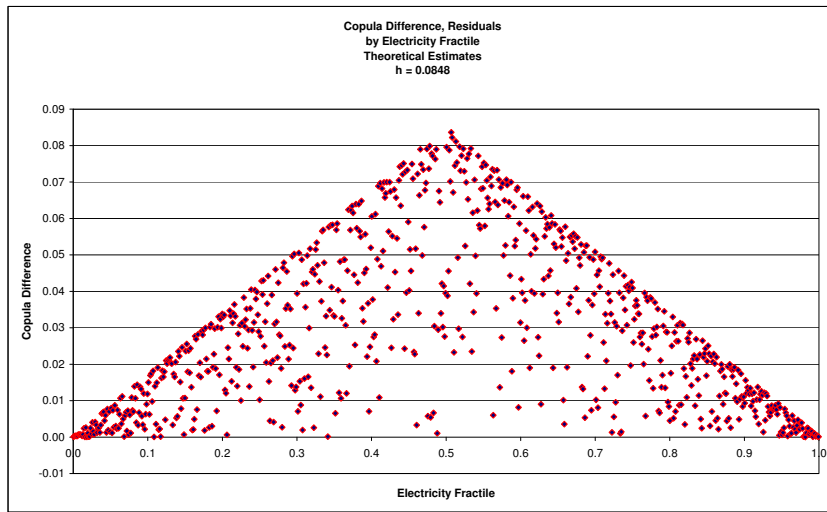
FIGURE 3. Copula Difference Electricity, Estimate
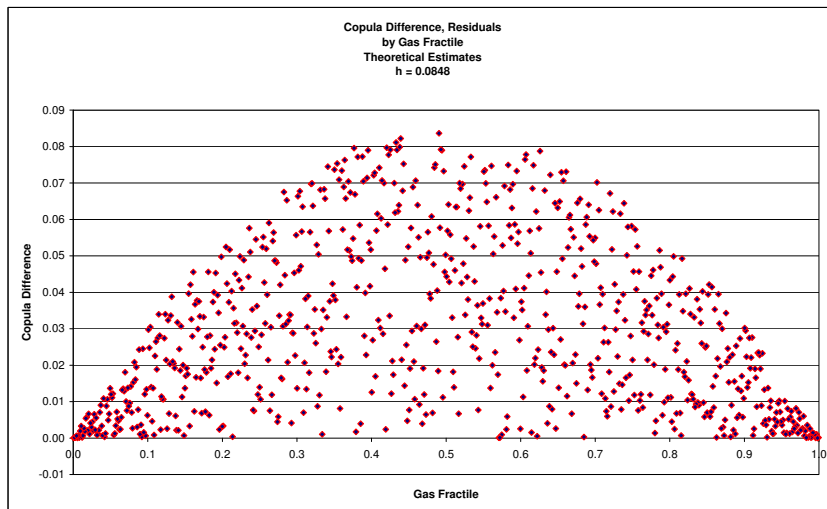
FIGURE 4. Copula Difference Gas, Estimate
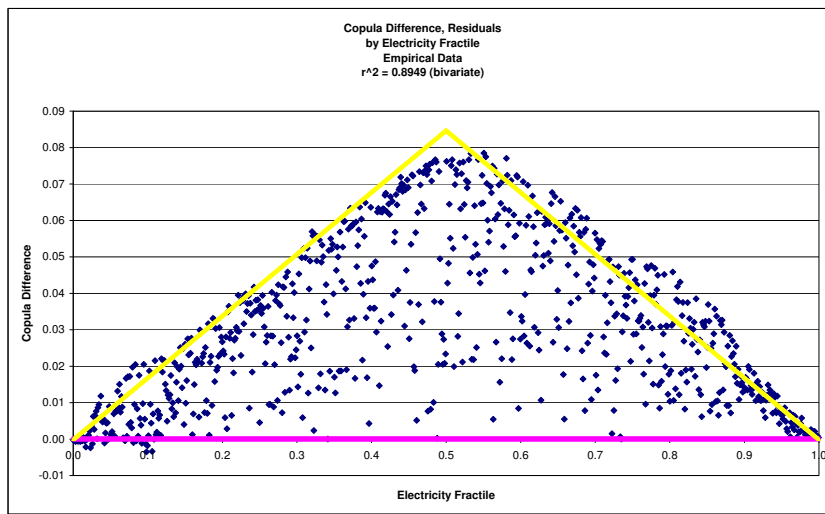
FIGURE 5. Copula Difference Electricity
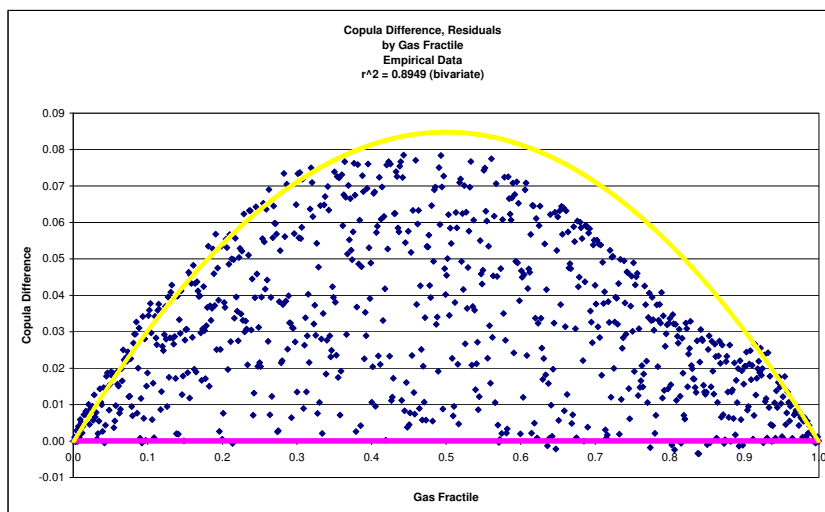
FIGURE 6. Copula Difference Gas

Empirical Copula Density, Electricity vs. Gas
Spearman's rank correlation 0.3378
Kendall's tau 0.2357
Blomqvist's beta 0.2932
Point count: LL&HH--258, LH&HL--141
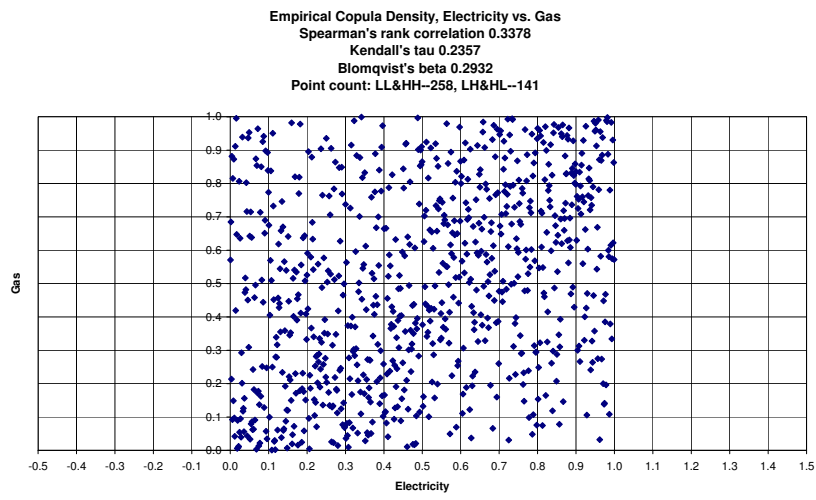
FIGURE 7.   Empirical Copula Density

Q-Q Electricity Normal
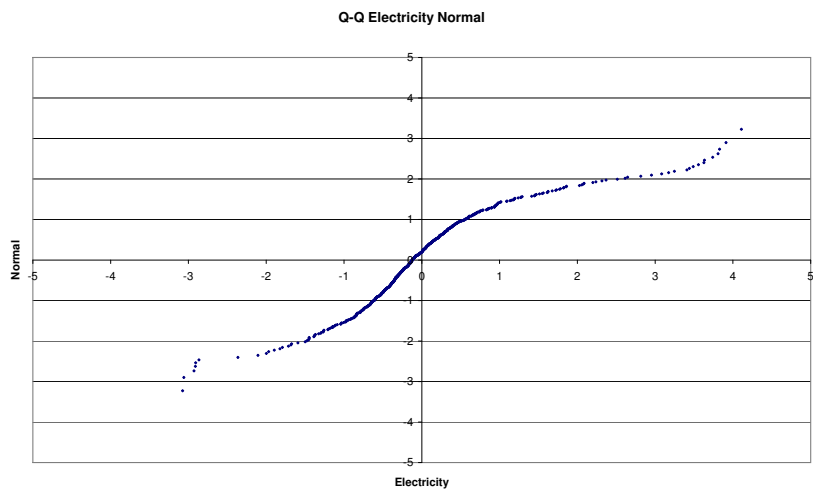
FIGURE 8.   Q-Q Electricity Normal
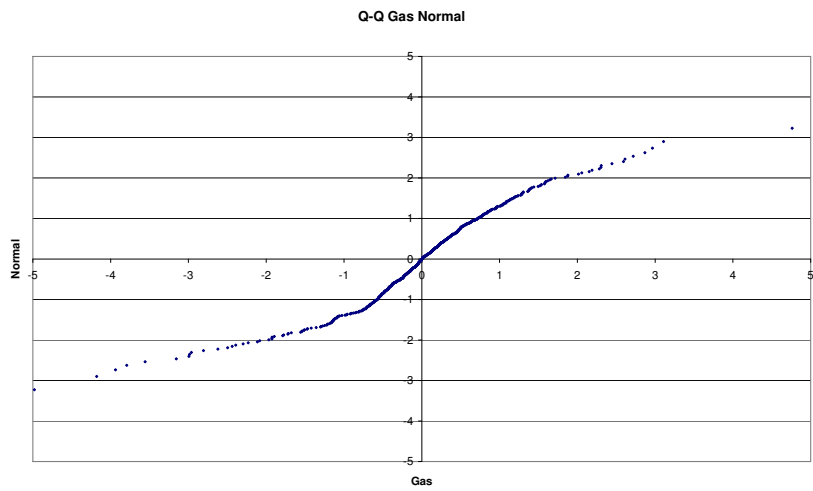
Q-Q Gas Normal

FIGURE 9.   Q-Q Gas Normal
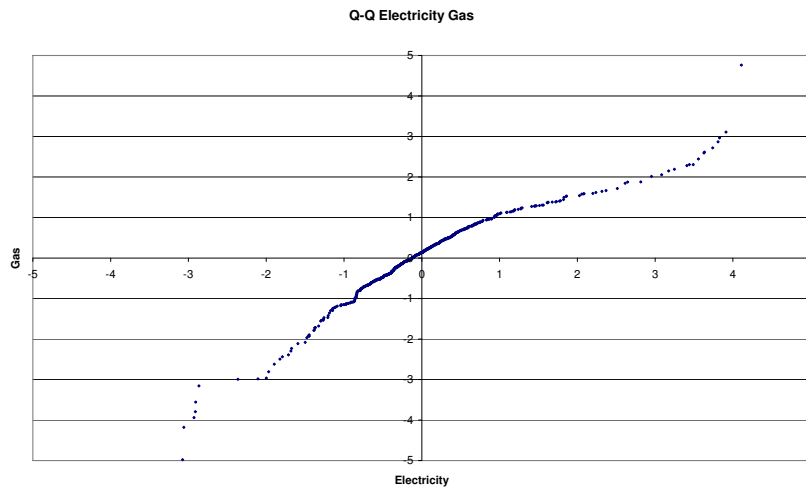
FIGURE 10. Q-Q Electricity Gas
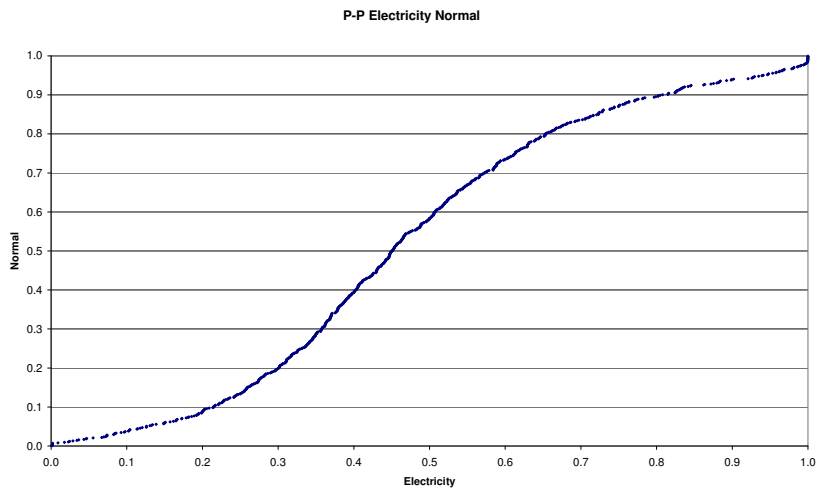
FIGURE 11. P-P Electricity Normal
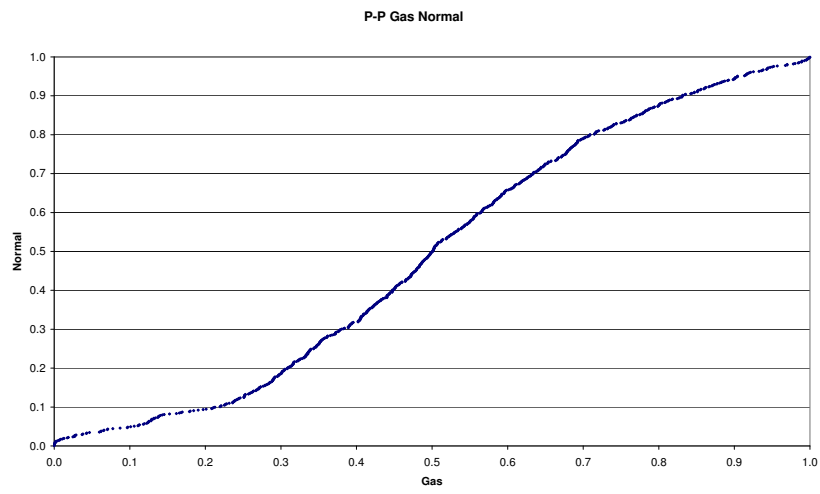
FIGURE 12.   P-P Gas Normal
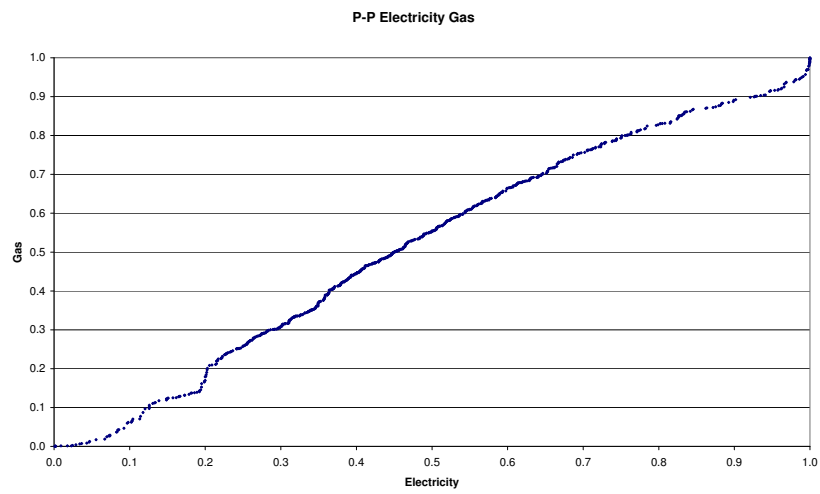
FIGURE 13.   P-P Electricity Gas
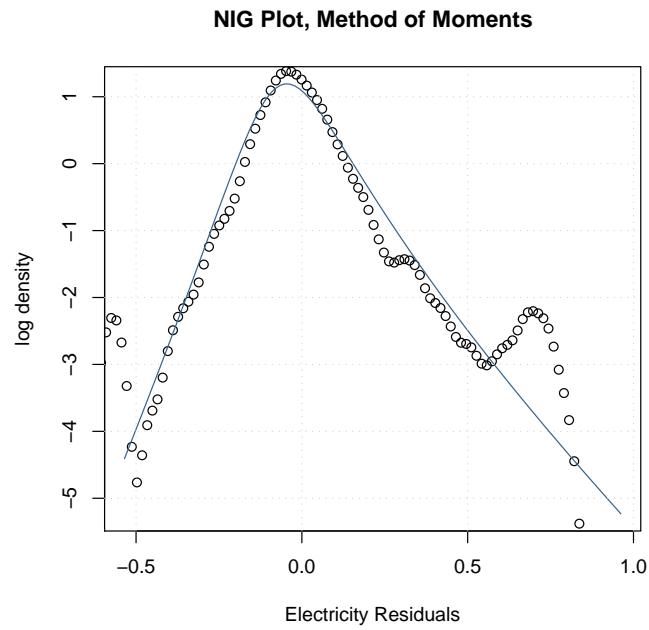
FIGURE 14. Electricity residuals, NIG fit, $(\alpha, \beta, \mu, \delta) = (6.9342, +2.8003, -0.0694, 0.1514)$
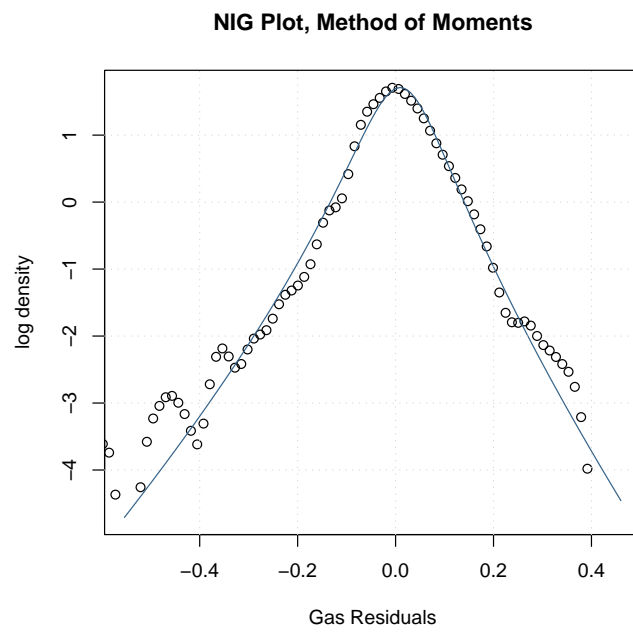
FIGURE 15. Gas residuals, NIG fit, $(\alpha, \beta, \mu, \delta) = (7.7740, -0.9982, +0.0122, 0.0831)$
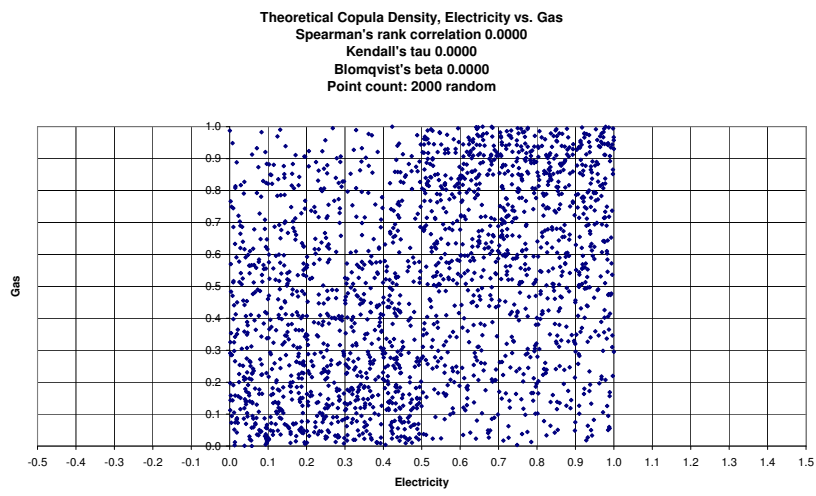
FIGURE 16.   Theoretical Copula Density

NIGs (Moment Estimates) Coupled
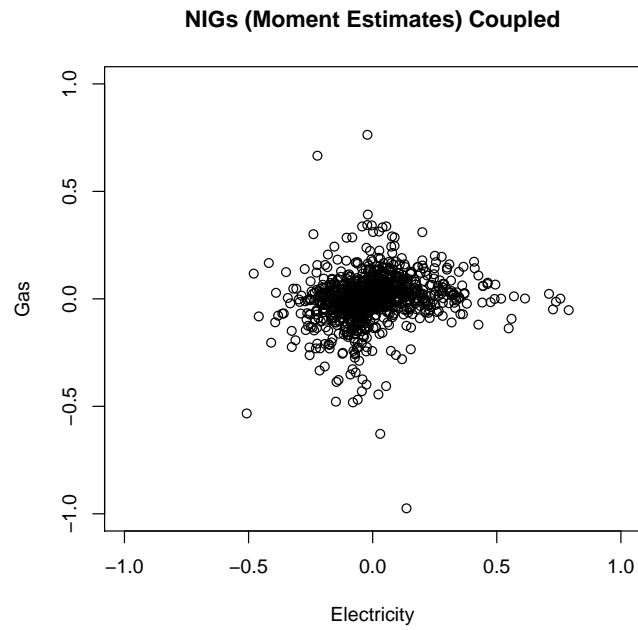
FIGURE 17. Pair estimates, 1000 points, NIG fit

Bivariate Normal (Mean, Covariance Estimates)
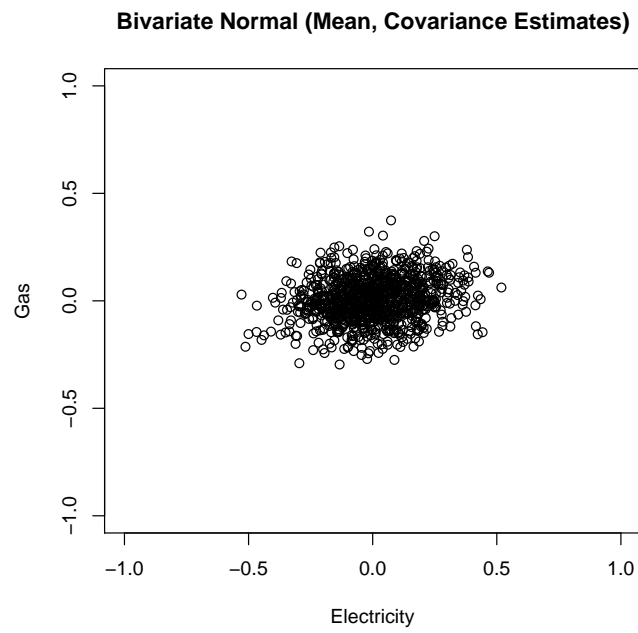
FIGURE 18. Pair estimates, 1000 points, Binormal fit

Histogram NIG 1000 Path Terminations
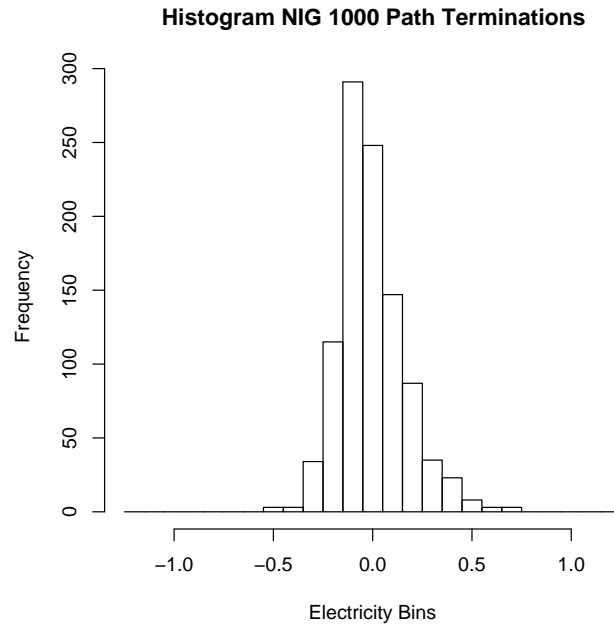
FIGURE 19.   Terminal distribution NIG, Electricity

Histogram NIG 1000 Path Terminations
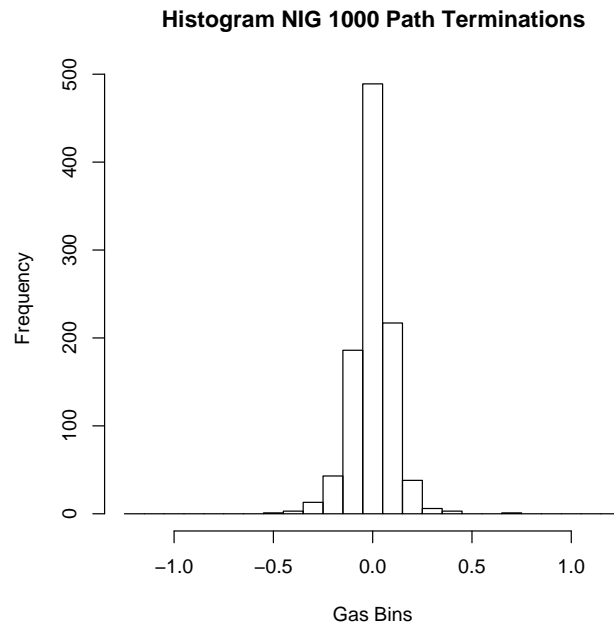
FIGURE 20.   Terminal distribution NIG, Gas

Histogram Binormal 1000 Path Terminations
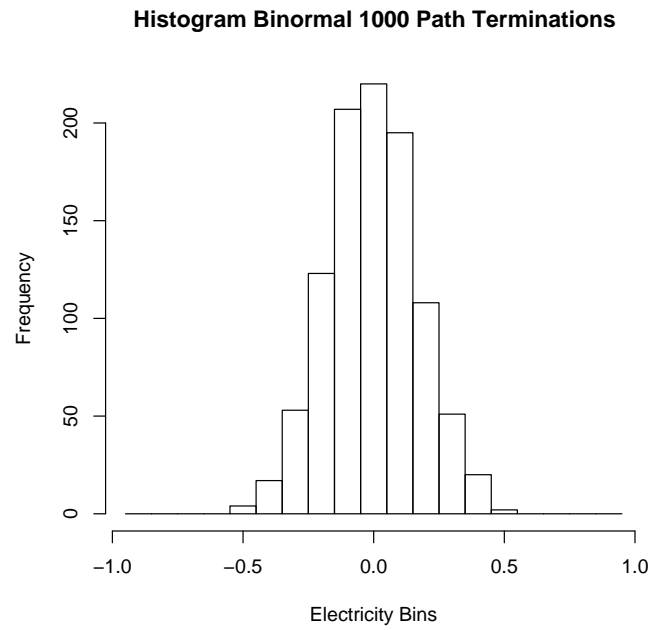
FIGURE 21. Terminal distribution Binormal, Electricity

Histogram Binormal 1000 Path Terminations
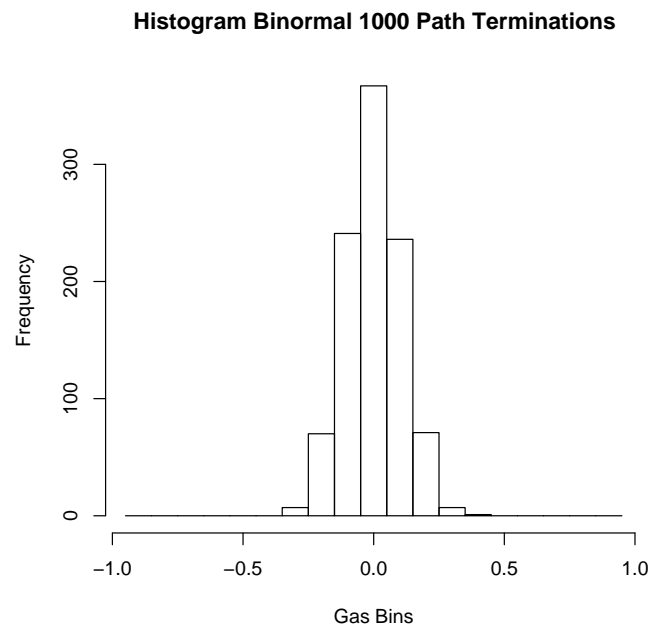
FIGURE 22. Terminal distribution Binormal, Gas

Distribution Comparison, Electricity

FIGURE 23.   NIG minus binormal terminal distributions, compared to binormal, electricity axis

Distribution Comparison, Gas
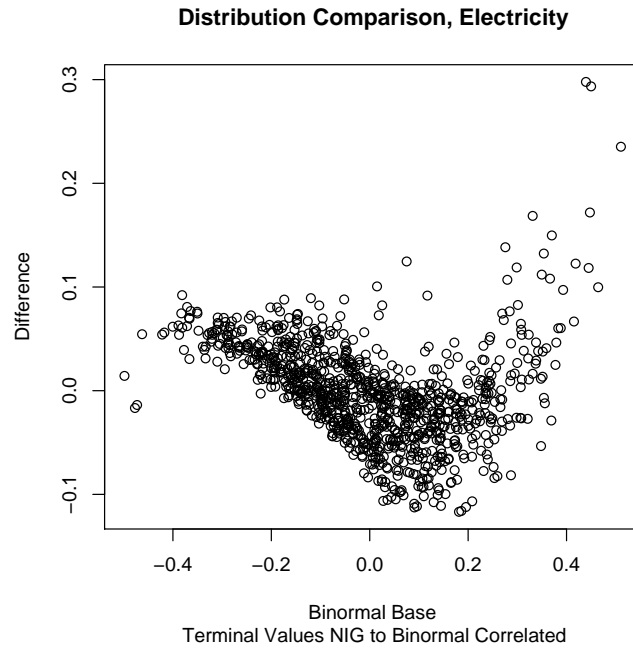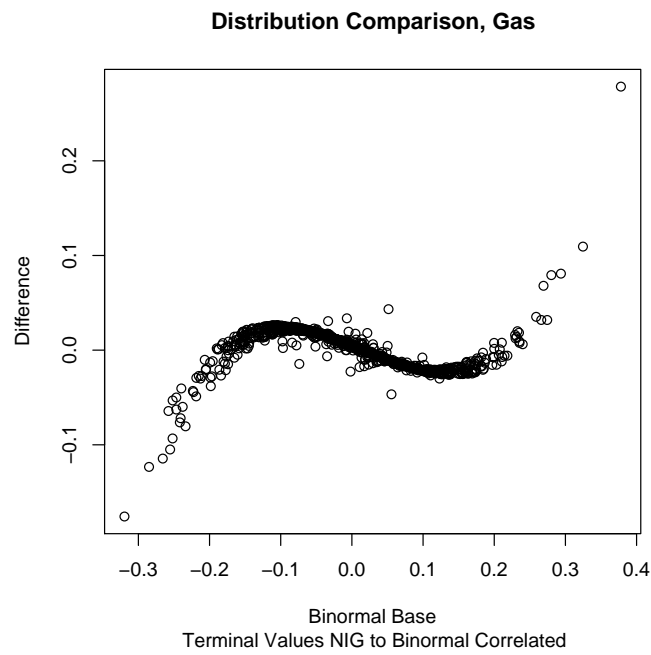

FIGURE 24.   NIG minus binormal terminal distributions, compared to binormal, gas axis
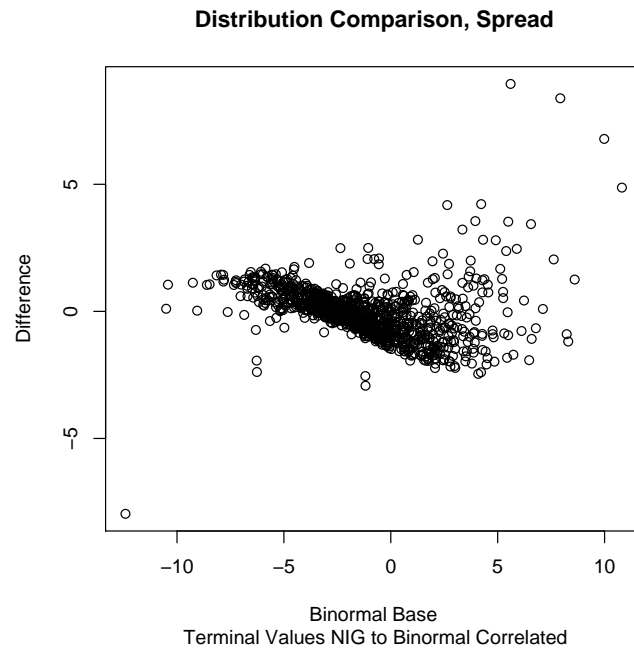
FIGURE 25. NIG minus binormal terminal spreads, compared to binormal spread, Q1

References

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

Benth, F. E. 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. To appear in special issue, "Nonlinear analysis of electricity prices," Stud. Nonlinear Dynam. Econometrics.

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

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

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

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). *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
Agder University College
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/