Cuadernos de Investigación UCEIF
Número 19/2016

COMMODITY MARKETS: ASSET ALLOCATION, PRICING AND RISK MANAGEMENT

Carlos González Pedraz

Fundación UCEIF

CANTABRIA CAMPUS INTERNACIONAL

Editorial Universidad Cantabria

COMMODITY MARKETS: ASSET ALLOCATION, PRICING AND RISK MANAGEMENT

Colección Difunde #233

Cuadernos de Investigación UCEIF Número 19/2016

COMMODITY MARKETS: ASSET ALLOCATION, PRICING AND RISK MANAGEMENT

Carlos González Pedraz

CANTABRIA CAMPUS INTERNACIONAL

Editorial Universidad Cantabria

González Pedraz, Carlos
Commodity markets : asset allocation, pricing and risk management / Carlos González Pedraz. – Santander : Editorial de la Universidad de Cantabria, D.L. 2016.
77 p. ; 24 cm. – (Difunde ; 233. Cuadernos de investigación UCEIF ; 19/2016)

En la port.: Cantabria Campus Internacional.
D.L. SA. 811-2016. – ISBN 978-84-8102-802-7

1. Mercado financiero. 2. Materias primas. I. Fundación de la Universidad de Cantabria para el Estudio y la Investigación del Sector Financiero .

330.322:620.21

IBIC: KCH, KCJ

Promueve: Fundación de la Universidad de Cantabria para el Estudio y la Investigación del Sector Financiero (UCEIF)

Coordinadora: Ana Fernández Laviada

Secretaria: Tamara Cantero Sánchez

ISBN: 978-84-8102-802-7

Depósito Legal: SA 811-2016

Imprenta Kadmos

Impreso en España. *Printed in Spain*

FUNDACIÓN DE LA UNIVERSIDAD DE CANTABRIA PARA EL ESTUDIO Y LA INVESTIGACIÓN DEL SECTOR FINANCIERO (UCEIF)

La Fundación UCEIF, promovida por la UC y Banco Santander, viene desarrollando proyectos de gran envergadura a lo largo de sus 10 años de historia, organizando su actividad en dos ámbitos de trabajo: banca y finanzas, por un lado, y actividad empresarial, con especial atención al emprendimiento, por otro. Ambas se articulan por medio de los dos centros creados en 2012: el Santander Financial Institute (SANFI) y el Centro Internacional Santander Emprendimiento (CISE).

Santander Financial Institute (www.sanfi.org)

SANFI es el centro de referencia internacional en la generación, difusión y transferencia del conocimiento sobre el sector financiero, promovido por la UC y el Banco Santander a través de la Fundación UCEIF. Desde sus inicios dirige actividades de gran calidad en áreas de formación, investigación y transferencia:

Máster en Banca y Mercados Financieros UC-Banco Santander. Constituye el eje nuclear de una formación altamente especializada, organizada desde la fundación en colaboración con el Banco Santander. Es Impartido en España, México, Marruecos y Brasil, dónde se están desarrollando la 21ª Edición, 18ª Edición, 10ª Edición respectivamente, además de clausurarse la primera promoción de la Edición Brasil. Recientemente se ha firmado el convenio de colaboración con la Pontificia Universidad Católica de Valparaíso y Santander Chile para que la Edición Chile.

Formación In Company. SANFI potencia sus actividades para desarrollar la formación de profesionales del sector financiero, principalmente del propio Santander, destacando también su actuación dentro de otros programas, como el realizado con el Attijariwafa Bank.

Archivo Histórico del Banco Santander. Situado en la CPD del Santander en Solares, comprende la clasificación, catalogación, administración y custodia, así como la investigación y difusión de los propios fondos de Banco como de otras entidades. Cabe destacar que posee más de 27.000 registros de fondo.

Educación Financiera: Finanzas para Mortales (www.finanzasparamortales.es). Proyecto educativo dirigido a fomentar la cultura financiera a través de sus plataformas online y sesiones presenciales, utilizando y aplicando las nuevas tecnologías y los medios actuales. Cuenta con más de 600 voluntarios procedentes de Banco Santander, distribuidos por los diferentes puntos de la geografía española. Han realizado, 560 sesiones formativas en 2016, donde se ha logrado acercar conocimientos financieros a más de 5.000 ciudadanos. Han colaborado con más de 50 instituciones, destacando colegios e institutos, Cáritas, Cruz Roja, Fundación del Secretariado Gitano, la ONCE, Fundación Integra, Ayuntamientos en las que han contribuido a mejorar la cultura financiera de beneficiarios y empleados.

Investigación

- Atracción del Talento, con diferentes acciones para el desarrollo de líneas de investigación estratégicas dedicadas al estudio de los “Mercados Globales”, al desarrollo e innovación de “Procesos Bancarios” al conocimiento de la “Historia Bancaria y Financiera”.

- Becas de investigación, con la finalidad de colaborar en la realización de Proyectos de Investigación, especialmente de Jóvenes Investigadores, que posibiliten el avance en el conocimiento de las metodologías y técnicas aplicables en el ejercicio de la actividad financiera, en particular las que llevan a cabo las entidades bancarias, para mejorar el crecimiento económico, el desarrollo de los países y el bienestar de los ciudadanos.

- Premios Tesis Doctorales, con el fin de promover y reconocer la generación de conocimientos a través de actuaciones en el ámbito del doctorado que desarrollen, impulsen el estudio y la investigación en el Sector Financiero.

- Y por último, la línea editorial, en la que se enmarcan estos Cuadernos de investigación, con el objetivo de poner a disposición de la sociedad en general, y de la comunidad académica y profesional en particular, el conocimiento generado en torno al Sector Financiero fruto de todas las acciones desarrolladas en el ámbito del Santander Financial Institute y especialmente los resultados de las Becas, Ayudas y Premios Tesis Doctorales.

ÍNDICE

I. INTRODUCTION

Since the mid-2000s, many investors have entered into commodity markets; investment banks, hedge funds, and other portfolio managers increasingly view commodities as an alternative asset class. Prior to the financial meltdown of 2008-2009, financial analysts recommended long-only investments in commodities as an alternative asset class, to decrease expected portfolio risk, increase expected portfolio returns, and hedge macroeconomic risk. During that period, commodity prices and volatility increased significantly, as we can see in Figure 1.1.

In addition to futures and options contracts, other financial products that allowed investors to gain exposure to commodities –such as commodity index funds, over-the-counter (OTC) swap agreements, and exchange traded funds– were also widely popularized. According to some estimates, index investment in commodities increased from around \$15 billion at the end of 2003 to more than \$200 billion in 2008, just prior to the financial crisis.

The notional amounts outstanding in OTC markets and the open interest of futures contracts (see Exhibits 1 and 2 of Figure 1.2, respectively) experienced rapid growth in the last decade. Furthermore, players other than traditional producers and retailers started to trade physical assets. For example, some investment banks owned power plants and pipelines and purchased other commodity assets to use as hedging tools.

The recent volatility of commodity prices and the rise in commodity investing also renewed academic interest in the behavior of commodity markets. For that reason, a better understanding of the behavior of commodity prices from a multivariate and univariate perspective, especially when extreme events occur, is of paramount importance for accurate asset valuation, risk management, and portfolio decisions. This monograph tries to contribute in this respect.

In this opening chapter, we describe commodity markets, paying special attention to energy commodities.

1.1. Commodity Markets

In a financial context, the term "commodity" refers to a relatively homogeneous consumption good. Commodities differ from stocks and bonds in that they do not generate a stream of future cash flows (Geman, 2005).

Figure I.1. Relative price and moving volatility of the CRB index

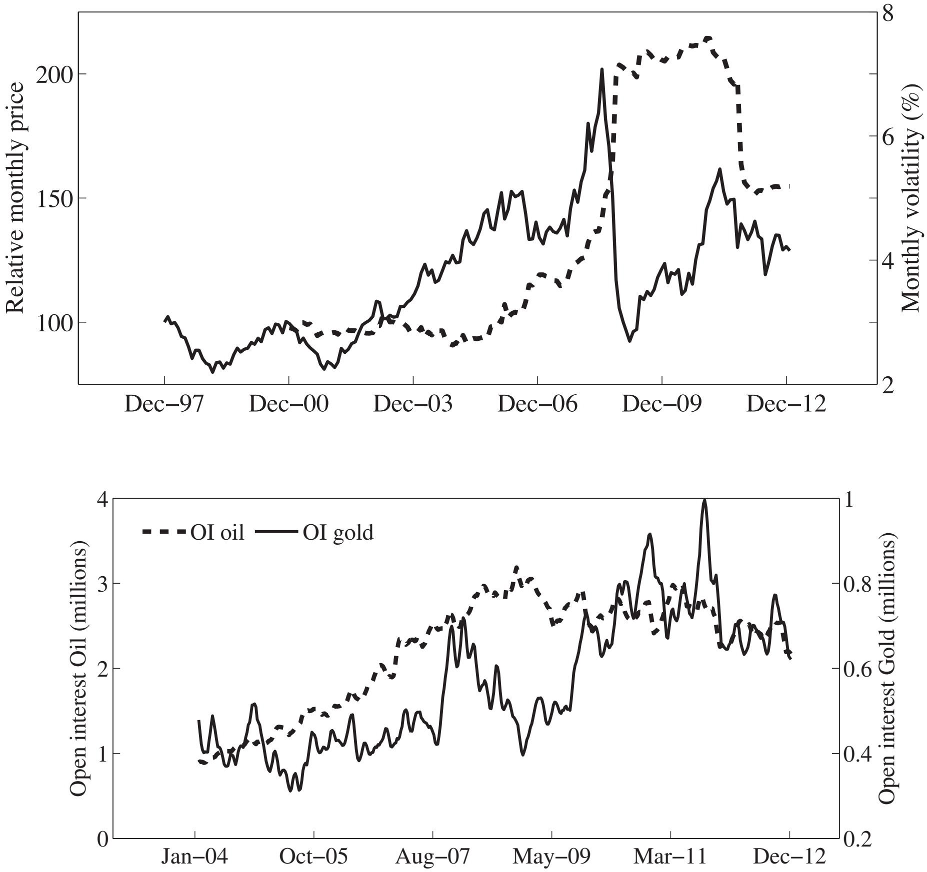

Grains, livestock, energy, metals, foodstuffs, and textiles are traditional examples of commodity classes. Weather, carbon dioxide emission allowances, and computing resources are examples of new commodity markets.

There is considerable diversity among commodities. Most are storable at some cost, but some, such as electricity are impossible or very costly to store. Supply and demand patterns also establish differences among commodities. Some commodities exhibit substantial seasonality in demand (e.g., natural gas, electricity), whereas other commodities are produced periodically (e.g., grains). Thus, storability, demand, and supply characteristics determine the different behavior of commodity prices and present a challenge to modelers (Pirrong, 2012).

Buyers and sellers can trade commodities in spot markets, where delivery is immediate or following a very small lag; or they can trade them using forward agreements with a given future delivery date, either in organized futures exchanges or OTC markets. Commodity forward contracts allow firms to obtain insurance for the future value of their outputs or inputs, whereas investors in these contracts receive compensation for bearing the risk of short-term commodity price fluctuations Gorton and Rouwenhorst (2006).

There are many active and liquid commodity futures markets, including crude oil, heating oil, natural gas, gold, silver, copper, and aluminum futures traded in the New York Mercantile Exchange (NYMEX); corn, soybean, and wheat futures traded in the Chicago Board of Trade (CBOT); non-ferrous metals in the London Metal Exchange (LME); and oil, natural gas, electricity, freight, and agriculturals in the Intercontinental Exchange (ICE) in London. Other commodity exchanges in emerging markets have gained importance in recent years, especially the Dalian Commodity Exchange and the Shanghai Futures Exchange for agriculturals and non-precious metals trading.

Figure 1.2. Notional amounts and open interest

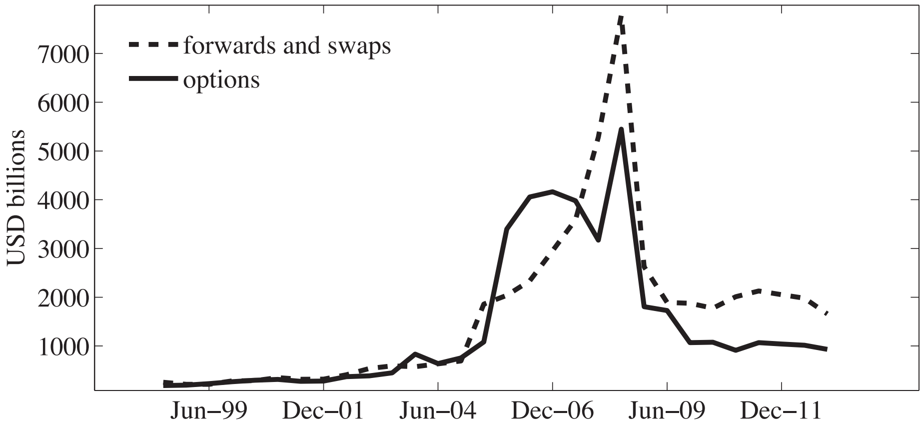

1.2. Energy spot and futures markets

The market for energy is huge. The world's population consumes about 15,000 gigawatts of power (1 gigawatt is the capacity of the largest coal-fired power station). That means a business of $6 trillion a year, one-tenth of the world's economic output. Energy markets have been liberalized in the recent years and are still developing. The crude oil market is the most liquid and global commodity market. Other important energy markets are power, natural gas, and coal. Natural gas is used for heating purposes and as an input for power generation. In Western Europe and the United States, coal is mainly used for power generation.

Since the end of last century, there has been a continuous process of liberalization and deregulation of energy markets. These developments have resulted in the separation of services into generation, transmission, distribution, and retail, with the goal of creating more competition and liquidity. They also have created new market risk exposures, which must be managed at every interface of the formerly integrated chain. The players in this growing risk management market consist of two major categories: physical and financial players. The first are engaged in the energy markets due to their physical exposure, such as utilities and oil and gas producers. Most actively engage in hedging and trading activities, both physically and financially. The second group comprises by investment banks, hedge funds, and trading houses, which act as risk managers and intermediaries or trade in their own accounts.

Energy markets exhibit some special features that differentiate them from other commodity markets. Natural gas or power markets have a great variety of traded contracts, including forwards with multiple maturities and different delivery periods and short-term contracts, such as weekly, day-ahead, and real-time prices. Electricity cannot be stored, and natural gas storage is costly and inflexible. Therefore, there is no clear price convergence among the different contracts. Location is also very important in these markets, because each gas or power hub has a different price every day and hour.

When modeling the evolution of gas and power prices, we also must consider the presence of extreme price spikes, seasonality, and mean reversion. Furthermore, energy markets are closely intertwined by substitution, complementary, and production relationships that complicate the modeling of the dependence among these commodities Casassus, Liu, and Tang (2013).

Some commodity derivatives and physical assets that are operated in these markets have complex payoffs. Asian, spread, and swing options are examples of exotic options that provide a hedge against price and volume risks. Such options also appear in the real option approach to the valuation of physical assets and contracts, such as power plants, interconnections, or gas storage.

1.3. Modelling and estimation methodologies

Because of the special characteristics of commodities, pure unconditional Gaussian models are not the most suitable framework to describe the multivariate and univariate behavior of these assets. Consequently, traditional approaches to deal with commodity market risk and price commodity derivatives have to be reconsidered. Furthermore, in this context of non-normality, nonlinear econometric techniques would be required to estimate realistic commodity pricing models. Some of the most promising distributions for modeling heavy-tailed and asymmetric returns are those belonging to the generalized hyperbolic (GH) class, which exhibit some attractive properties: They are closed under affine transformations, can display different tail behavior, and can be both symmetric and skewed.

When the joint distribution of returns is non-elliptical, the linear correlation is no longer sufficient to describe the dependence structure, and nonlinear dependence functions can be more informative. Copula theory allows us to obtain the dependence function, or copula, of a d-dimensional joint distribution, and use it to model the dependence between d arbitrary univariate densities. For example, we can distinguish between the tails of the marginal distributions and the presence of dependence in the tails, as well as between asymmetry in the distribution of individual returns and asymmetry in their dependence structure.

From a continuous-time perspective, we also consider the presence of heavy tails and skewness in the dynamics of commodity prices. A no-arbitrage spot price model, especially for energy commodities, should capture: large price spikes or jumps, strong mean reversion of large deviations, and the presence of a seasonal component. The presence of jumps in the dynamics of prices prevents us from estimating the model parameters using traditional techniques based on Gaussian hypothesis. Other approaches such as non-parametric jump filters or simulation methods have to be employed.

Because of the presence of discontinuous processes, establishing a link between the data generating measure and the risk-neutral measure is more difficult than in traditional pure Gaussian cases. Furthermore, when analyzing the behavior of commodity forward contracts, some questions arise. For example, how are the data generating and risk-adjusted measures related? Spot price dynamics are much easier to model than dynamics under the risk-adjusted measure, because the former are observed, whereas the latter can only be inferred from the price dynamics of instruments written on the spot commodity, such as forward contracts.

1.4. Organization of this monograph

This study is divided into two parts. The first part analyzes the multivariate distribution of commodity returns and its impact on portfolio selection and tail risk measures. Chapter 2 solves the portfolio selection problem of an investor with three-moment preferences when commodity futures are part of the investment opportunity set, providing a conditional copula model for the joint distribution of returns that allows for time-varying moments and state-dependent tail behavior. Chapter 3 approximates the exposure of physical and financial players to energy price risk using linear combinations of energy futures; it also analyzes the tail behavior of energy price risk using a dynamic multivariate model, in which the vector of innovations is generated by different generalized hyperbolic distributions.

The second part considers the valuation of real assets and commodity derivatives in the presence of non-Gaussian shocks in a continuous time framework.

Specifically, Chapter 4 employs a jump diffusion model for the price differentials and proposes a valuation tool for the connection between two electricity markets.

Chapter 5 proposes a reduced-form model for the data generating process of commodity prices together with a more flexible change of measure, capable of changing the mean-reversion rate of Gaussian and jump processes under the risk-adjusted probability measure.

2. PORTFOLIO SELECTION WITH COMMODITIES

Financial investors mainly take positions in commodity futures contracts as a natural way to gain exposure to commodity risk without owning the physical asset. Erb and Harvey (2006) and Gorton and Rouwenhorst (2006) find that historically, commodity futures exhibited little co-movements, zero or even negative correlations with stock returns, and Sharpe ratios fairly close to those of equities. Therefore, according to traditional portfolio theory, commodities should increase diversification when included in equity portfolios and may help enhance the portfolio's risk-return profile. Possibly boosted by the potential for such diversification benefits, investments in commodity futures indexes and related instruments grew quickly after the early 2000s (see Büyüksahin and Robe (2010), Etula (2013), Hong and Yogo (2012), and Tang and Xiong (2012) for some analysis about this recent boom).

Despite the growing interest in commodities as investment vehicles, few studies have analyzed the optimal portfolio allocation taking into account the stylized features of commodity futures. A standard mean-variance framework might not be appropriate for portfolios that contain commodity futures due to their returns' specific distributional characteristics, such as the presence of serial correlation, heavy tails, and skewness (Daskalaki and Skiadopoulos (2011); Gorton and Rouwenhorst (2006); Kat and Oomen (2007b); Börger, Cartea, Kiesel, and Schindlmayr (2009)). Instead, we propose a more flexible model to be used in the portfolio selection problem of a traditional equity investor when cash-collateralized commodity futures are part of the investment opportunity set. Our approach combines a three-moment preferences specification with time-varying multivariate density models that describe the statistical properties of commodities and equity returns, as well as their interactions.

With respect to the investor's preferences, we consider an allocation problem in which the investor's objective function is determined by the mean, variance, and skewness of portfolio returns (similarly to Guidolin and Timmermann (2008); Harvey, Liechty, Liechty, and Müller (2010); and Jondeau and Rockinger (2012)). With fairly general assumptions, investors show a preference for positive skewness in return distributions and aversion to downside risk (negative skewness). That is, in our proposed three-moment preferences specification, the investor is eager to decrease the chance of large negative deviations, which could reduce the final value of the portfolio. Skewness seems likely to play a role due to the specific features of commodity assets. For instance, the possibility of shortages in supply may produce jumps in prices, leading to skewness in the returns of futures contracts.

Regarding the multivariate density model, we offer a flexible approach to specify the joint distribution of returns using conditional copula models. Copula functions help disentangle the particular characteristics of the univariate distributions of equity and commodity returns from their dependence structure. We combine conditional copula theory, as presented in Patton (2006a, b), with the implicit copula functions of multivariate normal mixtures defined by Demarta and McNeil (2005) and Embrechts, Lindskog, and McNeil (2003). As our most general model, we propose a conditional skewed *t* copula with generalized Student's *t* marginal distributions. This copula model allows for asymmetric and tail dependence in a multivariate framework, and includes symmetric and linear dependence as special cases. Furthermore, the conditional set-up enables us to capture time-varying investment opportunities through time-varying moments and changes in the dependence parameters. These copula models are particularly easy to sample from, and therefore, we opt for solving numerically the optimization problem using Monte Carlo simulations.

We apply our theoretical approach to weekly data of crude oil and gold futures and the S&P 500 equity index, for the period from June 1990 to September 2010, reserving the observations from September 2006 to September 2010 for an out-of-sample performance evaluation. We examine four primary issues:

1. Is there asymmetric and tail dependence among commodities and equity returns?

2. Are there discrepancies in the optimal portfolio allocations between our conditional copula approach and other more traditional benchmarks, such as the equally weighted or Gaussian strategies?

3. Do these discrepancies translate into economically relevant performance differences among methods?

4. Is there a single key factor explaining these discrepancies?

First, our preliminary statistics and in-sample and out-of-sample results show evidence in favor of heavy tails and skewness in the univariate behavior and extreme and asymmetric dependence among oil, gold, and equity. Second, we also uncover substantial discrepancies between portfolio optimal weights of conditional *t* copulas and the portfolio weights provided by more conventional alternatives, especially for more aggressive investors and when there are no restrictions on short selling positions in equity. Third, in most cases, the discrepancies in portfolio weights translate into economically more profitable investment ratios and better relative performance measures with respect to the alternative procedures. Depending on the investor's preferences specification, the gains of considering the conditional copula model with tail and asymmetric dependence instead of the equally weighted portfolio represent up to 86 basis points per year for the period 2006-2010. When variance and loss aversion increase, portfolio strategies based on more flexible copulas are less likely to produce large performance differences. Fourth, no single factor offers a sufficient explanation of these differences. Rather, we find various explanatory elements, including, the specification of the univariate processes, in terms of conditional volatility, skewness, and fat tails; and the presence of tail and asymmetric dependence.

The remainder of this chapter is organized as follows: In Section 1, we present briefly the investor's objective function and the portfolio choice problem and describe the multivariate conditional copula model. Section 2 presents the in-sample estimations and the out-of-sample portfolio allocation results. We conclude in Section 3.

2.1. Investor's problem

In this section, we present the portfolio selection problem of an investor with mean-variance-skewness preferences that takes positions in commodity futures and other risky spot contracts, such as stocks.

No money changes hands when futures are sold or bought; just a margin is posted to settle gains and losses. Without any upfront payment, it is not clear how to define the rate of return. Following the common approach in the literature to analyze commodity futures as an asset class (Gorton and Rouwenhorst (2006); Hong and Yogo (2012)), we assume that long and short positions are fully collateralized. That is, the initial margin deposit corresponds with the overall notional value of the futures contract and indicates the initial capital investment related to that position (long or short). In this way, we control for the leverage involved in futures positions, and we can make fair comparisons with spot contracts. Therefore, taking collateral in futures contracts into account would affect the computation of their rates of return and the budget constraint of the investor's problem, as we will see.

Formally, our portfolio choice problem can be formulated in terms of an investor who maximizes expected utility at period *t+1* by building at time *t* a portfolio that includes two group of assets: a group with *n* commodity futures contracts, and another group with *N-n* spot contracts, such as stocks.

For this set of *N* investment opportunities, the wealth at time *t+1* equals the gross return of the portfolio over the period, defined as

$$1 + R_{t+1}(\omega_t) = 1 + \sum_{j=1}^{N} \omega_t^j (\exp(r_{j,t+1}) - 1) \tag{2.1}$$

where ω_t is the vector of portfolio weights (for spot and futures contracts), chosen at time *t*, and $r_{j,t+1}$ is the continuously compounded return of asset *j* over the period.

As is well known, returns on financial assets generally deviate from the Gaussian distribution, displaying heavy tails and skewness. This departure from normality is even greater in the case of commodities, magnified by the well-documented presence of positive and negative spikes in the data-generating process of commodity returns (see for example Cartea and Figueroa (2005) and Casassus and Collin-Dufresne (2005), among others). The fundamentals underlying commodity price formation are key determinants of these statistical properties. Accordingly, the presence of jumps can be explained by the convex relation between commodity prices and the balance among supply, inventories, and demand (see Routledge, Seppi, and Spatt (2000)).

For that reason, in our approach, the investor's objective consists of choosing a wealth allocation ω_t that maximizes the expected portfolio return penalized for the variance and negative skewness of the portfolio returns. That is, for each time t, the optimal weights are given by

$$\omega_t^* = \arg\max_{\omega_t \in \mathcal{D}} \Big(\mathrm{E}_t[R_{t+1}(\omega_t)] - \varphi_{\mathrm{V}} \operatorname{Var}_t[R_{t+1}(\omega_t)] + \varphi_{\mathrm{S}} \operatorname{Skew}_t[R_{t+1}(\omega_t)] \Big) \tag{2.2}$$

where E_t, Var_t and $Skew_t$ are the first three moments of the portfolio returns conditioned on the information set available at time t. The parameters $\varphi_v, \varphi_v \geq 0$ determine the impact of variance (traditional risk aversion) and skewness (loss aversion) on the investor's utility. By adding aversion to negative skewness, we acknowledge the possibility that an investor might accept a lower expected return if there is a chance of high positive skewness, such as in the form of a large probability of positive jumps.

Finally, in equation (2.2) the domain $\mathcal{D} \subset \mathbb{R}^N$ represents the budget constraint defined by

$$\mathcal{D} = \left\{ (\omega_t^1, \ldots, \omega_t^N) : \sum_{j=1}^{N-n} \omega_t^j + \sum_{i=N-n+1}^{N} \left|\omega_t^i\right| = 1 \right\} \tag{2.3}$$

Because both long and short positions in commodity futures contracts require the same initial collateral, we have to take the absolute value of the futures weights such that short positions in futures contracts cannot be used to increase holdings of other assets.

2.2. Conditional copula model

Once we have the set-up of the investor's problem, we need to define density forecasts of the returns joint distribution in order to compute the optimal portfolio weights. In particular, we employ multivariate conditional copulas to obtain a flexible model for the multivariate distribution of assets' log-returns vector r_{t+1} with dimension d (the number of risky assets). Every d-variate distribution consists of d marginal distribution functions or *margins* that describe each univariate behavior, as well as a joint dependence function that defines the relations among individual processes. Unlike traditional multivariate distributions, such as the Gaussian and Student's t distributions, copula models support the construction of multivariate distributions with arbitrary univariate processes and dependence.

Formally, a d-variate copula is a d-dimensional distribution function on the unit interval $[0,1]^d$, that is, a joint distribution with d uniform marginal distributions. Consider a multivariate conditional distribution $F_t(r_{1,t+1}, \ldots, r_{d,t+1})$ formed by d univariate conditional distributions $F_{i,t}(r_{i,t+1})$, where the subscript t denotes that joint and marginal distributions are conditioned on the information set available at time t. Following Patton (2006b), there must exist a function C_t that maps the domain $[0,1]^d$ toward the interval $[0,1]$, called the *conditional copula*, such that

$$F_t(r_{1,t+1}, \ldots, r_{d,t+1}) = C_t\left(F_{1,t}(r_{1,t+1}), \ldots, F_{d,t}(r_{d,t+1})\right) \tag{2.4}$$

Using the expression in equation (2.4), any copula C_t can be employed to define a joint distribution $F_t(\boldsymbol{r}_{t+1})$ with the arbitrary marginal distributions $F_{1,t}, \ldots, F_{d,t}$. Thus, using a bottom-up approach, we model the marginal distributions of asset returns, followed by the conditional copula function that describes their dependence structure.

Our multivariate copula model supports the use of various marginal distributions. Thus we can attend to the particular characteristics of each asset return, which is a useful feature when different types of assets appear in the portfolio, such as commodities and stocks. We present a marginal distribution model that captures individual skewness and heavy tails, as well as time-varying moments. We build on the autoregressive conditional density models of Hansen (1994), Harvey and Siddique (1999), and Jondeau and Rockinger (2003), and we propose a generalized Student's *t* distribution with possibly time-varying parameters. Thus, the univariate process for each asset returns $r_{i,t+1}$ ($i=1,\ldots ,d$) can be expressed as follows:

$$r_{i,t+1} = \mu_{i,t+1} + \sqrt{\sigma^2_{i,t+1}}\, z_{i,t+1} \tag{2.5}$$

$$z_{i,t+1} \sim g_{i,t}(z_{i,t+1}\,;\,\nu_{i,t+1}\,,\lambda_{i,t+1}) \tag{2.6}$$

The conditional mean $\mu_{i,t+1}$ is a linear function of lagged returns and other possible explanatory variables. This specification can capture the possible presence of autocorrelation and predictability in asset returns. As the exogenous regressors we consider explanatory variables employed in previous literature (Hong and Yogo (2012)) to predict variation in stocks and commodity futures returns, including the short rate, default spread, momentum, basis, and growth in open interest. For the conditional variance $\sigma^2_{i,t+1}$, we employ an asymmetric or leveraged GARCH dynamic. This specification is designed to account for volatility clustering and leverage effects, such as possible asymmetric responses to positive and negative shocks that have occurred in the previous period. Finally, the univariate innovations $z_{i,t+1}$ are drawn from a generalized Student's *t* distribution, $g_{i,t}$, which can capture heavy tails and individual skewness through the degrees of freedom ν_i and asymmetry parameter λ_i.

This general specification also includes some well-known univariate distributions as particular cases. For instance, if the asymmetry parameter goes to 0, we obtain the symmetric Student's *t* distribution; as degrees of freedom tend to infinity, we would converge to a Gaussian distribution.

Now, we present the copula functions that determine the dependence structure of our model. The copula function acts like a joint distribution of the probability transformed vector $(F_{1,t}(r_{1,t+1}), \ldots, F_{d,t}(r_{d,t+1}))'$, where $F_{i,t}(r_{i,t+1}))'$ are the marginal distribution functions of asset returns $r_{i,t+1}$, as described in equations (2.5)-(2.6). In particular, we employ three multivariate copula functions: two well-known elliptical copulas, the Gaussian and the *t* copula (Embrechts, Lindskog, and McNeil (2003)), and an asymmetric multivariate dependence, the so-called skewed t copula (Demarta and McNeil (2005)). They are all implicit dependence functions of various multivariate normal mixtures. More specifically, they are the parametric copula functions contained in the multivariate Gaussian, Student's *t*, and generalized hyperbolic skewed *t* distributions, respectively. We can obtain these implicit copulas by evaluating a given multivariate distribution (e.g., generalized hyperbolic skewed *t*) at the quantile functions of its corresponding marginal distributions.

For illustrative purposes, in Figure 2.1 we present the contour plots and probability density functions of these copulas for a two-dimensional case. Although the examples in Figure 2.1 are for a bivariate case, a useful property of all three copulas considered is that they can be employed directly to specify the dependence structure of an arbitrary number of risky assets.

As Figure 2.1 reveals, using these three copulas, we can model three different types of dependence. The Gaussian copula, $C^G(\cdot\,;\boldsymbol{P})$, defines linear, symmetric dependence, completely determined by the correlation matrix $\boldsymbol{P}$. Thus it is unable to capture tail dependence or asymmetries. The *t* copula, $C^T(\cdot\,;\boldsymbol{P},\nu)$, is also elliptically symmetric but allows for tail dependence through the degrees-of-freedom parameter, ν. The plots in Figure 2.1 show that the *t* copula assigns more probability to the extremes than does the Gaussian copula. The greater the degrees of freedom, the smaller the level of tail dependence, converging in the limit $\nu \to \infty$ to the Gaussian copula. Finally, the skewed *t* copula, $C^T(\cdot\,;\boldsymbol{P},\nu,\boldsymbol{\gamma})$, can capture extreme and asymmetric dependence of the asset returns. Through the d-dimensional vector of asymmetry parameters $\boldsymbol{\gamma}$, the skewed *t* copula can assign more weight to one tail than the other. For example, in Figure 2.1 all elements of the asymmetry vector are negative, and therefore, the density contour is clustered in the negative-negative quadrant. Eventually, if $\boldsymbol{\gamma} \to \boldsymbol{0}$, asymmetric dependence goes to 0, and we recover the symmetric *t* copula.

In addition, following pioneering works by Patton (2006a,b) we can parametrize time variation in the conditional copula function of our multivariate model. For that purpose, we allow that the dependence matrix $\boldsymbol{P}_t$ of our conditional copula may evolve over time, according to some GARCH-type process.

Our model structure, formed by the marginal distributions and the copula, allows for a two-step estimation procedure. In the first step, we obtain the maximum likelihood (ML) estimates of the individual processes; then, we determine the parameter estimates of the copula function. From this ML approach, we can compute the asymptotic and robust standard errors for the estimates.

Once we have estimated the model density function, we use this information to obtain the optimal portfolio. For our parametric density models, the integrals defining the portfolio return moments involved in the investor's optimization problem of equation (2.2) do not have a closed-form solution. Using Monte Carlo simulations to estimate the value of these integrals, we can solve numerically the optimization problem. In this respect, an advantage of our implicit copulas is that it is easy to sample from them, as long as we are able to sample from the normal mixture distribution from which they are extracted.

Figure 2.1. Copula functions

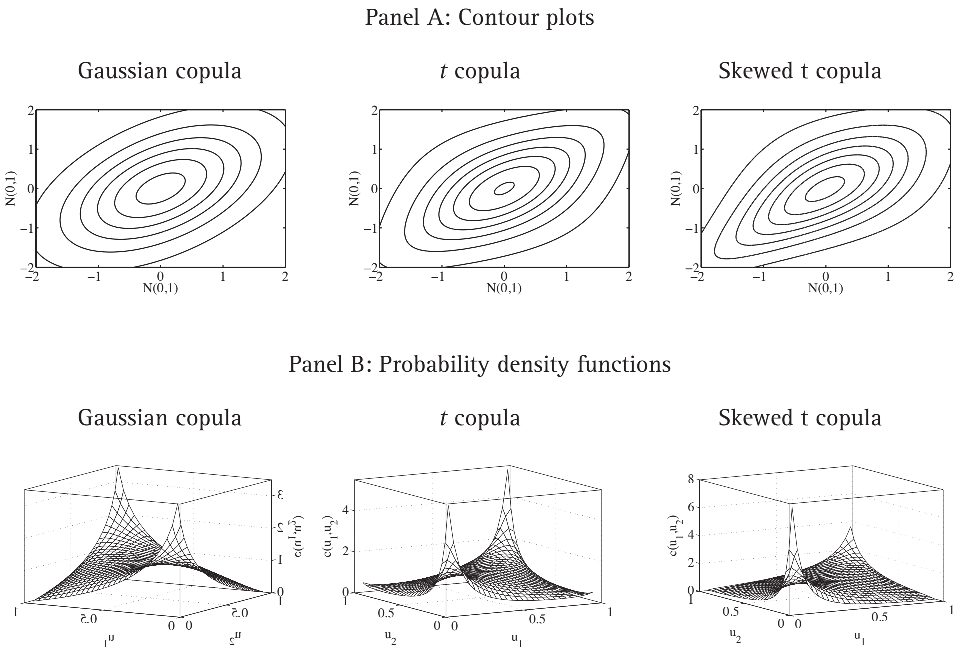

Panel A shows the contour plots of the distribution for three copulas. To compare just the copula function, all of them are evaluated using standard normal marginal distributions, N(0,1). Panel B shows a bivariate representation of the probability density function for the three copulas.

Table 2.1. Descriptive Statistics for oil, gold, and equity weekly returns

	Full sample: 1990–2010			In-sample: 1990–2006			Out-of-sample: 2006-2010		
Assets	oil	gold	equity	oil	gold	equity	oil	gold	equity
Mean	0.136	0.120	0.104	0.163	0.061	0.147	0.033	0.343	-0.059
	(0.341)	(0.082)	(0.145)	(0.285)	(0.353)	(0.045)	(0.929)	(0.114)	(0.766)
Std. Dev.	4.634	2.236	2.326	4.401	1.896	2.129	5.441	3.210	2.958
Min.	-36.53	-13.21	-16.45	-36.53	-11.04	-9.04	-16.63	-13.21	-16.45
Max.	23.98	12.88	10.18	14.55	12.88	10.18	23.98	10.92	9.639
Sharpe	0.029	0.054	0.045	0.037	0.032	0.069	0.006	0.107	-0.020
VaR 5%	6.892	3.294	3.744	6.494	2.737	3.483	7.717	4.753	4.931
Skewness	-0.598	0.007	-0.552	-0.929	0.102	-0.134	0.114	-0.205	-1.024
	(0.000)	(0.927)	(0.000)	(0.000)	(0.229)	(0.115)	(0.489)	(0.215)	(0.000)
Kurtosis	8.259	7.343	7.238	9.860	8.258	5.000	4.914	4.536	8.012
	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
JB	1280	829.8	843.9	1759	964.3	141.9	34.06	23.18	268.7
	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.002)	(0.000)
KS	0.446	0.467	0.469	0.450	0.473	0.470	0.442	0.459	0.467
	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
LB(10)	27.93	25.97	29.98	14.27	29.12	24.05	35.99	17.08	22.55
	(0.002)	(0.004)	(0.001)	(0.161)	(0.001)	(0.007)	(0.000)	(0.073)	(0.013)
LM(10)	90.38	189.6	138.7	54.76	72.11	84.29	39.40	64.06	31.52
	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)

2.3. Empirical Results

Our empirical application relies on three risky assets: two commodity futures, oil and gold, and the S&P 500 equity index. The oil futures correspond to West Texas Intermediate (WTI) crude oil from the New York Mercantile Exchange (NYMEX). The gold futures correspond to the gold bar, with a minimum of 0.995 fineness, from the New York Commodities Exchange (COMEX). These futures are two of the most actively traded commodity contracts in the world, and they do not have tight restrictions on the size of daily price movements. In both cases, we employ the most liquid futures contracts, measured by daily trading volume, of all maturities available. The sample period considered ranges from June 20, 1990 to September 8, 2010, for a total of 1056 weekly observations. We divided the sample in two sub-periods, such that the period from June 20, 1990 to June 20, 2006 (836 observations) supported the in-sample estimation analyses of the models, and the remaining 220 observations from June 20, 2006 to September 8, 2010 were reserved for the out-of-sample portfolio performance exercise.

According to the Jarque-Bera and Kolmogorov-Smirnov tests, normality in the returns' unconditional distribution is strongly rejected for all samples (see Table 2.1). Besides, skewness and kurtosis of returns differs across assets and sample periods. There is also evidence of serial correlation in the returns and squared returns for all time series.

To check for the presence of asymmetric dependence between asset returns in our sample, we analyzed the exceedance correlation and tail dependence. For each pair of asset returns, Figure 2.2 plots the exceedance correlation function proposed in Longin and Solnik (2001), which depicts the correlation between returns above or below a given quantile. In the case of symmetric dependence, the correlation for both extremes

Figure 2.2. Exceedance correlation

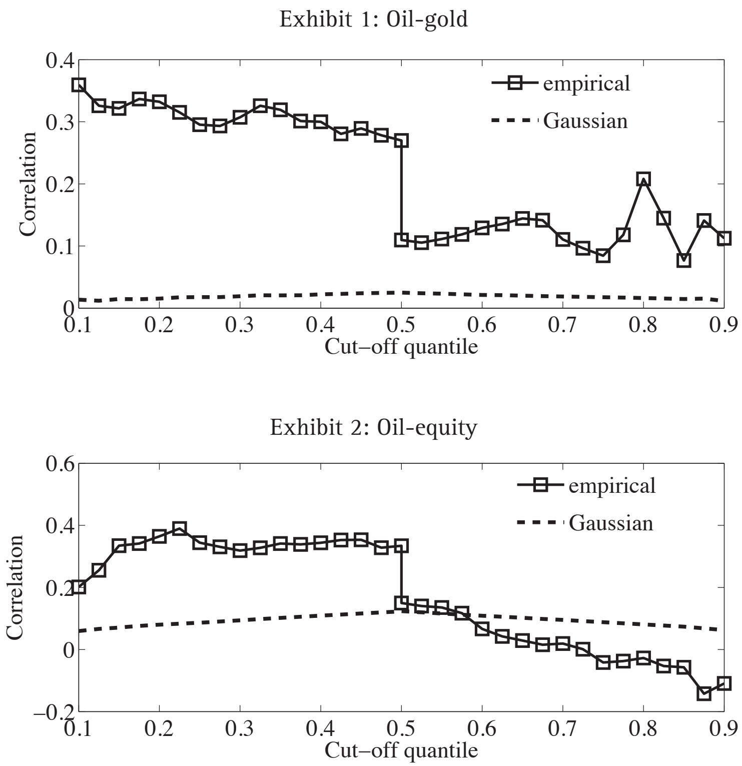

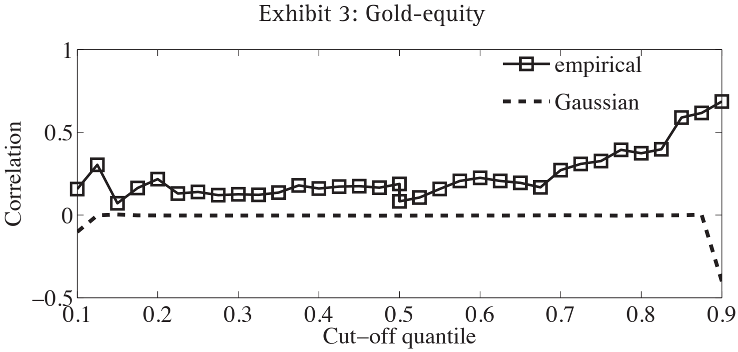

should be similar and equal to zero for Gaussian dependence. According to these plots, any assumptions of normality or symmetry seem unrealistic for our sample. Oil and gold do not display the same level of diversification for bear and bull markets, and correlation between oil and equity is highly positive for large negative returns but smaller for large positive returns. The correlation between gold and equity is close to 0 for large negative returns and significantly positive for very large positive returns. Although oil and gold are very positively correlated for large negative returns, are not or even are negatively correlated for large positive returns.

To obtain the optimal portfolio decisions based on our copula models over the out-of-sample period, we need the forecasts of the different parameters at play over the 2006-2010 period. For that purpose, we recursively re-estimate the marginal and copula models throughout the out-of-sample period (220 weekly observations) using a rolling window scheme that drops distant observations as more recent ones are added and therefore keeps the size of the estimation window fixed at 836 observations. Once we re-estimate the model for each point in the out-of-sample period, we construct the time-series of one-period-ahead parameter forecasts needed for the allocation stage.

Figure 2.3 shows the output of the forecasts of the conditional mean, volatility, and skewness of each return process throughout the out-of-sample period. The volatility forecasts of all asset returns are relatively high, especially around October 2008. Conditional skewness is negative for equity and oil returns during the 2006-2010 period, but it is positive for gold returns during that period.

Figure 2.4 presents the forecasts of the conditional dependence parameters. It is worth noting that there is an increase in the fitted correlation coefficients among oil, gold, and

equity from October 2008, especially for oil and equity returns (see Exhibit 1). In addition, the dependence coefficients seem to evolve more similarly in the latter part of the sample. The degrees-of-freedom forecasts decrease after August 2007, indicating rising tail dependence since then (see Exhibit 2). In addition, the asymmetry parameter of oil ranges between -0.6 and -0.2, which implies that extreme dependence seems to be stronger during large depreciations of oil, compared with large drops in gold or equity, whose asymmetry parameters range between -0.2 and +0.2 (see the forecast of the asymmetry parameter vector in Exhibit 3).

In summary, the skewed *t* copula provides a more informative measure of the dependence between commodities and equity-index returns, even taking into account that part of the tail behavior is captured by the skewed fat-tailed marginal distribution models. Therefore, possibly univariate tail behavior and asymmetric dependence are key factors not taken into account in a standard elliptical, *à la Markowitz*, approach. The extent to which these factors have a significant impact on the portfolio choice decision is addressed in the next section.

We now investigate the optimal portfolio decisions based on six model-driven portfolio strategies that can be analyzed from the perspective of copula models.

First, we consider the unconditional multivariate Gaussian model (Markowitz strategy), a constant Gaussian copula with unconditional Gaussian marginal distributions. Second, we generalize this case by considering two conditional multivariate Gaussian distributions: the constant conditional correlation (CCC) and the dynamic conditional correlation (DCC). Both CCC and DCC specifications are formed by conditional Gaussian marginal distributions with conditional means and variances. Third, we compute portfolio strategies using the conditional copula models introduced previously. Thus, we consider the generalized Student's *t* distribution for the marginal models and three types of conditional dependence functions: Gaussian, *t*, and skewed *t*.

Figure 2.3 Conditional parameters of the marginal distribution model

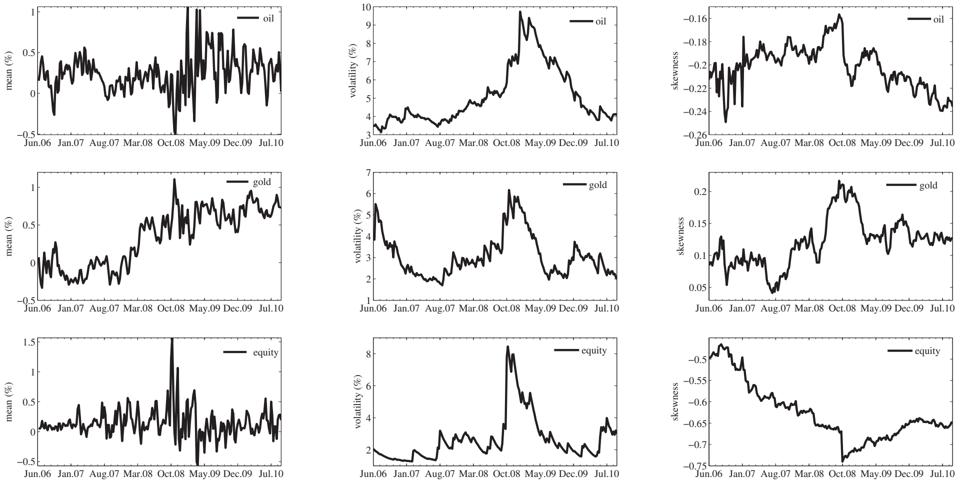

Figure 2.4 Conditional parameters of the conditional skewed *t* copula

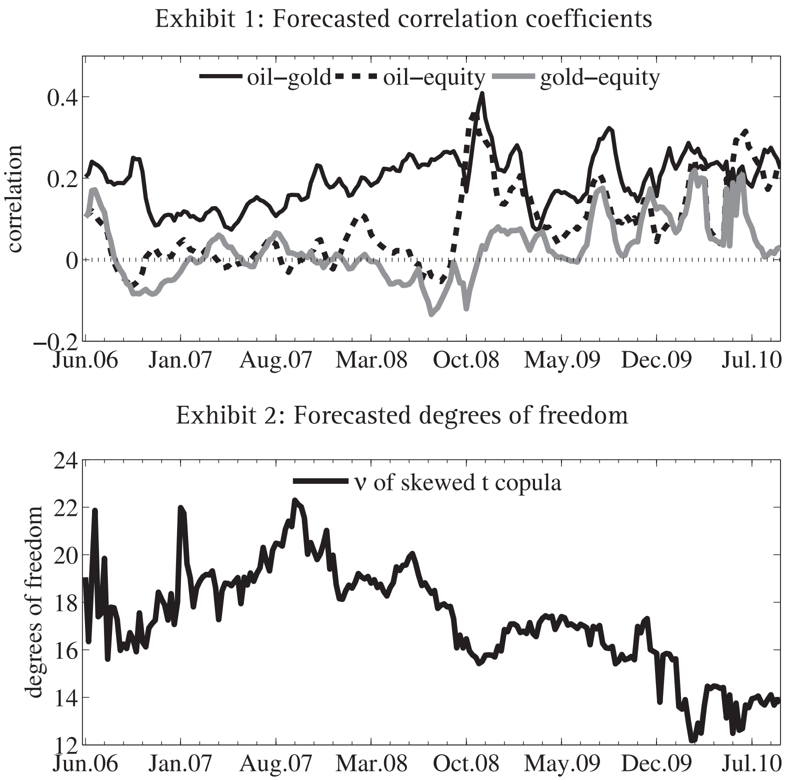

Exhibit 3: Forecasted asymmetry vector

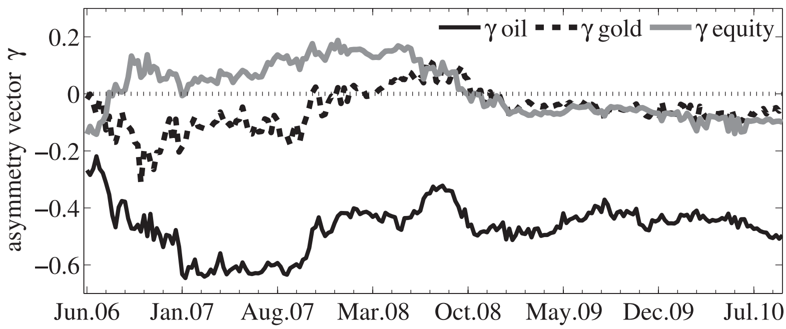

With this set of alternatives, we can compare the gains of including more flexible models as a means to compute portfolio decisions. In addition, we include in the analysis the equally weighted portfolio, as a common benchmark used in prior literature. Moreover, we analyze the portfolio allocations for different parameterizations of the investor's three-moment preferences, defined by of φ_V and φ_V.

In Panel A of Figure 2.5 we plot, for two of these preferences specifications, the time-series of portfolio weights resulting from the portfolio decisions made using our most general model, the conditional skewed t copula. Panel B of Figure 2.5 shows the allocation differences between the unconditional Gaussian model and the conditional skewed t model for the same risk aversion parameterization.

Figure 2.5 Optimal portfolio weights

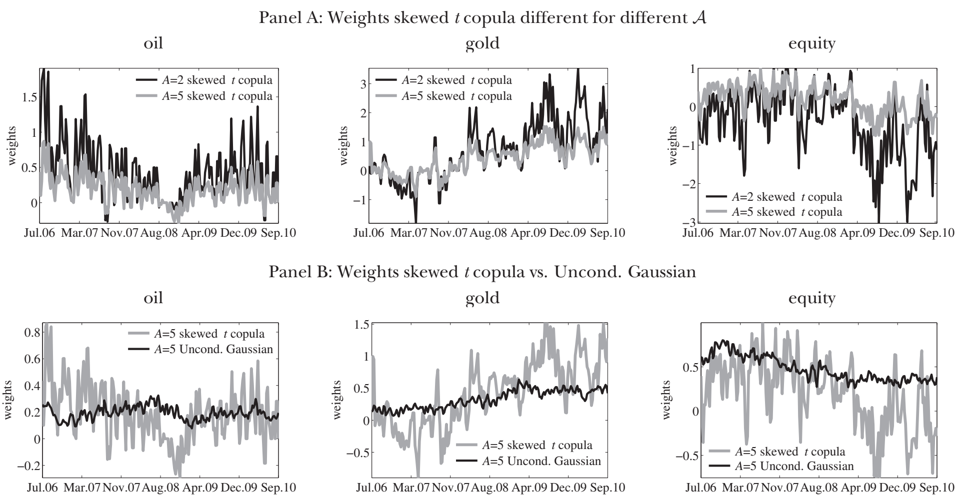

The results show that the bulk of the difference between portfolios strategies depends largely on the use of different marginal distribution models. The first significant discrepancies arise when using time-varying Gaussian marginal distributions (CCC and DCC models) instead of unconditional Gaussian margins (Uncond. Gaussian model).

A second source of allocation differences is driven by the various types of dependence captured with our copula models. These discrepancies in optimal portfolio weights arise, first, from introducing a time-varying conditional dependence (e.g., CCC vs. DCC); and second, from considering tail dependence (e.g., *t* copula vs. Gaussian copula) and asymmetric dependence (e.g., skewed *t* vs. *t* copula).

2.4. Conclusions

This chapter investigates the portfolio selection problem of an investor with time-varying three-moment preferences when commodity futures are part of the investment opportunity set. In our specification, the portfolio returns' skewness provides a measure of the investor's loss aversion. We model the joint distribution of asset returns using a flexible multivariate copula setting that can disentangle the specific properties of each asset process from its dependence structure. The more general model we posit consists of a conditional skewed *t* copula with generalized Student's *t* marginal distributions and time-varying moments. Thus we can capture the specific distributional characteristics of commodity-futures returns and focus on their implications for the portfolio selection problem.

The empirical application employs weekly data for oil and gold futures and for the S&P 500 equity index, from June 1990 to September 2010. We find substantial discrepancies between the holdings obtained from our conditional copula models and those from more traditional Gaussian models.

The key factors underlying these differences are the different specifications of the time-varying marginal distributions, the presence of dynamic conditional dependence among the univariate processes, and the modeling of tail and asymmetric dependence. The univariate higher moments and the type of tail dependence are more relevant for aggressive investors. These discrepancies translate into economical differences in terms of better investment ratios and relative performance measures for the different specifications considered.

3. TAIL RISK IN ENERGY PORTFOLIOS

The growth of energy markets has been sustained by continued deregulation processes, which have encouraged the separation of the formerly integrated value chains. This process has increased market risk exposures at every stage of the chain, including the purchase and sale of fuels, electricity generation, and obtaining gas or electricity for retail supply. In addition to the physical resource holders, financial players, such as banks and hedge funds, increasingly participate in energy markets to satisfy their customers' demands to gain or hedge energy risk exposure, as well as to trade on their own behalf. In this context, energy-related companies and financial players experience greater exposures to energy price risk, which has particular characteristics that make it different from other market risks and requires clearer explication.

In this chapter, we therefore analyze the energy price risks from a multivariate perspective. In particular, we study the aggregate tail risk of different linear energy portfolios using an asset-level approach. Accordingly, we can propose a multivariate model for the vector of energy risk factors; using the portfolio exposures to each factor, we in turn can calculate the aggregate tail behavior of the portfolio. Next, we compute the corresponding portfolio risk measures and evaluate the extent to which the tail pattern of the model is important in practice.

With this asset-level approach, we can capture the entire structure of energy risk factors in a portfolio and their interdependence relationships. This multivariate behavior (univariate and joint structure) of energy risk factors depends on the special characteristics of energy markets. In particular, the pricing of energy commodities relies largely on an equilibrium among supply, demand, and inventories, subject to various operational constraints (for example, due to infeasible or overly costly storage). These characteristics cause deregulated energy markets to exhibit substantial volatility, price spikes, time-varying correlation, dependence in the extremes, and mean-reversion patterns (e.g., Cartea and Figueroa (2005), Escribano, Peña, and Villaplana (2011), Pirrong (2012), and Routledge, Seppi, and Spatt (2001)).

We therefore employ a multivariate density model to depict the energy risk factors, in which we seek to include all the stylized features of the data generating process. For this purpose, we consider an econometric specification with time-varying conditional means, volatilities, and correlations, in which the innovation vector follows a multivariate generalized hyperbolic (GH) distribution. The GH class is a very flexible family of distributions that accommodates excess kurtosis, skewness, and dependence in the extremes (see Börger, Cartea, Kiesel, and Schindlmayr (2009), Eberlein and Stahl (2003), and Giot and Laurent (2003), for previous theoretical and empirical studies that employ some distributions within this class).

We apply our multivariate GH specification to model the returns vector formed by the four most important commodities in the U.S. energy market: crude oil, natural gas, coal, and electricity. These commodities constitute the elements of our linear energy portfolios, which represent the exposure of any given energy company or financial player to energy price risk. We use daily data from August 2005 to March 2012 to estimate the multivariate models and evaluate the tail risk of the portfolio profit-and-loss (P&L) distribution. Then using data from March 2010 to March 2012, we conduct out-of-sample forecast evaluations of the risk measures.

We address the analysis of the aggregate tail risk by calculating two risk measures, the value at risk (VaR) and the expected shortfall (ES), for long and short trading positions in the energy portfolios. The VaR corresponds to the quantile of the portfolio loss distribution for a given probability or confidence level. The ES is defined as the conditional average loss beyond a given quantile, and it better describes the behavior of the portfolio losses in the tail. We estimate both measures for different confidence levels, which define how far out in the tails the risk measures are calculated, as well as for several day horizons, to obtain a short-term surface of risk. Whereas most equity risk studies have focused on the left tail of long positions, the presence of positive jumps in the data generating process of energy commodities, especially for the natural gas and electricity markets, suggests that the analysis of the right tail of the possibly asymmetric P&L distribution could be relevant for traders who are worried about increases in energy prices (i.e., those with short positions).

Finally, using different backtest procedures, we monitor, for the out-of-sample period, the performance of the risk measure estimates that correspond to the GH models. We pay special attention to the backtesting of the ES estimates, because this measure offers more information about aggregate tail behavior.

Empirical results show that there are more VaR violations across models for short positions than for long ones, confirming the positive asymmetry of the P&L distribution of the energy portfolios. We also observe that the ES exceedances are quite high for (conditional and unconditional) Gaussian models, especially for the two utility portfolios. The heavy-tail models behave much better than alternative versions, with regard to the tail risk of short positions. Therefore, the extent of the underestimations of the tail risk of the portfolio loss distribution depends on whether we are analyzing short or long positions in the energy portfolio, the type of portfolio, the horizon, and how far out in the tail the risk is being analyzed.

3.1. Energy portfolios and returns

We approximate a given exposure to energy price risk using a corresponding portfolio of energy futures. Thus, changes in the energy price risk factors can be mapped linearly to changes in the value of the energy futures portfolio. A portfolio of futures contracts can also be considered a first-order approximation of more general energy asset portfolios with non-linear payoffs (Tseng and Barz, 2002; Cartea and González-Pedraz, 2012). For example, a linear portfolio could represent directly the energy futures positions of an institutional investor or the energy price exposure of an electricity producer with fuel-fired power plants. In this chapter, we consider four energy commodities: crude oil, natural gas, coal, and electricity, identified by subscripts i equal to 1, 2, 3, and 4, respectively. These four commodities substantially represent any general exposure to energy price risk.

The h-period return (in dollars) on an energy portfolio at time t is given by

$$\Delta W_t(h) = W_t - W_{t-h} = R_t(h)W_{t-h} = w'_{t-h}(\exp(r_t(h)) - \imath_4) \qquad (3.1)$$

where $R_t(h)$ is the h-period net return on the portfolio, $r_t(h) = \sum_{k=0}^{h-1} r_{t-k}$ is the 4x1 vector of h-period log-returns at time t. The energy log-returns $r_{i,t}$ constitute the vector of risk factors.

According to the previous expression, the portfolio's profit-and-loss (P&L) distribution is determined by the multivariate density model of energy risk factors r_t and the positions in the energy commodities, w_t. The multivariate model of risk factors describes the joint behavior of the four commodities. The different positions in energy futures define the size and sign of the exposure to each energy commodity, mapping the multivariate model onto a specific P&L distribution. In this chapter, we consider four portfolios: two related to power utilities and two others more related to financial players. As representative portfolios of electricity producers, we include the energy portfolio of a utility with a diversified mix of generation that operates in the Pennsylvania-Jersey-Maryland (PJM) Interconnection and a linear portfolio corresponding to a gas-fired power plant. In addition, we account for two typical portfolios of financial players that seek exposure to energy commodities: an equally weighted portfolio and the minimum variance portfolio.

Using r_t as the vector of the four energy assets log-returns at time t, we assume that its data generating process is given by

$$r_t = m_t + H_t^{1/2} x_t \qquad (3.2)$$

where m_t is the vector of conditional means; $\mathrm{H}_t^{1/2}$ is the 4x4 Cholesky factor of the time-varying covariance matrix H_t such that $H_t = H_t^{1/2}(H_t^{1/2})'$; and x_t is the vector of independent innovations, which follows a four-variate generalized hyperbolic (GH) distribution with zero mean and an identity covariance matrix.

To capture the possible presence of serial correlation in energy returns, we consider a diagonal vector autoregressive (VAR) process with up to 5 lags for the vector of returns. We want to capture possible persistence and asymmetry in conditional variances and correlations. For that purpose, we assume univariate asymmetric GARCH(1,1) processes for the conditional variances and a modified version of the asymmetric dynamic conditional correlation (ADCC) model of Cappiello, Engle, and Sheppard (2006) for the time-varying correlation matrix. With this specification, we investigate, in the conditional correlation, the presence of asymmetric responses to positive shocks.

Motivated by the presence of jumps and spikes in energy prices, we employ multivariate GH distributions to model the conditional distribution of the vector of innovations x_t. These GH distributions are flexible enough to accommodate different tail behaviors and types of asymmetry (e.g., thin or heavy tails, symmetric or positive/negative skewness). The GH family can be obtained using the following normal mean-variance mixture representation (see McNeil, Frey, and Embrechts, 2005):

$$x_t \overset{\mathrm{d}ist.}{=} \mu + \omega_t \gamma + \omega_t^{1/2} A z_t\,, \quad z_t \sim \mathrm{N}_4(0, I_4) \tag{3.3}$$

where μ and γ are the 4x1 location and skewness parameter vectors, respectively, and $\Sigma = AA'$ is the 4x4 dispersion matrix. The random vector z_t follows a four-variate Gaussian distribution with zero mean and identity covariance, and ω_t is a non-negative random variable independent of z_t. The mixing random variable ω_t can be understood as a shock that affects the covariance of energy assets, due to the arrival of new information in the markets (e.g., shortages in future supply, unexpected increases in demand). Conditioned on ω_t the vector of innovations is normally distributed.

In the case of GH distributions, the mixing random variable ω_t follows a generalized inverse Gaussian (GIG) distribution. The very flexible GIG distribution includes as special boundary cases the gamma and inverse gamma distributions. We consider five particular cases of multivariate GH distributions: the normal inverse Gaussian (NIG) distribution, the variance-gamma (VG) distribution, the skewed t (skT) distribution, the Student's t (T) distribution, and the Gaussian (G) distribution (for which $\gamma = 0$ and $\omega_t = 1$).

The proposed GH distributions have the advantages of exhibiting different tail patterns (Bibby and Sørensen (2003)). On the one hand, the tails of the NIG and VG distributions decay exponentially, such that their probability density functions behave, when $x_t \to \pm\infty$, proportionally to an exponential function. This pattern is intermediate between the behavior of the Gaussian distribution, which decays more rapidly, and other, more extreme, polynomial decays. For this reason, NIG and VG distributions are sometimes referred to as semi-heavy tailed. The tails of the T distribution instead are symmetric and behave as polynomials, such that they decay slower than those of the NIG and VG distributions. Finally, the skT distribution offers the special property of possessing, for each component of the vector of innovations, one heavy (polynomial decay) and one semi-heavy (exponential decay) tail. Thus, when $\gamma_i > 0$, the right tail $(x_{i,t} \to \infty)$ is the heaviest, while the left tail $(x_{i,t} \to -\infty)$ decays exponentially; these roles switch for $\gamma_i < 0$.

3.2. Risk measures and numerical implementation

Measuring conditional risk is a natural and direct way to analyze the tail behavior of energy portfolios. From the point of view of risk managers, reliable estimates of risk measures depend on a proper understanding of the portfolio's tail behavior. In our approach, we study both long and short positions in the energy portfolios. Thus, we focus on the two tails of the P&L distribution. For short positions, the portfolio holder loses money when the portfolio value increases, so we attend to the right side of the distribution. For long positions, we focus on the left tail.

We consider two measures of risk: the value at risk (VaR) and the expected shortfall (ES). The VaR is widely used in the financial industry to monitor risk exposures for regulatory purposes and to establish trading constraints in investment decisions. In the energy industry, especially for producers, VaR is becoming more popular, with increasing relevance for corporate decisions. For example, VaR provides insights to determine hedging policies or, in the case of utilities, to obtain an optimal selection in the generation mix. Formally, for a certain horizon h and confidence level α, the VaR is defined as the α-quantile of the conditional distribution of portfolio changes $\Delta W_t(h)$. That is, the probability of incurring losses greater than a certain threshold value, called the VaR, is equal to α:

$$P(\,\Delta W_t(h) \leq \mathrm{VaR}_t(\alpha, h)\,|\,\mathcal{F}_{t-1}\,) = \alpha \tag{3.4}$$

Despite its widespread use, the VaR also has been subject to substantial criticism, particularly because diversification does not always reduce risk when it is measured by VaR. In addition, the VaR ignores important information related to the tails of the loss distribution beyond the α-quantile, disregarding the risk of extreme losses. In contrast, ES measures cope well with such shortcomings and describe tail risk better (Artzner, Delbaen, Eber, and Heath (1999)). The ES is defined as the expected loss, conditional on the loss exceeding the VaR over a certain horizon h,

$$\mathrm{ES}_t(\alpha, h) = E[\,(\Delta W_t(h)\,|\,\Delta W_t(h) \leq \mathrm{VaR}_t(\alpha, h))\,|\,\mathcal{F}_{t-1}\,] \qquad (3.5)$$

that is, the average portfolio loss in the of α% worst cases.

We employ an asset-level approach to measure the tail risk of the energy portfolios. We begin by modeling the joint distribution of energy returns under the dynamic econometric models proposed in the previous section. Then we aggregate these results for each portfolio according to its exposures to each commodity. To aggregate the risk factors, it is convenient to represent the portfolio's P&L as a linear function of the individual energy log-returns. Because the GH distributions are closed under linear transformations, when we aggregate the energy risk factors in a given portfolio, the linearized P&L distribution still belongs to the same class of GH distributions as does the vector of risk factors. Finally, we can adopt two alternative numerical implementations for calculating the risk measures. We can compute VaR and ES under the GH model by solving the integrals implicit in equations (3.4) and (3.5) numerically for the portfolio P&L distribution. Alternatively, we can apply Monte Carlo simulations, which are usually more effective and preferred in this context.

3.3. Model estimation

Our energy portfolios consist of energy commodity futures for crude oil, natural gas, coal, and electricity. These four commodities effectively represent a wide range of exposures to energy price risk. In all cases, we employ daily series of one-month ahead monthly futures contracts traded on the New York Mercantile Exchange (NYMEX), which are the most liquid contracts for the four energy commodities analyzed.

The full-sample period runs more than six years from August 2005 to March 2012, and includes 1,640 daily observations. We consider all data since the launch of the PJM electricity futures in the NYMEX (April 2003) until the day of the analysis (March 2012), but we drop the first observations (from April 2003 to August 2005), for which liquidity of electricity and coal futures was very scarce. To avoid in-sample over-fitting and spurious findings, we reserve the last two years of data, from March 2010 to March 2012 (504 observations), for the out-of-sample investigation of the tail risk.

Figure 3.1 Relative Prices and QQ-plots

Exhibit 1: Relative prices

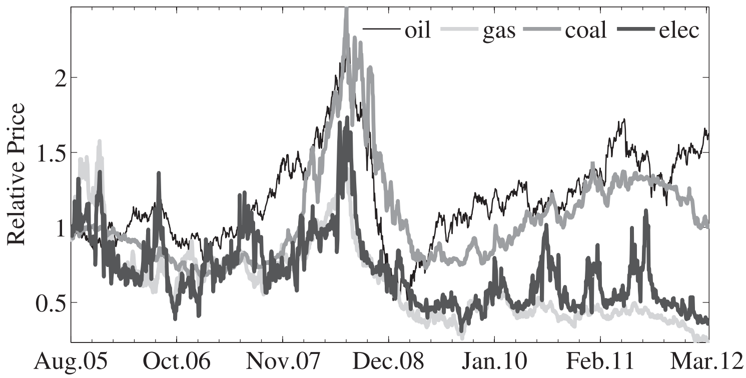

Exhibit 2: Quantile-quantile plots

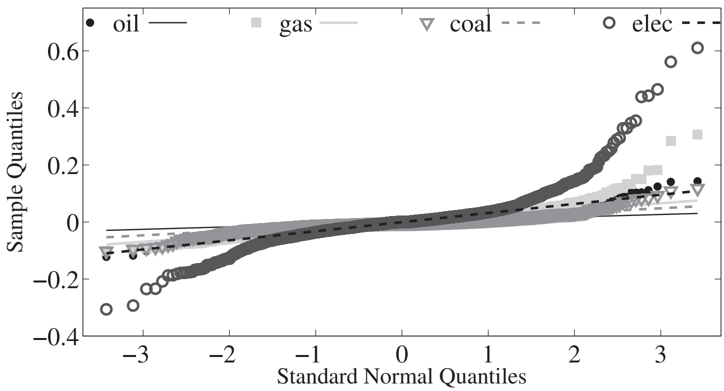

Figure 3.1 presents the relative prices and quantile-quantile plots for the four energy commodities from August 2005 to March 2012. Electricity has the highest volatility, kurtosis, and risk measures over the entire sample. It also shows extreme positive and negative daily moves, some larger than 30%. We observe non-negligible skewness across commodity returns. Electricity and natural gas returns exhibit significantly positive skewness for all periods, suggesting that positive moves are more frequent than negative ones in these markets.

The correlation coefficients between fuels (oil, gas, and coal), which to some extent represent substitute goods, range from 19% to 32% over the 2005-2010 period. In the 2010-2012 period, the correlations of oil and gas with coal increase to greater than 31%, whereas the correlation between oil and gas decreases from 28% to 12%. The linear dependence between electricity and fuels is less than 10% and only significant for oil and natural gas during 2005-2010. In the last period, correlation with coal increases to 10%.

We also conduct Mardia's test of multivariate normality (not reported here). This test is based on multivariate measures of skewness and kurtosis. The large values that we obtain for the test statistics, corresponding to multivariate skewness and kurtosis measures, reject the null hypothesis of joint normality of energy returns.

The estimation of the multivariate GH models for the energy returns is carried out in two stages. The large dimension of the model prompts us to use this sequential approach to estimate the set of parameters. In the first stage, we carry out the quasi-maximum likelihood (QML) estimation of the dynamic regression model for the conditional mean and covariance. In the second stage, we obtain the ML parameter estimates of the different multivariate GH conditional distributions.

We observe different patterns in the variance equation, especially with respect to the leverage effect. For crude oil and natural gas, the parameter corresponding to the leverage effect, is positive and significant, which suggests that negative shocks have a stronger effect on variance than do positive ones. Coal and electricity do not indicate any such asymmetry in terms of the response of volatility to negative moves, which suggests that positive shocks could have more impact on variance. Volatility persistence also is very large (>0.95) for fossil fuels but smaller for electricity variance (around 0.50).

When we consider the time-varying evolution of the correlation matrix for the vector of four energy returns, we find that dependence dynamics are strongly persistent. When we consider the dynamics of the correlation between pairs of energy returns, we also obtain significant, positive estimates of the asymmetry parameter; in particular, for natural gas and electricity. Such positive asymmetry in correlation seems sound from an economic perspective, because an increase in gas prices has a strong positive impact on the generation costs of peak-load electricity. We study this relationship between gas and electricity in depth when we analyze the tail risk of the gas-fired power plant portfolio.

The results of the fit of the conditional distributions offer strong evidence against multivariate normality, as we expected. First, the shape parameter estimates of the mixing distributions point to the presence of fat tails in the different GH models. In particular, for the T and skT distributions, the small value of the degree-of-freedom parameter indicates the existence of jumps and tail dependence. Similar arguments apply to the VG and NIG parameter estimates. Second, the asymmetry parameter estimates γ for the three skewed GH distributions (skT, VG, and NIG) are positive for all vector components, suggesting positive skewness in the multivariate conditional distribution of daily energy returns. We reach similar results when we re-estimate the GH conditional distributions throughout the out-of-sample period.

To analyze the tails of the returns distribution of energy portfolios, we consider the following examples: a utility with different generation units, a gas-fired power plant, and equally weighted and minimum variance portfolios. Using the multivariate GH models previously estimated and knowing the portfolio weights, we can obtain a fitted distribution of portfolio returns for each GH model. Then, we compare the in-sample tail

Figure 3.2 Tail plots of energy portfolios

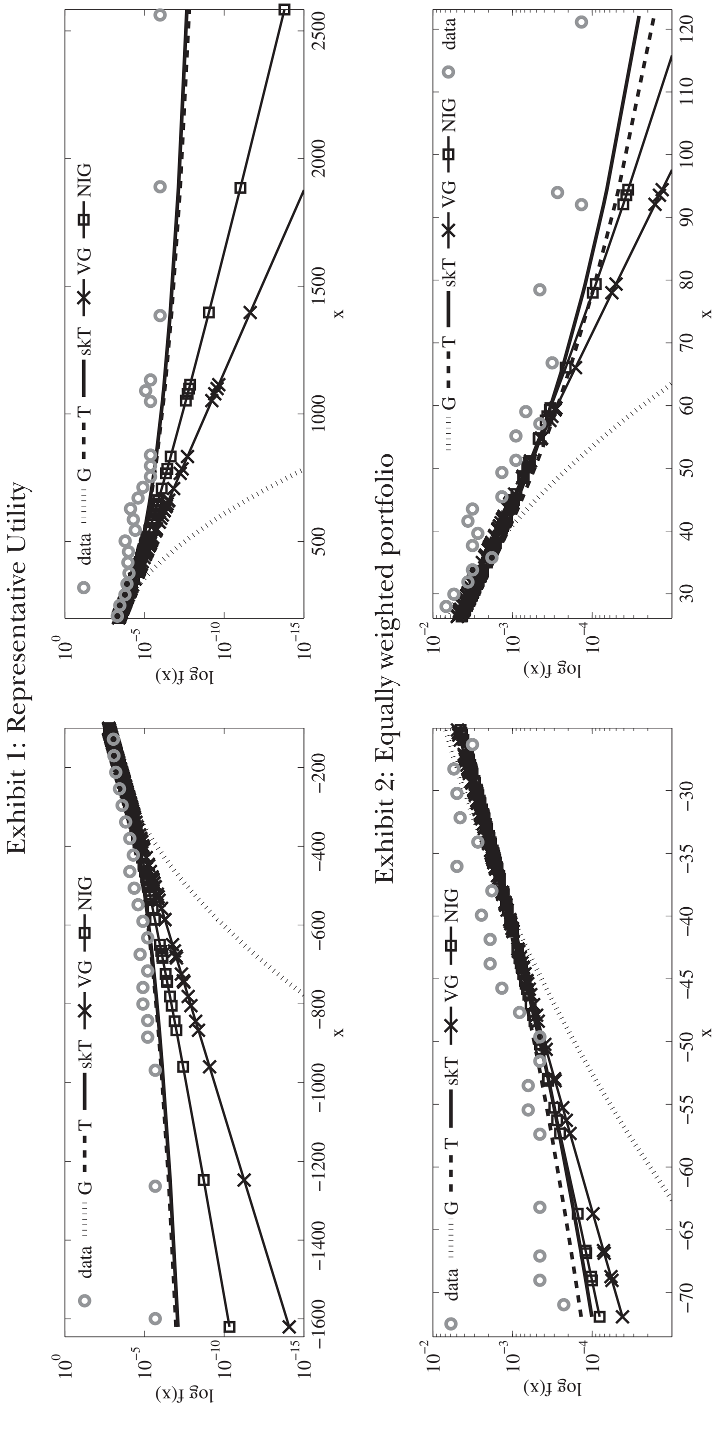

fit of the estimated models graphically, by plotting the estimated logarithmic density functions and the empirical log-density function of the portfolio.

To focus on the aggregate tail risk behavior, we display in Figure 3.2 enlarged sections of the left and right tails of the energy portfolios. The circles represent the empirical probability density of portfolio return innovations. The left panel of each exhibit presents the left tail of a long position in the portfolio, and the right panel is the corresponding right tail. As expected from previous multivariate results, the distributions of portfolio returns show positive skewness and fat tails. We find that the Gaussian model (dotted line) clearly underestimates the extent of both tails, that is, the probability of extreme realizations.

The T and skT models better estimate the aggregate tail risk behavior, according to the plots in Figure 3.2. The slower tail decay of T and skT distributions (solid and dashed lines, respectively) causes them to outperform the tail fit of the VG and NIG models (marked with crosses and squares, respectively), especially for the right tail, which corresponds to losses of a short position in the energy portfolio. We also observe slight differences between the tail fits of the T and skT distributions, partially due to the asymmetric tail behavior of the skT distribution. Similar exhibits for equally weighted and minimum variance portfolios show that the left tail of the estimated skT distribution is above the T distribution, whereas the right tail is below it.

In the next section, we further consider the aggregate tail behavior of the energy portfolios' loss distribution, looking at the out-of-sample performance of the VaR and ES measures.

3.4. Out-of-sample performance of risk measures

In the final part of the chapter, we compute VaR and ES over the out-of-sample period for different horizons and confidence levels, characterizing the term structure of these risk measures for each GH model. Then, we test the out-of-sample performance of the forecast risk measures, assessing the relative ability of the various multivariate models at hand.

In addition to our GH models, we also calculate the portfolio risk measures using several approaches: a traditional variance-covariance method with multivariate unconditional Gaussian distribution (VC); the Riskmetrics procedure or exponentially weighted moving average model (EWMA), as first introduced by J.P. Morgan; the multivariate Gaussian GARCH with constant conditional correlation (CCC); and the non-parametric historical simulation method (HS).

Using the various multivariate approaches, we calculate the conditional risk measures (VaR and ES) of the four energy portfolios for horizons extending from 1 to 22 days. By way of a sensitivity exercise to the cut-off point selection, we also consider in our analyses different confidence levels $\alpha\alpha$: 0.1%, 0.5%, 1%, and 5%. Thus, we can analyze the possible bias in the risk measure estimates due to the fixing of the confidence level.

Figure 3.3 1-day VaR

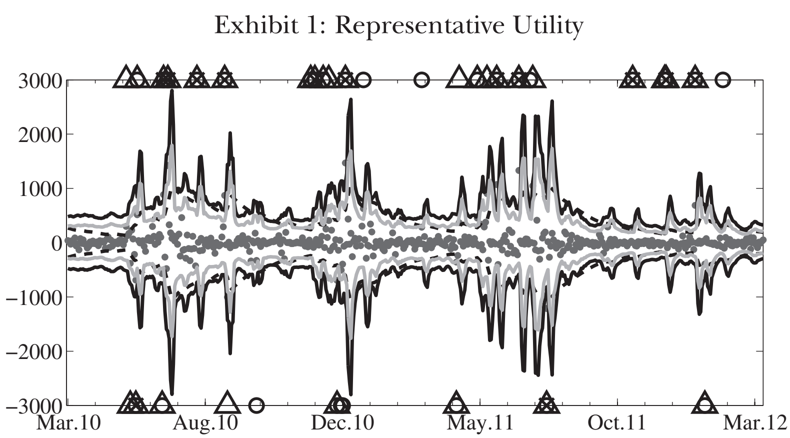

Exhibit 2: Equally weighted

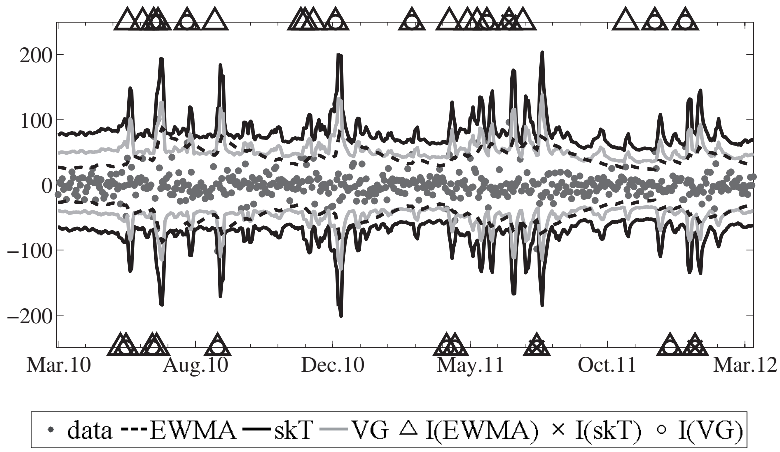

The smallest ES estimates are generally produced by the EWMA or Gaussian models (CCC and VC are not reported here, in the interest of clarity). The largest ES estimates among the GH models correspond to the distributions with polynomial (slower) tail decays, that is, to the T and skT distributions. The asymmetric pattern of the skT model produces slightly larger tail risk estimates than the T model for the long positions of the energy portfolios, especially for the equally weighted and minimum variance portfolios.

The ordering of the ES estimates across GH methods is invariant to the forecast horizon.

In general, the tail risk estimates of the nonparametric HS are close to those of the T and skT models. The estimation windows characterized by turbulent periods of fuel returns are responsible for these large risk measure estimates of the HS approach; as we observe, only heavy-tailed distributions are able to produce similar tail risk patterns.

In practice, our interest lies in comparing (backtest) the h-horizon risk measures forecasts for long and short positions with the actual portfolio losses during the two-year out-of-sample period, from March 2010 to March 2012. Thus we can assess the differences in tail risk patterns, controlling for over-fitting and other spurious findings.

Figure 3.3 shows 1% 1-day VaR over the out-of-sample period according to three different approaches: the EWMA, VG, and skT models. In the lower side of the figure, we draw the VaR violations of each model, corresponding to the long position in the portfolio. In the upper side, we mark the VaR violations for the short position. In general, there are more violations across models for short positions than for long ones, in support of the positive asymmetry of the P&L distribution of the energy portfolios.

We also observe that the VG and skT estimates (grey and black lines, respectively) respond more quickly to changing volatility than does the EWMA estimate (dashed line), which tends to be violated several times in a row during more turbulent periods (violations of the EWMA risk measure are marked with triangles). In addition, the VaR violations of the skT estimate (crosses) are fewer than those of the VG estimate (circles), suggesting again the importance of modeling the presence of heavy tails to produce conservative tail risk measures.

Using the number of VaR violations for a given confidence level over the tested period, as well as the proportion of losses beyond that VaR estimate, we can build a series of backtests to monitor the out-of-sample performance of tail risk estimates.

The results show that the more traditional parametric approaches, such as VC, EWMA, CCC, and Gaussian-DCC models, tend to underestimate the VaR, especially for short positions and for utility portfolios. The non-parametric HS produces better coverage probabilities for these cases.

Comparing all models jointly to determine whether the differences in the tail patterns are statistically significant, we reject the claim that the (conditional and unconditional) Gaussian models, such as VC, EWMA, CCC, and G, perform as well as the best competing alternative model, with the possible exception of the long portfolio positions at 10-day horizons. These tests support our previous findings, namely that models with exponential tail decay (i.e., VG and NIG models) yield inferior tail estimates for short portfolio positions, especially for the far tail (alpha=1%) of utility portfolios at the 1-day horizon.

3.5. Conclusions

In this chapter, we have characterized the tail behavior of energy price risk using a dynamic multivariate model. We approximated exposure to energy price risk for physical and financial players using linear combinations (portfolios) of crude oil, natural gas, coal, and electricity futures.

To model the stylized features of the vector of energy risk factors, we have proposed a flexible econometric specification. With respect to the conditional distribution, we considered the possibility that the vector of innovations may be generated by a multivariate GH distribution. With these distributions, we can model different dependence patterns (e.g., dependence in the extremes, positive or negative skewness) and tail decays (e.g., exponential vs. polynomial).

Our in-sample and out-of-sample results showed the importance of fat tails and positive skewness in the multivariate distribution of energy risk factors. We also proposed comparing the tail risk estimates corresponding to the GH models and other more traditional procedures.

Regarding the tail risk of short positions, our backtest results confirmed that distributions with polynomial tail decay (heavy-tailed) outperformed alternative versions, especially for the utility portfolios. Ultimately, the extent to which we underestimate the tail risk of the portfolio loss distribution depends on the portfolio weights of the different energy commodities, whether we are analyzing the short or long trading position, and the horizon and confidence level considered.

It is worth mentioning that many power firms in liberalized markets have two main lines of business: electricity generation and electricity distribution. These days, and given the chronic generation overcapacity afflicting many developed markets (United States, Europe) most firms tend to focus more on the distribution business which implies an aggregate short position in electricity. The evidence we present suggests that conventional market risk measures (Gaussian VaR and ES) severely underestimate market risk under these circumstances. This fact should be taken into account not only by the company's shareholders and creditors but also by market regulators and supervisors.

4. VALUING ENERGY REAL ASSETS: THE CASE OF AN INTERCONNECTOR

Electricity markets have undergone a series of fundamental changes sparked by the liberalization of this industry. The first stage of liberalization required privatization of all or most of the generation assets, as well as privatization of the transmission grid which transports electricity from the generation points to the end consumer. Another important step in the development of the wholesale electricity markets is to exploit price differentials between locations by building interconnectors which are bidirectional transmission lines connecting the grids of two locations or the grids of two countries. Although interconnecting different grids is at the top of the political agenda in many countries, the decision to build them depends on their financial value.

Electricity prices are characterized by exhibiting extreme volatility and by undergoing abrupt changes (large upward spikes and large downward jumps), as well as fast mean reversion to a seasonal trend. This extreme behavior is also present in the difference between prices of two locations and explains why interconnecting two markets could be profitable. The main question we address in this chapter is how to value an interconnector. One of the key features that drives the financial value of an interconnector is that the owner has the right, but not the obligation, to transmit electricity between two locations. Therefore, once it has been built, the financial value of an interconnector is given by a series of real options which are written on the price differential between two electricity markets.

In this chapter we propose a valuation tool that uses real options theory to consider the problem and we employ market data of five pairs of European neighboring countries to value hypothetical interconnectors under realistic assumptions. The value of an interconnector is given by a strip of European-style options (Bull Call Spreads) written on the spread between the two markets and the valuation formula is in closed-form and is quick to implement. Our model for the spread captures the main characteristics of the dynamics of price differentials: jumps in both directions, high seasonal volatility, and fast mean reversion to a seasonal trend. We propose an algorithm to detect jumps where the emphasis is placed on avoiding misclassifying mean reversion as jumps. We estimate the parameters of the spread model and find that the introduction of jumps in the model delivers gains in the in-sample performance of between 20% and 48% with respect to a misspecified or "naïve" model in which jumps are not included.

We show valuations under different liquidity caps, which proxy for the depth of the interconnected power markets. We also derive no-arbitrage lower bounds for the value of the interconnector in terms of electricity futures contracts of the respective power markets.

We find that, depending on the depth of the market, the jumps in the spread can account for between 1% and 40% of the total value of the interconnector. The two markets where an interconnector would be most (resp. least) valuable are Germany and the Netherlands (resp. France and Germany). The markets where off-peak transmission between the two countries is more valuable than transmission during peak times are: France and Germany, France and UK, and the Netherlands and UK. We also provide "rules of thumb" to summarize the different drivers of the interconnector value.

4.1. Literature review and the market for interconnectors

In energy markets there are many projects whose value depends on the flexibility of being able to delay decision-making until more information becomes available. These decisions can include delaying or accelerating production, postponing entry, scaling production, changing technology, etc. In many cases the flexibility embedded in some types of project is what drives most of their value. For example, some electricity plants are only economically viable to operate when market prices are very high, otherwise they must be "switched off". Moreover, gas-fired plants are very valuable because relative to other plants (for instance nuclear and coal-fired ones) it is easier to ramp up or ramp down according to the level of market prices. Neglecting these embedded real options may seriously undervalue some projects to the extent that they might seem to deliver a negative NPV when in fact they are viable.

Real options in electricity markets are key components in project valuation. Power plants that offer operational flexibility derive most of their value from the option to produce electricity when prices are high. These options are valuable because wholesale electricity prices are extremely volatile, but the extreme behavior of power prices makes electricity prices a difficult commodity to model. Modeling power prices, and other contracts such as futures and forwards, can be found in Roncoroni (2002), Cartea and Figueroa (2005), and Escribano, Peña, and Villaplana (2011), among others.

An important feature common to all energy commodities is that their market value depends on the location and the date that the delivery of the commodity takes place. This is particularly important for electricity where date and location are crucial determinants of market clearing prices because electricity must be consumed immediately upon delivery, while consumption of other energy commodities such as gas and oil can be deferred by either postponing delivery or by storing them. In fact, as a consequence of the non-storability of electricity, one can think of electricity delivered over different intervals of the day, or throughout periods of the year, as different goods.

A further consequence of not being able to store electricity is that, strictly speaking, there are no electricity spot prices as commonly understood. Market clearing prices must be agreed prior to delivery at a time when production and demand are not known for sure; this uncertainty is resolved at the time when the physical transaction occurs. Therefore, the convention in the market and the literature is to treat the day-ahead prices as the spot prices, although their structure is more akin to that of a forward contract. Another standard way in which blocks of electricity are bundled is peak and off-peak. Peak hours correspond to a fixed interval of hours for business days characterized by high electricity demand, normally between 8am and 8pm. Off-peak hours belong to the interval between the end of a peak block and the beginning of the next one, and include the 24 hours of weekends' days and holidays.

The owner of the interconnector capacity needs to schedule the flows according to prevailing market prices and the transmission costs in the two interconnected locations. In practice these decisions are generally taken on the day-ahead market. Thus, we assume that the decision to use the interconnector to dispatch electricity from A to B, or vice versa, is based on the peak and off-peak market prices observed in the day-ahead market, net of transmission costs. Therefore, every day the owner of the interconnector capacity faces various alternatives. To commit to dispatching electricity the following day from A to B, or from B to A, during the peak and off-peak hours. To decide not to dispatch electricity in any direction during the peak and/or off-peak period.

4.2. Valuing interconnection capacity: a strip of real options

Our objective is to price the optionality provided by an interconnector that can exploit the wholesale electricity spot price differential between two markets. Every day the owner of the interconnector exercises the right to use the capacity to simultaneously buy electricity in market A, to sell the same quantity of electricity in market B, or vice versa. In other words, the owner of the capacity holds four daily European options: two options on the spread between A and B; and two options on the spread between B and A (one option for peak and the other for off-peak). Since each individual option is only for one day, we cannot cast the valuation problem in terms of futures contracts since the delivery period for these will be at least one month. Nevertheless the information provided by futures contracts can be used to determine no-arbitrage bounds for the European options on the spread.

The valuation problem thus reduces to being able to price European capped options. For ease of presentation let us focus on the spread between A and B, which we denote $S^{A,B}(t)$, and assume that it is for peak electricity, without specifying the particular hour during the peak segment.

Let $C_p^{A,B}(S^{A,B}, M, t; T, K^{A,B})$ denote the price of a European call at time t, written on the spread $S^{A,B}(t)$ during peak time, but capping the maximum value at $M > 0$, and expiring at a future date T with strike price $K^{A,B} < M$. The option gives the right to transmit 1 MWh of electricity, during a designated hour of the day, but for ease of notation we do not specify the particular hour of the day. The strike price represents the transmission costs between locations A and B, and time T represents the time in future periods when the decision will be made whether to use the interconnector capacity. Then, the price of the call is given by

$$C_p^{\mathrm{A,B}}\left(S^{\mathrm{A,B}}, M, t; T, K^{\mathrm{A,B}}\right) = \\ = e^{-\rho(T-t)}\mathbb{E}_t\left[\max\left(\min\{S^{\mathrm{A,B}}(T), M\} - K^{\mathrm{A,B}}, 0\right)\right] \tag{4.1}$$

where ρ is the risk-adjusted discount rate. The valuation problem of the capped European call (4.1) is also known in the literature as a Bull Call Spread. Note that capping the states of nature where the value of the call exceeds the cap M is equivalent to being long a standard European call option with strike $K^{A,B}$ and short a standard European call option with strike M written on the underlying $S^{A,B}(t)$. Hence

$$C_p^{\mathrm{A,B}}\left(S^{\mathrm{A,B}}, M, t; T, K^{\mathrm{A,B}}\right) = C_p^{\mathrm{A,B}}\left(S^{\mathrm{A,B}}, \infty, t; T, K^{\mathrm{A,B}}\right) \\ - C_p^{\mathrm{A,B}}\left(S^{\mathrm{A,B}}, \infty, t; T, M\right) \tag{4.2}$$

Where the standard European call is given by Equation (4.1) with $M = \infty$.

Generally, rights to interconnector capacity are sold over a period of time that covers a number of years and represents a significant proportion of the life of the interconnection assets. For expository purposes we will assume that the rights are in the form of a one-year lease and we value a lease for capacity of 1 MWh during peak times and 1 MWh during off-peak times. The value of the interconnector lease is given by the sum of all the capped European call options (one for every day of transmission from A to B and from B to A) between time t and expiry of the lease contract.

4.3. Modelling electricity spot price differentials

Modeling electricity prices, and other financial instruments related to this market, is quite recent in the academic literature. For instance, the work of Schwartz (1997) and Schwartz and Smith (2000) which considered storable commodities served as a platform for a number of articles that proposed no-arbitrage models for the dynamics of electricity prices.

Since the valuation of the call options embedded in the interconnector capacity is cast within the real options framework, the emphasis must be placed on a model that is specified under the statistical measure. Instead of estimating the parameters for the two markets A and B, we can value the interconnector capacity by modeling the difference in prices directly. Therefore, we can estimate the parameters of the spread model and use it as the departure point to value the European call options on the spread.

Here we propose a model for the spread in the spirit of the no-arbitrage spot price models which captures the most important features of the price dynamics, that is: large price spikes or jumps, strong mean reversion of large deviations and the presence of a seasonal component. In addition, we obtain the following three desired properties. First, the spread model also exhibits the stylized characteristics observed in the price difference between two locations, specifically large positive and negative deviations that mean revert very quickly to a seasonal trend. Second, the estimation of the spread model parameters can be achieved with the usual techniques. Third, the spread model specification enables us to calculate the price of European-style options by employing standard tools.

We propose, under the statistical measure, the following arithmetic model for the price differences between locations A and B,

$$S^{\mathrm{A,B}}(T) = f(T) + X(T) + Y(T) \tag{4.3}$$

Figure 4.1 Detecting jumps in the jump diffusion process

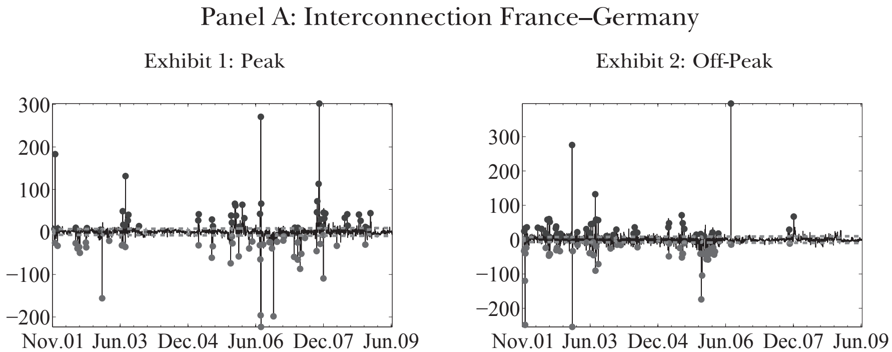

Panel B: Interconnection France–United Kingdom

Exhibit 1: Peak Exhibit 2: Off-Peak

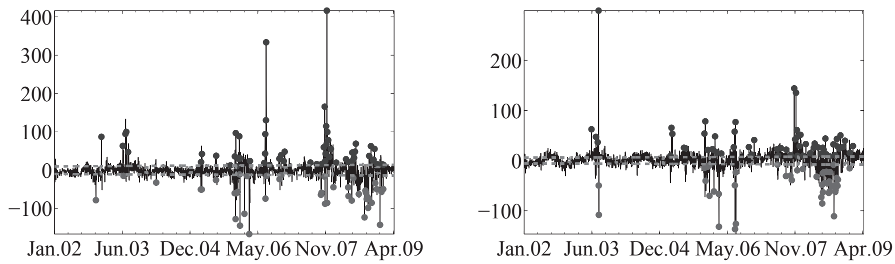

Panel C: Interconnection Nord Pool–Germany

Exhibit 1: Peak Exhibit 2: Off-Peak

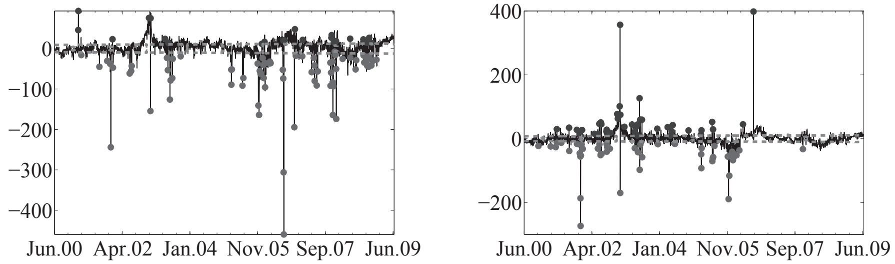

where $f(t)$ is a deterministic seasonal pattern (i.e. the long term trend of the spot) evaluated at time T, $X(T)$ is a mean reverting stochastic process at time T and $Y(T)$ is a zero-mean reverting pure jump process at time T with exponentially distributed jumps. We consider seasonal patterns in the trend, as well as in the volatility of the Gaussian component and in the intensity of the jumps.

Once we have the model for the price differences, we can proceed to value the European call options on the spread. The value of the call option is expressed in closed-form in Fourier space (see details in Cartea and González-Pedraz (2011)).

We apply a recursive semi-parametric filter to identify the calendar position of the jumps in the spread. The procedure identifies a hypothetical arrival of a jump when the detrended spread difference deviates, in absolute value, by more than three standard deviations from its mean. Once we have identified the jumps, we estimate their intensity by maximum likelihood. We analyze two scenarios, one where intensities are constant, and the other where intensities are time-dependent and may exhibit a seasonal pattern. In 16 out of 20 cases, we reject the model with constant intensity at a 10% significance level. In Figure 4.1 we show the positions of negative and positive jumps for some of the markets we study.

The last step of the estimation consists of estimating the mean-reversion rates of the Gaussian process X(t) and jump process Y(t), and the volatility parameters of X(t) by minimizing the mean-squared errors which are given by the average of the squared differences between the observed and the modeled spreads. Results show that spreads show significant mean reversion in jumps and in the Gaussian deviations. The half-life of the jumps ranges between 1 and 15 days approximately.

The introduction of jumps in the model delivers gains in the in-sample performance of between 20% and 48% compared to a "naïve" version, for which we do not include the jump process.

4.4. The market value of interconnectors

In this section we discuss the results of valuing interconnection capacity in neighboring European countries. We calculate the market value of a one-year lease of an interconnector that gives the lessee the right, but not the obligation, to transmit 1 MWh of electricity between two markets during peak and off-peak times.

We provide different values of the interconnector, which result from different assumptions about: the seasonal function of the spread; the liquidity and depth in both markets; and how jumps affect the extrinsic value of the real options used to calculate the value of the interconnector.

As discussed above, it does not seem plausible to exploit large price differentials due to liquidity reasons in the two markets. We cap the maximum price differentials that can be profited from at different levels: M in $\{10, 20, 30, 40, 50, \infty\}$ Euros/MWh, where we allow $M = \infty$ to include the hypothetical case where there are no liquidity constraints in the day-ahead market.

The values of the one year-lease can be broken into the four options available to the manager of the lease: transmit electricity from A to B and from B to A for both on-peak and off-peak segments of the day. Note that these values are for the use of the interconnector during the 365 days of the one-year lease. The total value of the lease is given by the sum of the four options.

The effect of the liquidity cap is different across the markets we study. For example, if the cap between Germany and the Netherlands is reduced from $M = \infty$ to $M = 50$ Euros/MWh, the value of the interconnector decreases by almost 75%. If we draw the same comparison in the UK–Netherlands market, the value of the interconnector only decreases by 8%. These different effects of the liquidity cap are due to the particular characteristics of the spread in each market: seasonal component, volatility of the OU process, jump intensities and jump sizes. Depending on the depth of the market, results

Figure 4.2 Interconnector values (dashed lines) and no-arbitrage bounds

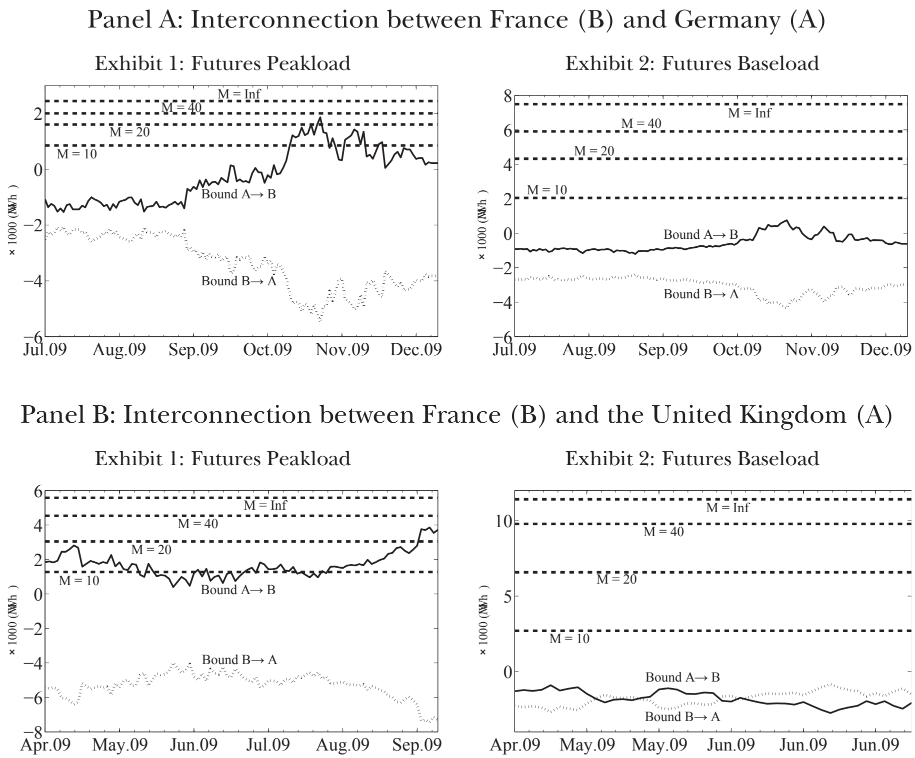

Panel C: Interconnection between Nord Pool (B) and Germany (A)

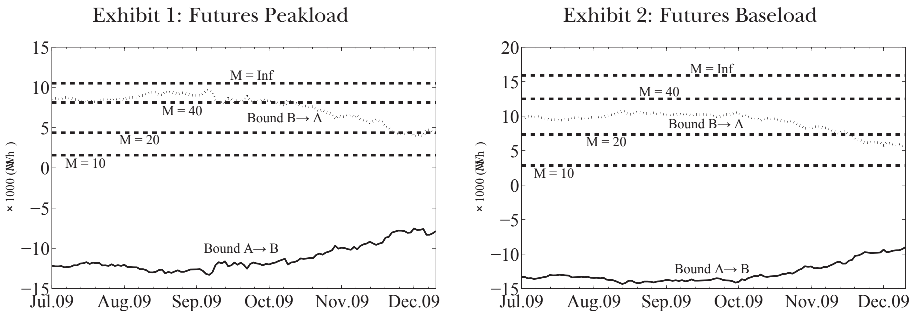

show that the jumps in the spread can account for between 1% and 40% of the total value of the interconnector.

Figure 4.2 shows the corresponding interconnector value for different liquidity constraints. Figure also shows the no-arbitrage bounds (Eur/MWh) resulting from buying a forward contract in location B, and shorting a forward contract for location A (bound $A \rightarrow B$), which are depicted by a solid line. Similarly, the dotted line represents the bound resulting from being long a forward contract in location A, and short a forward in location B (bound $B \rightarrow A$).

We can observe that the bounds exhibit considerable variation over time. These pronounced changes in the bounds are a reflection of changes in futures prices due to changes in: market expectations, fuel prices, changes in risk-premia, weather predictions, etc.

4.5. Conclusions

In this chapter we show how to value an electricity interconnector as an asset that gives the owner the optionality to manage electricity flows between two markets. In financial terms, the value of the interconnector is the same as a strip of real options written on the spread between the power prices of two markets. We model the spread prices based on a: seasonal trend, mean-reverting Gaussian process, and mean-reverting jump process.

As a first contribution, we express the value of these real options in closed-form in the presence of mean-reverting jumps processes. We also propose a methodology to detect jumps in the spread that addresses the possible miss-classification of mean reversion as jumps. We estimate the parameters of the spread model and find that the introduction of jumps in the model delivers significant gains in the in-sample performance.

Although we cast the problem in terms of real options, where the statistical distribution of the spread and the risk-adjusted discount rate are key ingredients in the valuation, we also derive no-arbitrage lower bounds for the value of the interconnector in terms of electricity futures contracts.

5. RISK PREMIUM IN COMMODITY MARKETS

Trying to understand the price behavior of commodities has a long tradition in the finance literature and is a long standing issue for stakeholders in commodity markets. On the one hand, there are market participants with exposure to spot price risks because they produce or consume the commodity. On the other hand, there are those that have no need to purchase or sell the commodity, but enter the market for speculative purposes. Either way, both types need to understand not only the behavior of spot prices but also the dynamics of the financial instruments written on the commodities so that decisions about bearing spot price risk, hedging, and speculation can be made.

From a reduced-form modeling perspective there are two possible ways to model price behavior. One way is first to build a model for commodity prices that tries to capture the main features of the price dynamics under both the data generating measure and risk-adjusted measure. The alternative way is to specify a reduced-form model under the risk-adjusted measure and place less importance on the dynamics of the commodity under the data generating measure. This second approach, although desirable in some cases, is built at the expense of not capturing some of the characteristics exhibited by price dynamics of the commodity under the data generating measure which in some cases, such as in energy commodities, might be an undesirable feature.

Our departure point in this chapter is that an important proportion of market participants are exposed to spot price commodity risk and it is their needs to hedge their positions the key factor which brings them to market to trade derivatives instruments to manage their exposures. Therefore, understanding the dynamics of commodities under the data generating measure is as important as understanding the dynamics of prices under the risk-adjusted measure. The main questions we set out to answer here are: how are the data generating and risk-adjusted measures related? How can we reconcile the behavior of the physical dynamics of spot prices with those of the different forwards with different expiries?

To answer this question, we first need to understand: what are the key elements that market participants price according to their risk preferences? What are the main risks that different stakeholder wish to offload? What happens to the risks that are being managed across different time horizons? What are futures contracts insurance for? Although our discussion could be applied to a wide variety of commodities, especially those whose prices tend to show a degree of mean reversion, we focus on two energy commodities: gas and electricity.

In general, establishing the link between the data generating measure and the risk-neutral measure in asset pricing is a difficult task because of market incompleteness. In energy markets in particular, the connection between these two measures is less well understood and has been overlooked in most cases, especially in electricity. So far, most of the reduced-form models for gas and electricity lack either a more realistic representation of prices under the data generating measure, or a better specification of the risk-adjusted measure to reconcile the dynamics of derivative instruments, for instance futures or forward contracts across different maturities.

This chapter contributes to the literature on the pricing of risk in commodities by proposing a parsimonious reduced-form model that can capture the main characteristics of commodity prices under the data generating measure and show that there is a family of risk-adjusted measures capable of capturing the fact that market participants may overstate (understate) the probability of occurrence of undesirable (desirable) events.

In particular we show that participants in energy markets price financial instruments under the risk-adjusted measure by modifying how long deviations from the seasonal component may last. In the most general version of our model there are three factors out of which two are mean reverting: one factor is an Ornstein-Uhlenbeck (OU) driven by Brownian motion and the other factor is a mean reverting jump process with positive and negative jumps.

Until now, all reduced-form models that specify jumps in prices under the data generating measure assume that under the risk-adjusted measure the jump component has all or most of the following characteristics: a) jumps are non-systematic or that the market-price of jump risk is zero; b) jumps have the same arrival structure under both measures; c) the distribution of jump sizes is the same under both measures; and d) the speed at which jumps mean revert is the same under both the data generating and risk-adjusted measures.

Assumptions a) to d) are at odds with the evidence observed in gas and power markets because forward contracts are in most cases bought (resp. sold) by consumers (resp. producers) of the commodity as insurance against large upward (resp. downward) price deviations that would have an adverse effect on their profits. For instance, assuming that jumps are non-systematic implies that in the cross section of futures prices the presence of jumps does not affect futures prices.

Assuming that spot prices mean revert at the same speed under both measures makes it difficult to reconcile the spot and forwards model dynamics with observed market prices. The family of risk-adjusted measures that we propose allow for the mean reversion of the jump component of spot prices to be different between the two measures.

In the empirical part, we estimate our model using the Bayesian inference for two types of energy commodities, natural gas and power. Specifically, we implement a Markov Chain Monte Carlo (MCMC) estimation scheme, which accounts for parameter uncertainty. Our results suggest that the degree of mean reversion under the physical and the risk-adjusted measures differ.

5.1. Literature review and a model for the spot

Reduced-form models for storable commodities have been around for a long time. Gibson and Schwartz (1990) propose a two-factor model where spot prices follow a geometric Brownian motion and the stochastic convenience yield follows an OU mean reverting process under the data generating measure. They propose a risk-adjusted measure which results from introducing a market-price of convenience yield risk in the form of a linear shift in the distribution of the convenience yield under the data generating measure. In their model the mean reversion of spot prices under both the data generating and risk-adjusted measure is induced by the mean reversion in the convenience yield. This model is extended in Schwartz (1997) and applied to oil, gold and copper. Hilliard and Reis (1998) further extend the model to include jumps in the spot price process. Another common feature to all these models is that when one of the factors of the model is mean reverting then the speed of mean reversion will be the same under both the data generating and risk-adjusted measure.

Casassus and Collin-Dufresne (2005) propose a three-factor model of spot prices, convenience yields and interest rates. The factor dynamics are driven by OU Brownian motion processes. The connection between the data generating measure and the risk-adjusted measure is introduced via a state dependent market price of risk for each factor. The immediate implication is that under the data generating measure the speed of mean reversion of spot prices, convenience yields, and interest rates can be different from the speed of mean reversion under the risk-adjusted measure.

We propose a reduced-form, arithmetic model for commodity spot prices; and a new, more flexible change of measure, for which pricing and calibration of basic building blocks such as futures contracts can be performed analytically.

Let $X(t)$ be the vector of state variables $(X_1(t), X_2(t), X_3(t))'$. The evolution of the state variables $X(t)$ under the real probability measure P is given by

$$dX(t) = -A_P X(t)dt + C_P\, dL(t) \tag{5.1}$$

where $dL(t) = \left(dL_1(t), dL_2(t), dL_3(t)\right)'$; and $L_1(t)$ and $L_2(t)$ are independent Brownian motions and $L_3(t)$ is an independent compound Poisson process with intensity parameter λ and jump sizes distributed as a normal distribution. Here A_P is a 3x3 diagonal matrix that reflects the mean-reversion rates of the state variables under the physical measure P. The 3x3 lower triangular matrix C_P defines the dependence between Gaussian state variables.

Under the physical measure, the spot price process $S(t)$ can be decomposed into a stochastic component defined by the state variables and a deterministic component. Notice that the model is arithmetic and allows eventually for possible negative prices in the spot.

We now consider the following change of measure Q:

$$dL(t) = dL^Q(t) + \Lambda(t)dt \tag{5.2}$$

with

$$\Lambda(t) = \Phi_Q + C_P^{-1}B_Q X(t) \tag{5.3}$$

Avoiding the technical details, whit this change of measure, we are considering possible changes in the mean-reversion rate under Q for the mean-reverting Gaussian and mean-reverting jump processes. Then, under the Q-measure, the dynamics of state variable vector $X(t)$ is given by

$$dX(t) = (C_P\Phi_Q - A_Q X(t))dt + C_P dL^Q(t) \tag{5.4}$$

where $A_Q = A_P - B_Q$; and B_Q and is a diagonal matrix of constant parameters.

The effect of the stochastic terms $X_i(t)$ in the measure change can be seen as a complete shift by a *constant*. When $X_i(t)$ is low or for $B_{Q,i} = 0$, we essentially have the usual transform rescaling. But if, for example, the process $X_i(t)$ exhibits large shocks, we shift the whole compensator measure by a factor. In the case of the jump process, this shift could be interpreted as if we rescale the jump intensity in the compound Poisson process. That is, if this is a negative (positive) rescaling, we *lessen* (*intensify*) the compound Poisson process under Q.

The futures price $f(t,T)$ at time t of a contract to deliver one unit of commodity at time $T > t$ is defined as the expected spot price under the risk-adjusted probability measure Q. For energy's futures contracts, instead of a single maturity time, every contract specifies a delivery period, which are typically a month, quarter, season, or even year. Hence, if $F(t,T_1,T_2)$ denotes the market price for an energy's futures contract with time T_1 until maturity and delivery period $[T_1,T_2]$, then

$$F(t,T_1,T_2) = \frac{1}{T_2 - T_1} E_t^Q\left[\int_{T_1}^{T_2} S(u)du\right] \tag{5.5}$$

Where $t < u$ and $T_1 \leq u \leq T_2$. Therefore, by definition, we have that $S(t) = f(t,t) = F(t,t,t+1)$.

Furthermore, lets define the risk premium for Q as the difference between the futures price and the predicted spot price. The theory of normal backwardation says that producers are willing to pay a premium for having their production hedged, implying a positive risk premium. In electricity, there is some empirical evidence for negative risk premium in the short end, that is, for contracts close to maturity. Under our specification, we may observe a change in the sign of the risk premium, for example, from positive to negative along the forward curve, a change which is stochastically dependent on the price factors $X_i(t)$.

5.2. Econometric methodology and data

This section describes the estimation problem when both spot and futures prices are observed. Our implementation is based on an MCMC simulation that provides inference for unobserved state variables and model parameters given information under both probability measures.

Figure 5.1 Spot and futures prices

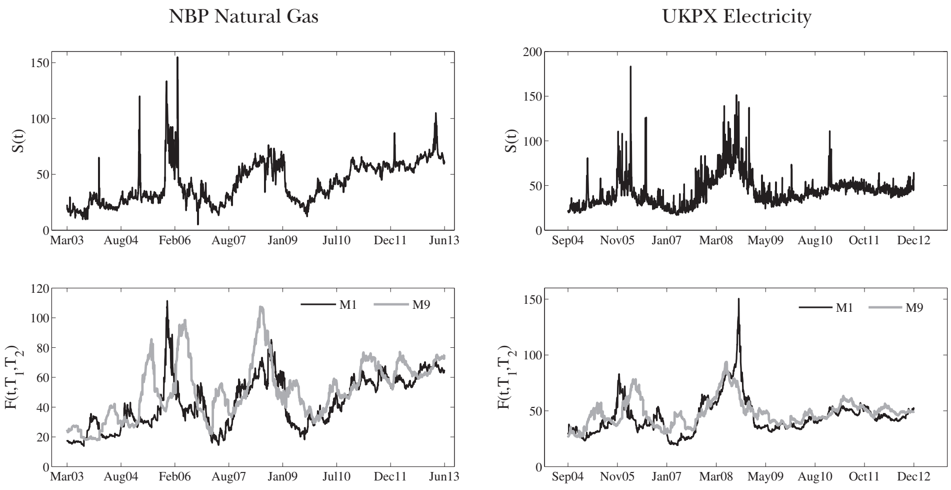

Consider the observations of spot prices and monthly futures prices for N different maturities $Y_{t,j}$ with $j = 1, \ldots, N+1$. We assume that market prices for these instruments are observed with independent pricing errors with respect to the theoretical values; which, in our model specification, are determined by the vector of state variables X_t and the set of model parameters Θ. There are mainly two motivations for adding an additive pricing error: to consider the possibility that the model can be misspecified and that there can be measure errors related with noisy price observations.

The specification of the pricing errors allows us to describe the conditional likelihood of the set of observations, $p(Y|X,\Theta)$. Then, the inference problem consists of the computation of the joint posterior density of latent variables and parameters:

$$p(X,\Theta|Y) \propto p(Y|X,\Theta)p(X|\Theta)p(\Theta) \tag{5.6}$$

The MCMC approach provides a method to sample parameters and latent variables from their joint posterior density (Johannes and Polson, 2009).

The empirical analysis of this chapter is based on a sample of spot (day ahead) and futures prices of natural gas and electricity. In particular, we use daily data of monthly futures contracts of natural gas and electricity, both traded at the Intercontinental Exchange (ICE) in pence/therm and pounds/MWh, respectively. We consider contracts with maturities 1, 3, 6, and 9 months. Delivery is made equally each day throughout delivery period, in this case, one month. Figure 5.1 exhibits the price of spot, M1, and M9 contracts on natural gas and electricity over their respective sample periods.

5.3. Empirical results

Given the samples of the posterior distributions for the parameter vector and the state variables, we can obtain straightforward the Monte Carlo estimate of the mean of parameters and latent variables. Our estimates take into account parameter

Figure 5.2 Estimated jump probabilities

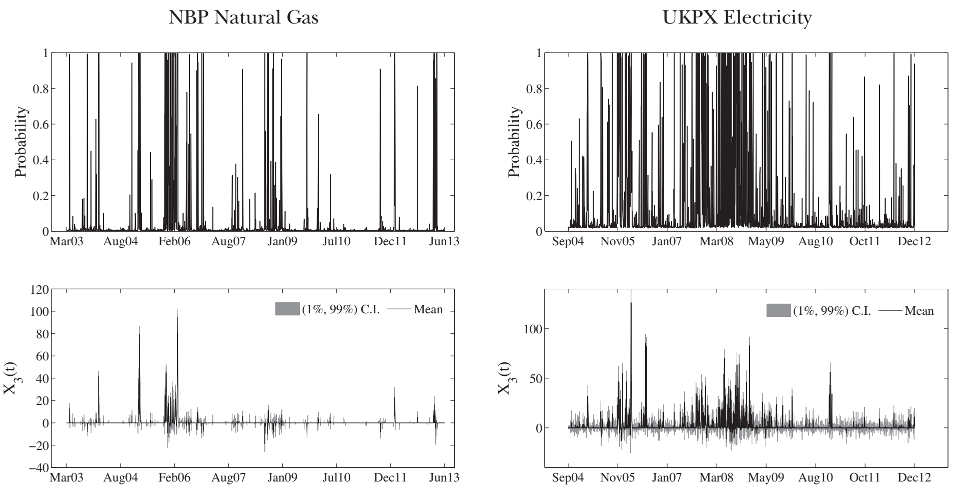

uncertainty, when we estimate the mean of the posterior state-variable distribution, we are considering the fact that parameter estimates are random variables. We also compare the more general model with other nested specifications with less state variables: the one-factor and two-factor Gaussian Models (OU model and OU-AB model, respectively); and the previous two Gaussian models plus a mean-reverting jump diffusion factor (OU-MRJ and OU-AB-MRJ models).

Examining the jump intensity and the jump size parameters (see Figure 5.2); we observe that jumps are more frequent than in other asset classes, such as equity. The intensity estimates show that on average we can expect around 5 jumps in 100 trading days for NBP natural gas, and about 15 jumps for UKPX electricity. The estimates of the mean jump size are largely positive and significant.

According to OU-AB-MRJ estimates, the average jump size is in the range [6.4, 14.8] with 95\% probability for natural gas, and in the range [11.7, 16.5] for electricity; while the posterior mean of the standard deviation of jump sizes is 16.8 and 16.9 for gas and electricity, respectively. These results indicate that, when jumps occur, positive size jumps are more common than negative ones.

The mean reversion rates of the Gaussian OU variable and the jump factor under the physical measure are significant for both commodities and for all models. The estimates of speed of mean reversion of the jump factor under P for the OU-MRJ and OU-AB-MRJ show that are around 1.44 and 1.53 for natural gas, and 1.69 and 1.70 for electricity. That means the observed half-life of the price deviations corresponding to the jump factor is less than 1 day. For both jump models, the estimates of the mean-reversion rate of the Gaussian state-variable corresponds to an estimated half-life between 75 and 1.0 days, and 57 and 1.2 days, for natural gas and electricity, respectively.

Under the Q measure, the estimates of the speed of mean-reversion of the Gaussian factor decrease. For instance, for the jump diffusion model OU-AB-MRJ, the 95%-probability range of expected half-life for natural gas is [124, 433], in days. On the contrary, the estimates of the mean-reversion rate corresponding to the jump factor are much higher under Q than under the physical measure P. For electricity, the observed half-life values are very close to zero. These results suggest that the presence of a jump is largely vanished under the risk-adjusted measure right after the jump occurs. Further analysis will be required to understand all the implications of these findings.

Finally, we find that the inclusion of mean-reverting jumps reduces on average the estimates of the measure error variances. In particular, for natural gas, the pricing error variances of the spot prices diminish 54% and 51% when using the OU-MRJ and OU-AB-MRJ models instead of the OU and OU-AB models. For electricity, this reduction is even greater, 64% and 54%, respectively.

Conclusions

In this chapter, we propose a new way of thinking about the market price of risk so that market participants bearing spot commodity risk are compensated for: jump arrival risk; jump size risk; and speed of mean reversion risk of both diffusions and/or jumps. When pricing under the risk-adjusted measure agents will: over-state the time it takes to return to the seasonal trend; alter the mean of the process; and change the intensity of the jumps and their average size. Our approach can also be viewed as a special case of stochastic discount factors that not only affect the mean of the process but also its variance via the persistence of shocks to the economy.

Using a panel data on natural gas and electricity futures, we empirically estimate the model using a MCMC methodology. We are able to analyze the behavior of a mean-reverting jump component under both probability measures. We find evidence of different speeds of mean reversion under the physical and the risk-adjusted measures.

Although our specification does not include specifically a dynamic for the convenience yield, we can infer from the instantaneous interest rate dynamics the stochastic process for the convenience yield or, in terms of market price of risk, we can infer an implicit stochastic market price of risk.

REFERENCES

Artzner, P., F. Delbaen, J.M. Eber, and D. Heath (1999): "Coherent risk measures", *Mathematical Finance*, 9(3), 203-228.

Bibby, B., and M. Sørensen (2003): "Hyperbolic processes in finance", in *Handbook of Heavy Tailed Distributions in Finance*, ed. by S. Rachev, pp. 211-248. Elsevier, Amsterdam.

Börger, R.H., Á. Cartea, R. Kiesel, and G. Schindlmayr (2009): "Cross-commodity analysis and applications to risk management", *Journal of Futures Markets*, 29(3), 197-217.

Büyüksahin, B., and M. Robe (2010): "Speculators, commodities, and cross-market linkages", *Working paper*, Commodity Futures Trading Commission.

Cappiello, L., R. Engle, and K. Sheppard (2006): "Asymmetric dynamics in the correlations of global equity and bond returns", *Journal of Financial Econometrics*, 4(4), 537-572.

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

Cartea, Á., and C. González-Pedraz (2012): "How much should we pay for interconnecting electricity markets? A real options approach", *Energy Economics*, 34(1), 14-30.

Casassus, J., and P. Collin-Dufresne (2005): "Stochastic convenience yield implied from commodity futures and interest rates", *Journal of Finance*, 60(5), 2283-2331.

Casassus, J., P. Liu, and K. Tang (2013): "Economic linkages, relative scarcity, and commodity futures returns", *Review of Financial Studies*, 26(5), 1324-1362.

Daskalaki, C., and G. Skiadopoulos (2011): "Should investors include commodities in their portfolios after all? New evidence", *Journal of Banking and Finance*, 35(10), 2606-2626.

Demarta, S., and A.J. Mcneil (2005): "The t copula and related copulas", *International Statistical Review*, 73(1), 111-130.

Eberlein, E., and G. Stahl (2003): "Both sides of the fence: A statistical and regulatory view of electricity risk", *Energy and Power Risk Management*, 8(6), 32-36.

Embrechts, P., F. Lindskog, and A. Mcneil (2003): "Modelling dependence with copulas and applications to risk management", in *Handbook of heavy tailed distributions in finance*, ed. by S. T. Rachev, pp. 329-384. Elsevier, Amsterdam.

Erb, C.B., and C.R. Harvey (2006): "The strategic and tactical value of commodity futures", *Financial Analysts Journal*, 62(2), 69-97.

Escribano, A., J.I. Peña, and P. Villaplana (2011): "Modelling electricity prices: International evidence", *Oxford Bulletin of Economics and Statistics*, 73(5), 622-650.

Etula, E. (2013): "Broker-dealer risk appetite and commodity returns", *Journal of Financial Econometrics*, 11(3), 486-521.

Geman, H. (2005): *Commodities and commodity derivatives: Modelling and pricing for agriculturals, metals, and energy*. Wiley, New York, NY.

Gibson, R., and E. Schwartz (1990): "Stochastic convenience yield and the pricing of oil contingent claims", *Journal of Finance*, 45(3), 959-976.

Giot, P., and S. Laurent (2003): "Market risk in commodity markets: A VaR approach", *Energy Economics*, 25(5), 435-457.

Gorton, G., and K.G. Rouwenhorst (2006): "Facts and fantasies about commodity futures", *Financial Analysts Journal*, 62(2), 47-68.

Guidolin, M., and A. Timmermann (2008): "International asset allocation under regime switching, skew, and kurtosis preferences", *Review of Financial Studies*, 21(2), 889-935.

Hansen, B.E. (1994): "Autoregressive conditional density estimation", *International Economic Review*, 35(3), 705-730.

Harvey, C.R., J.C. Liechty, M.W. Liechty, and P. Müller (2010): "Portfolio selection with higher moments", *Quantitative Finance*, 10(5), 469-485.

Harvey, C.R., and A. Siddique (1999): "Autoregressive conditional skewness", *Journal of Financial and Quantitative Analysis*, 34(4), 465-487.

Hilliard, J., and J. Reis (1998): "Valuation of commodity futures and options under stochastic convenience yields, interest rates, and jump diffusions in the spot", *Journal of Financial and Quantitative Analysis*, 33(1), 61-86.

Hong, H., and M. Yogo (2012): "What does futures market interest tell us about the macro-economy and asset prices?", *Journal of Financial Economics*, 105(3), 473-490.

Johannes, M., and N. Polson (2009): "Markov chain Monte Carlo", in *Handbook of Financial Time Series*, ed. by T. Mikosch, J.-P. Kreiß, R. Davis, and T. Andersen, pp. 1001-1013. Springer Verlag, Berlin.

Jondeau, E., and M. Rockinger (2003): "Conditional volatility, skewness, and kurtosis: existence, persistence, and comovements", *Journal of Economic Dynamics and Control*, 27(10), 1699-1737.

Jondeau, E., and M. Rockinger (2012): "On the importance of time variability in higher moments for asset allocation", *Journal of Financial Econometrics*, 10(1), 84-123.

Kat, H.M., and R.C.A. Oomen (2007b): "What every investor should know about commodities. Part II: Multivariate return analysis", *Journal of Investment Management*, 5(3), 16-40.

Longin, F., and B. Solnik (2001): "Extreme correlation of international equity markets", *Journal of Finance*, 56(2), 649-676.

Mcneil, A.J., R. Frey, and P. Embrechts (2005): *Quantitative risk management: Concepts, techniques, and tools*. Princeton University Press, Princeton, NJ.

Miltersen, K., and E. Schwartz (1998): "Pricing of options on commodity futures with stochastic term structures of convenience yields and interest rates", *Journal of Finance*, 33(1), 33-59.

Patton, A.J. (2006a): "Estimation of multivariate models for time series of possibly different lengths", *Journal of Applied Econometrics*, 21(2), 147-173.

Patton, A.J. (2006b): "Modelling asymmetric exchange rate dependence", *International Economic Review*, 47(2), 527-556.

Pirrong, C. (2012): *Commodity price dynamics: A structural approach*. Cambridge University Press, New York, NY.

Roncoroni, A. (2002): "A class of marked point processes for modeling electricity prices", *Doctoral thesis*, Université Paris IX Dauphine.

Routledge, B.R., D.J. Seppi, and C.S. Spatt (2000): "Equilibrium forward curves for commodities", *Journal of Finance*, 55(3), 1297-1338.

Schwartz, E.S. (1997): "The stochastic behavior of commodity prices: Implications for valuation and hedging", *Journal of Finance*, 52(3), 923-973.

Schwartz, E.S., and J.E. Smith (2000): "Short-term variations and long-term dynamics in commodity prices", *Management Science*, 46(7), 893-911.

Tang, K., and W. Xiong (2012): "Index investment and financialization of commodities", *Financial Analysts Journal*, 68(6), 54-74.

Tseng, C.L., and G. Barz (2002): "Short-term generation asset valuation: A real options approach", *Operations Research*, 50(2), 297-310.

Diciembre 2016

The increasing presence of investors and financial intermediaries in commodity markets, together with the huge increase in the volatility of commodity prices, have renewed the interest in commodities and commodity derivatives. In the last decade, a better understanding of the behavior of commodity prices and their idiosyncratic statistical features has emerged as a relevant financial and policy topic. This book tries to provide new insights, first, to analyze the multivariate distribution of commodity returns and its impact on portfolio selection and tail risk measures; and, second, to price commodity derivatives under the presence of non-Gaussian shocks in a continuous time framework.

Promotora editorial:

Fundación UCEIF
Fundación de la Universidad de Cantabria para el Estudio y la Investigación del Sector Financiero

SANFI SANTANDER FINANCIAL INSTITUTE

UC UNIVERSIDAD DE CANTABRIA

Santander

978-84-8102-802-7 9 €

www.editorialuc.es

IBIC: KCH, KCJ