Hedging downside risk of oil refineries: A vine copula approach

Kunlapath Sukcharoen *, David J. Leatham

Department of Agricultural Economics, Texas A&M University, United States

A B S T R A C T

The financial health of an oil refinery greatly depends on its refining margin or the difference between the prices of its refined products (typically, gasoline and heating oil) and the cost of crude oil. The refinery may hedge against the downside risk of unfavorable price movements using crude oil, gasoline, and heating oil futures. This paper examines the use of a vine copula approach to estimate multiproduct hedge ratios that minimize the downside risk of the refinery. The advantage of the vine copula approach is that it allows us to capture important characteristics of petroleum price changes, including skewness and fat-tailedness in the marginal distributions of individual price change series as well as heterogeneous (tail) dependence patterns between different pairs of price changes. The out-of-sample hedging effectiveness of two popular classes of vine copula models – canonical (C-) and drawable (D-) vine copula models – are evaluated and compared with that of the widely used nonparametric method and three standard multivariate copula models. The empirical results reveal that the D-vine copula model is a good and safe choice in managing the downside risk of the refinery.

© 2017 Elsevier B.V. All rights reserved.

1. Introduction

A typical oil refinery purchases crude oil and sells refined products (e.g., gasoline and heating oil). Its refining or profit margin is then related to the spread between the prices of refined products and the price of crude oil. Thus, the refinery faces downside risk in both crude oil and refined product markets. As can be seen from Fig. 1, since late 2005, a large decline in the refining margin (due to the simultaneous adverse movements in the petroleum prices) has appeared to be quite common. The risk of losses because of unfavorable petroleum price movements clearly signifies the importance of hedging the joint downside risk of input and output prices. Accordingly, the goal of this paper is to develop a multiproduct futures hedging model that minimizes the downside risk of the refinery.[1]

Solving for the minimum-downside risk hedge ratios requires the estimation of the entire joint distribution of spot and futures price movements. For single-product hedging, the standard practice is to rely on a nonparametric method – in particular, the empirical distribution or historical simulation method (Lien and Tse, 2000; Demirer and Lien, 2003; Harris and Shen, 2006). This approach is very flexible and could be easily extended to the case of multiproduct hedging. However, it often produces inaccurate estimates of extreme quantiles due to its heavy dependence on historical data (McNeil and Frey, 2000; Pritsker, 2006; Cao et al., 2010). Recently, Barbi and Romagnoli (2014) propose a standard bivariate Archimedean copula model for estimating downside-risk hedge ratios in a single-product setting. They show that their proposed method produces greater downside risk reductions than the nonparametric approach. The superior performance is likely due to the model's ability to capture important characteristics of asset returns, including skewness and fat-tailedness in the distributions of individual asset returns as well as their nonlinear and asymmetric dependence relationship. These characteristics are also found in crude oil and refined product markets (Hammoudeh et al., 2003; Grégoire et al., 2008; Chang et al., 2010; Ji and Fan, 2011; Serra and Gil, 2012; Aloui et al., 2014).

While hedging models that incorporate these characteristics (in particular, the nonlinear and asymmetric dependence relationship between asset returns) lead to better hedging outcomes, they have been limited to the case of single-asset hedging. This is because, when dealing with more than two random variables (i.e., when hedging more than

---

* Corresponding author.
  E-mail address: kunlapath.s@gmail.com (K. Sukcharoen).

[1] Multiproduct hedging involves the use of multiple futures contracts to hedge exposures to price risks in multiple commodities. In this study, crude oil, gasoline, and heating oil futures are used simultaneously to hedge the refining company's exposures to adverse price movements in the crude oil, gasoline, and heating oil spot markets. In contrast, single-product hedging uses a single futures contract to hedge a spot position in a particular commodity market.
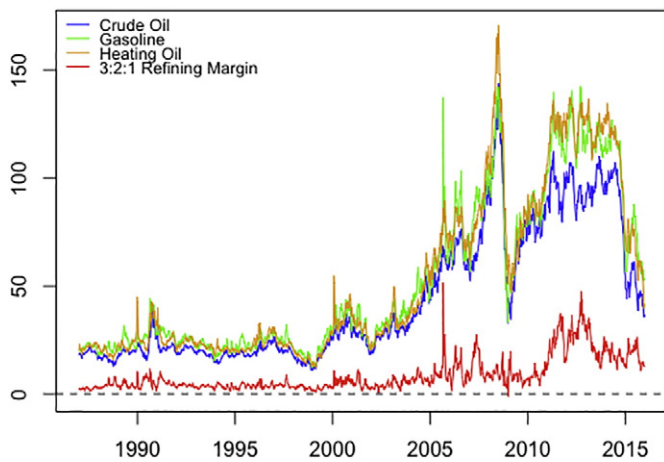
Fig. 1. Weekly crude oil spot prices, gasoline spot prices, heating oil spot prices, and 3:2:1 refining margin (unhedged). Notes: The 3:2:1 refining margin approximates the profitability of a typical U.S. refinery which is able to convert 3 barrels of crude oil to 2 barrels of gasoline and 1 barrel of heating oil.

one asset), standard multivariate copulas are less flexible as they restrict the degree of tail dependence (or comovements during extreme market conditions) between all pairs of variables to be identical. For example, suppose a standard multivariate Archimedean copula is used to model the dependence structure of crude oil, gasoline, and heating oil returns. This means that the degree of tail dependence between crude oil and gasoline returns is assumed to be the same as that between crude oil and heating oil returns and also the same as that between gasoline and heating oil returns. This is clearly too restrictive. Instead of relying on the standard multivariate copulas, one could model the dependence relationship of multiple variables using more advanced multivariate copulas (known as "vine copulas").

The vine copula model, initially introduced by Joe (1996) and first estimated by Kurowicka and Cooke (2006), is a relatively new class of multivariate copula models. Similar to the standard multivariate copula models (e.g., the standard Gaussian, Student's t, and Archimedean copula models), the vine copula model is able to account for both skewness and fat-tailedness in the univariate marginal distributions. This is because the model allows us to separate the modeling of the marginal distributions from the dependence structure that links these marginal distributions to form a joint distribution. However, while the standard copula models require all pairs of variables to have the same tail dependence patterns, the vine copula model permits different tail dependence specifications for different pairs of variables (Czado, 2010; Brechmann and Schepsmeier, 2013). Accordingly, this presents an important opportunity for developing a new multiproduct hedging model that is able to capture the potentially complex (nonlinear, asymmetric, and heterogeneous) dependence patterns among multiple petroleum markets. As such, we propose to combine a vine copula model with Monte Carlo simulation to construct the joint distribution of spot and futures price changes.[2]

In particular, the proposed hedging model builds the joint distribution of multiple variables using an empirical distribution function for the marginal distributions and two different classes of vine copulas – the canonical (C-) and drawable (D-) vine copulas (Kurowicka and Cooke, 2005) – for the dependence structure. The C– and D-vine copula models are estimated using a sequential maximum likelihood procedure proposed by Aas et al. (2009), and the joint distribution is generated using Monte Carlo simulation. The optimal hedge ratios are then derived through a numerical optimization method for four alternative downside-

risk hedging objectives: the minimization of Semivariance (SV), Lower Partial Moment (LPM), Value at Risk (VaR), and Expected Shortfall (ES) of the refinery's hedged margin. The usefulness of the proposed model is evaluated through an extensive out-of-sample hedging exercise. Its performance is also compared with that of the widely used nonparametric method and three standard multivariate copula models (namely, the standard Gaussian, Student's t, and Clayton copula models).[3]

This paper contributes to the literature by estimating multiproduct hedge ratios for oil refineries in a downside-risk framework. Previous studies in this area have mainly focused on deriving either minimum-variance or mean-variance hedge ratios.[4] However, it is well known that the variance is not a proper risk measure when asset returns are non-normal because businesses and investors are only concerned with downside risks but not upside risks (Lien and Tse, 1998; Unser, 2000; Veld and Veld-Merkoulova, 2008). Despite the awareness of the non-normality of asset returns, studies on downside risk hedging in a multi-product setting are still scarce.[5] One of the few studies is Power and Vedenov (2010) who estimate the minimum-LPM hedge ratios for a feedlot operator (whose profit depends on the prices of corn, feeder cattle, and fed cattle) and compare them with the minimum-variance hedge ratios. Another is Awudu et al. (2016) who consider a hedging problem of a corn-based ethanol producer and derive the mean-VaR hedge ratios based on two distributional specifications: multivariate normal and Gaussian copula distributions. The other two studies are Chen et al. (2016) and Liu et al. (2017); the former derives mean-VaR hedge ratios for grain processors using standard multivariate copulas, whereas the later estimates minimum-LPM hedge ratios for oil refineries. This paper also develops a multiproduct hedging model in a downside-risk framework. Similar to Liu et al. (2017), we focus on the oil refining industry. However, we consider four (not just one) alternative measures of downside risk. This allows us to examine the sensitivity of the results vis-à-vis the downside risk measures used. In addition, unlike other studies, this paper analyzes the usefulness of the proposed model through an extensive out-of-sample hedging exercise. The out-of-sample performance of different hedging objectives for the best performing hedging model is also evaluated using various hedging effectiveness measures. Moreover, while the vine copula methodology has been applied to study the dependence structures of financial and asset markets (Allen et al., 2013; Zhang, 2014; Zimmer, 2015), to forecast VaR and ES of financial portfolios (Weiß and Supper, 2013; Brechmann et al., 2014; Zhang et al., 2014), and to analyze asset allocation problems (Low et al., 2013; Riccetti, 2013; Bekiros et al., 2015), this is the first study to examine the use of vine copula approach in the context of hedging downside risk. Our findings would benefit oil refineries (as well as other multiproduct hedgers), and provide a richer understanding of the usefulness of vine copulas in energy risk management.

The remainder of this paper is organized as follows. Section 2 describes a methodology. Section 3 presents data and preliminary analysis. Section 4 reports and discusses the empirical results. Section 5 concludes the paper.

2. Methodology

2.1. Oil Refinery's hedging problem

In the empirical analysis, the stylized problem of a typical oil refinery whose profit depends on the refining margin is considered. We focus on a 3:2:1 refining margin, which approximates the profitability of a typical

---

[2] Following Haigh and Holt (2002) and Alexander et al. (2013), our hedging analysis is based on the price changes. The reasons for why the price changes should be used instead of the log returns or percentage returns are discussed in Alexander et al. (2013).

[3] The standard Clayton copula model is a commonly used Archimedean copula model due to its ability to capture lower tail dependence among variables.

[4] See, for example, Haigh and Holt (2002), Ji and Fan (2011), and Alexander et al. (2013) for previous studies on multiproduct hedging of an oil refinery.

[5] Non-normality of petroleum prices and returns are documented in many studies such as Hammoudeh et al. (2003), Chang et al. (2010), Ji and Fan (2011).

U.S. refinery that converts 3 barrels of crude oil to 2 barrels of gasoline and 1 barrel of heating oil. The refinery may hedge its exposures to downside risk in the three petroleum markets using crude oil, gasoline and heating oil futures.

Following Haigh and Holt (2002), we assume that the refinery takes futures positions in period $t-1$ (long crude oil futures, and short gasoline and heating oil futures) and liquidates all futures positions in period t (when the purchase of crude oil and the sales of refined products occur). Accordingly, the refinery's hedged margin (or profit per barrel) at time t is:

$$\pi_t(\boldsymbol{b}) = -S_t^C + \frac{2}{3}S_t^G + \frac{1}{3}S_t^H + b_C\left(F_t^C - F_{t-1}^C\right) + \frac{2}{3}b_G\left(F_{t-1}^G - F_t^G\right) + \frac{1}{3}b_H\left(F_{t-1}^H - F_t^H\right) \tag{1}$$

where superscripts and subscripts C, G and H refer to crude oil, gasoline and heating oil, respectively; S_t and F_t denote spot and futures prices at time t, respectively; and $\boldsymbol{b} = \{b_C, b_G, b_H\}$ are hedge ratios determined at time $t-1$. For simplicity, we assume that other costs are deterministic and thus do not affect hedging decisions. Prices at time $t-1$ are known at time t, whereas prices at time t are random (stochastic) variables.

The hedged margin in Eq. (1) can be rewritten in terms of spot and futures price changes:

$$\pi_t(\boldsymbol{b}) = -\Delta S_t^C + \frac{2}{3}\Delta S_t^G + \frac{1}{3}\Delta S_t^H + b_C\Delta F_t^C - \frac{2}{3}b_G\Delta F_t^G - \frac{1}{3}b_H\Delta F_t^H + S_{t-1}^{CS} \tag{2}$$

where $\Delta S_t = S_t - S_{t-1}$ denotes the changes in spot prices; $\Delta F_t = F_t - F_{t-1}$ denotes the changes in futures prices; and $S_{t-1}^{CS} = -S_{t-1}^C + \frac{2}{3}S_{t-1}^G + \frac{1}{3}S_{t-1}^H$. The last term in Eq. (2), S_{t-1}^{CS}, is known at the time the hedge is initiated, and hence does not cause a variation in the refiner's profit margin at time t. Therefore, similar to Alexander et al. (2013), we focus on hedging the risky portion of the hedged margin at time t, denoted by:

$$y_t(\boldsymbol{b}) = -\Delta S_t^C + \frac{2}{3}\Delta S_t^G + \frac{1}{3}\Delta S_t^H + b_C\Delta F_t^C - \frac{2}{3}b_G\Delta F_t^G - \frac{1}{3}b_H\Delta F_t^H \tag{3}$$

where $y_t(\boldsymbol{b}) = \pi_t(\boldsymbol{b}) - S_{t-1}^{CS}$. In Alexander et al. (2013), $y_t(\boldsymbol{b})$ is known as the hedged (portfolio) profits and losses (P&Ls).

The refinery's objective is then to select the optimal hedge ratios $\boldsymbol{b}^*$ that minimize the downside risk of the hedged P&Ls. Mathematically,

$$\boldsymbol{b}^* = \arg\min_{\boldsymbol{b}} Risk(y_t(\boldsymbol{b})) \tag{4}$$

where $Risk(y_t(\boldsymbol{b}))$ is the measure of downside risk defined on $y_t(\boldsymbol{b})$. In this study, we consider four standard measures of downside risk: the SV, LPM, VaR, and ES, which we describe in more detail in the next section.

2.2. Downside risk measures

The first downside-risk measure considered is the Semivariance (SV). The SV, introduced in Roy (1952), measures the variability of P&Ls that fall below the target level. It is defined as:

$$SV = \int_{-\infty}^{c} (c-y_t)^2 dF(y_t) \tag{5}$$

where c is the target P&L; y_t is the random P&L; and F is the distribution function of y_t. As the basic goal of hedging is to avoid loss (i.e., y_t being less than zero), we select the target P&L equal to zero (that is, $c = 0$).

The second measure is the nth-order lower partial moment (LPM_n). The LPM_n, proposed by Fishburn (1977), is a generalization of the SV, and is defined as:

$$LPM_n = \int_{-\infty}^{c} (c-y_t)^n dF(y_t) \tag{6}$$

where c is the target P&L; $n > 0$ is the level of hedger's risk tolerance; y_t is the random P&L, and F is the distribution function of y_t. Fishburn (1977) shows that $0 < n < 1$ reflects risk-seeking behavior, $n = 1$ captures risk-neutral behavior, and $n > 1$ corresponds to risk-averse behavior. For the similar reason as above, we assume $c = 0$. In addition, we consider $n = 3$ to focus on a risk-averse hedger.

The third measure is Value-at-Risk (VaR). The VaR measures the largest potential loss over a certain period of time (for this study, over one week) for a particular confidence level (p). More generally, VaR at the confidence level p is given by:

$$VaR_p = -F^{-1}(1-p) \tag{7}$$

where F is the distribution of y_t. In this study, the VaR is calculated for three different confidence levels: $p = 0.90, 0.95$ and 0.99.

The forth risk measure is Expected Shortfall (ES). It measures the expected loss given that losses exceed the VaR. The ES at the confidence level p is given as:

$$ES_p = -E[y_t | y_t \leq -VaR_p] \tag{8}$$

Similar to the VaR, the ES is calculated for $p = 0.90, 0.95$ and 0.99.

2.3. Empirical procedure

Solving for the minimum-SV, minimum-LPM, minimum-VaR, and minimum-ES hedge ratios is technically very demanding. This is because the calculation of SV, LPM, VaR, and ES depends on the entire joint distribution of the six random variables in Eq. (3). In this study, we use a multivariate copula approach to model the joint distribution of random variables.

The copula approach has been widely used in a variety of empirical work to model joint distributions of random variables.[6] The models are essentially based on the Sklar's theorem (Sklar, 1959), which states that any n-dimensional multivariate distribution can be decomposed into n individual marginal distributions and a copula that describes the dependence structure. More formally,

$$F(x_1, x_2..., x_n) = C(F_1(x_1), F_2(x_2), ..., F_n(x_n)) \tag{9}$$

where F is a joint distribution of $x_1, x_2, ..., x_n$ with marginal distributions $F_i = F_i(x_i)$ for $i = 1, 2, ..., n$, and $C : [0,1]^n \to [0,1]$ is a copula function. Suppose that F_i and C are differentiable. Then, the joint density function is defined as:

$$f(x_1, x_2..., x_n) = f_1(x_1)f_2(x_2)\cdots f_n(x_n)[c(F_1(x_1), F_2(x_2), ..., F_n(x_n))] \tag{10}$$

where $f_i = f_i(x_i)$ is the (unconditional) density of F_i and c is the density of the copula.

The most important practical implication of the decompositions in Eqs. (9) and (10) is that, to construct a joint distribution, we can separate the modeling of the n marginal distributions from the modeling of the dependence structure. For the marginal distributions, they can be modeled parametrically or nonparametrically. Given a rich variety of univariate marginal distributions available, features such as skewness and fat-tailedness in each marginal distribution of price movements can be captured easily. As for the choice of copula families or dependence

---

[6] We refer the reader to Joe (1997) and Nelsen (2006) for details on construction, properties, and applications of copulas.

structures, a natural starting point might be any standard Archimedean copulas (typically, a standard Clayton copula) as they allow us to capture nonlinear and asymmetric dependence between random variables. However, the standard Archimedean copulas use only one or two parameters to describe the dependence structure among the n random variables and thus may not be able to adequately capture the dependence structure when $n \geq 3$. For instance, if a standard Clayton copula is used to model the dependence structure of the six random variables in Eq. (3), all possible pairs of random variables are assumed to have the same degree of comovements during both normal and extreme market conditions. This is very restrictive.

As a result, a common approach to building a joint distribution of more than two variables is to restrict attention to the elliptical copulas such as standard Gaussian and Student's t copulas. This is because, at least, standard Gaussian and Student's t copulas permit different pairs of variables to have different degree of comovements during normal market conditions. Nevertheless, they still restrict the tail dependence parameters (i.e., the degree of comovements during extreme market conditions) between all pairs of variables to be identical. In addition, unlike standard Archimedean copulas, they assume a symmetric tail dependence structure. That is, they require the degree of comovements during extreme market upturns to be the same as that during extreme market downturns.

We could, however, go beyond these standard multivariate copulas by using a vine copula approach, which is a more advanced and flexible alternative method of modeling the dependence structure (Joe, 1996; Bedford and Cooke, 2001; Aas et al., 2009). The key advantage of this approach over the standard copula approach is that it allows different pairs of variables to have heterogeneous dependence patterns during both normal and extreme market conditions. It can also account for nonlinearity and asymmetry in the dependence structure of each pair of variables. Therefore, a potentially complex (nonlinear, asymmetric, and heterogeneous) dependence relationship among multiple variables can be modeled.

Technically, a vine copula is a multivariate copula that is generated via a cascade of standard (conditional) bivariate copulas (called pair-copulas) and marginal distribution functions. In other words, the vine copula, which describes the dependence among multiple variables, is constructed by mixing a group of different standard bivariate copulas with each bivariate copula characterizing the dependence pattern of each pair of variables. The idea of the vine copula construction (also known as the pair-copula construction) can be easily illustrated using a three-dimensional case. Without loss of generality, the multivariate density of x_1, x_2 and x_3 can be represented as a product of unconditional and conditional densities:

$$f(x_1, x_2, x_3) = f_1(x_1) f_{2|1}(x_2|x_1) f_{3|1,2}(x_3|x_1, x_2) \quad (11)$$

where $f_{i|j,k} = f_{i|j,k}(x_i|x_j, x_k)$. Using the Sklar's theorem in (10), the first conditional density in (11) can be written as:

$$f_{2|1}(x_2|x_1) = \frac{f(x_1, x_2)}{f_1(x_1)} = c_{1,2}(F_1(x_1), F_2(x_2)) f_2(x_2) \quad (12)$$

where $c_{1,2}$ is a copula function linking x_1 and x_2. In a similar manner, the second conditional density can be written as:

$$f_{3|1,2}(x_3|x_1, x_2) = \frac{f_{2,3|1}(x_2, x_3|x_1)}{f_{2|1}(x_2|x_1)} \quad (13)$$
$$= c_{2,3|1}\left(F_{2|1}(x_2|x_1), F_{3|1}(x_3|x_1)\right) f_{3|1}(x_3|x_1)$$

where $f_{3|1}(x_3|x_1) = c_{1,3}(F_1(x_1), F_3(x_3)) f_3(x_3)$. Accordingly, the joint density function in (11) can be decomposed further as

$$f(x_1, x_2, x_3) = f_1 f_2 f_3 c_{1,2}(F_1, F_2) c_{1,3}(F_1, F_3) c_{2,3|1}\left(F_{2|1}, F_{3|1}\right) \quad (14)$$

with the conditional distribution functions $F_{i|j}(x_i|x_j)$ defined as:

$$F(x|\mathbf{v}) = \frac{\partial C_{x, v_j | \mathbf{v}_{-j}}\left(F(x|\mathbf{v}_{-j}), F(v_j|\mathbf{v}_{-j})\right)}{\partial F(v_j|\mathbf{v}_{-j})} \quad (15)$$

where $C_{x, v_j | \mathbf{v}_{-j}}$ is a conditional bivariate copula and $\mathbf{v}_{-j}$ is the vector $\mathbf{v}$ with the component v_j removed (Joe, 1997). That is, the joint density function can be expressed in terms of individual marginal distributions and a group of bivariate copulas – collectively referred to as a vine copula. Because each of the bivariate copulas (two of them are unconditional, $c_{1,2}(F_1, F_2)$ and $c_{1,3}(F_1, F_3)$, and one is conditional on x_1, $c_{2,3|1}(F_{2|1}, F_{3|1})$) that collectively forms a vine copula does not have to come from the same copula family, the vine copula approach allows for different types of the dependence patterns for each pair of variables. This provides an enormous flexibility in modeling high-dimensional dependence structure and could possibly lead to better hedging outcomes.

The decomposition in Eq. (14) is not unique. More specifically, $f(x_1, x_2, x_3)$ can also be represented as $f_1 f_2 f_3 c_{1,3} c_{2,3} c_{1,2|3}$ or as $f_1 f_2 f_3 c_{1,2} c_{2,3} c_{1,3|2}$. Consequently, a difficulty lies in selecting a vine copula structure – the specification indicating which pair-copulas are conditional on which other variables – from a large number of possible vine copula constructions. In this study, we consider two popular classes of vine copula structures: canonical (C-) and drawable (D-) vine structures (Kurowicka and Cooke, 2005). In essence, the C- and D-vine copula structures for n variables $(x_1, x_2, ..., x_n)$ can be represented graphically as a sequence of $(n-1)$ connected trees or vine trees $(T_1, T_2, ..., T_{n-1})$.[7] Fig. 2 (upper panel) represents a C-vine copula structure for six variables. In every tree of the C-vine copula structure, there is one variable that is connected to all the other variables. More specifically, tree T_1 indicates that the dependence patterns between x_1 and all the other variables (x_2, x_3, x_4, x_5) are modeled by unconditional pair-copulas. Tree T_2 indicates that the dependence patterns between x_2 and all other variables (except x_1) are modeled by conditional pair-copulas with x_1 as a conditioning variable. Tree T_3 indicates that the dependence patterns between x_3 and all the other variables (except x_1 and x_2) are modeled by conditional pair-copulas with x_1 and x_2 as conditioning variables, and so on. Accordingly, the joint density function associated with the six-dimensional C-vine copula structure is given by:

$$f(x_1, x_2, ..., x_6)$$
$$= f_1 f_2 f_3 f_4 f_5 f_6 c_{1,2} c_{1,3} c_{1,4} c_{1,5} c_{1,6} c_{2,3|1} c_{2,4|1} c_{2,5|1} c_{2,6|1} c_{3,4|1,2} c_{3,5|1,2} c_{3,6|1,2}$$
$$c_{4,5|1,2,3} c_{4,6|1,2,3} c_{5,6|1,2,3,4}$$
$$(16)$$

Fig. 2 (lower panel) represents a D-vine copula structure for six variables. In every tree of the D-vine copula structure, each variable is connected to at most two other variables. Specifically, tree T_1 indicates that the dependence patterns between any adjacent variables (x_1 and x_2; x_2 and x_3; x_3 and x_4; x_4 and x_5; and x_5 and x_6) are modeled by unconditional pair-copulas. Tree T_2 suggests modeling the dependence pattern between x_1 and x_3 conditional on x_2, x_2 and x_4 conditional on x_3, and so on. In the same manner, the dependence pattern between any two variables x_i and x_j in the remaining trees is modeled conditional on the variables that lie between the variables x_i and x_j in tree T_1 as conditioning variables. For example, the dependence pattern between x_1 and x_4 is modeled conditional on x_2 and x_3 (refer to T_3). Accordingly, the joint density function associated with the six-dimensional D-vine copula is given by:

$$f(x_1, x_2, ..., x_6)$$
$$= f_1 f_2 f_3 f_4 f_5 f_6 c_{1,2} c_{2,3} c_{3,4} c_{4,5} c_{5,6} c_{1,3|2} c_{2,4|3} c_{3,5|4} c_{4,6|5} c_{1,4|2,3} c_{2,5|3,4} c_{3,6|4,5}$$
$$c_{1,5|2,3,4} c_{2,6|3,4,5} c_{1,6|2,3,4,5}$$
$$(17)$$

---

[7] We refer the reader to Bedford and Cooke (2001), Kurowicka and Cooke (2005), and Aas et al. (2009) for details of these vine copulas.

C-vine copula structure
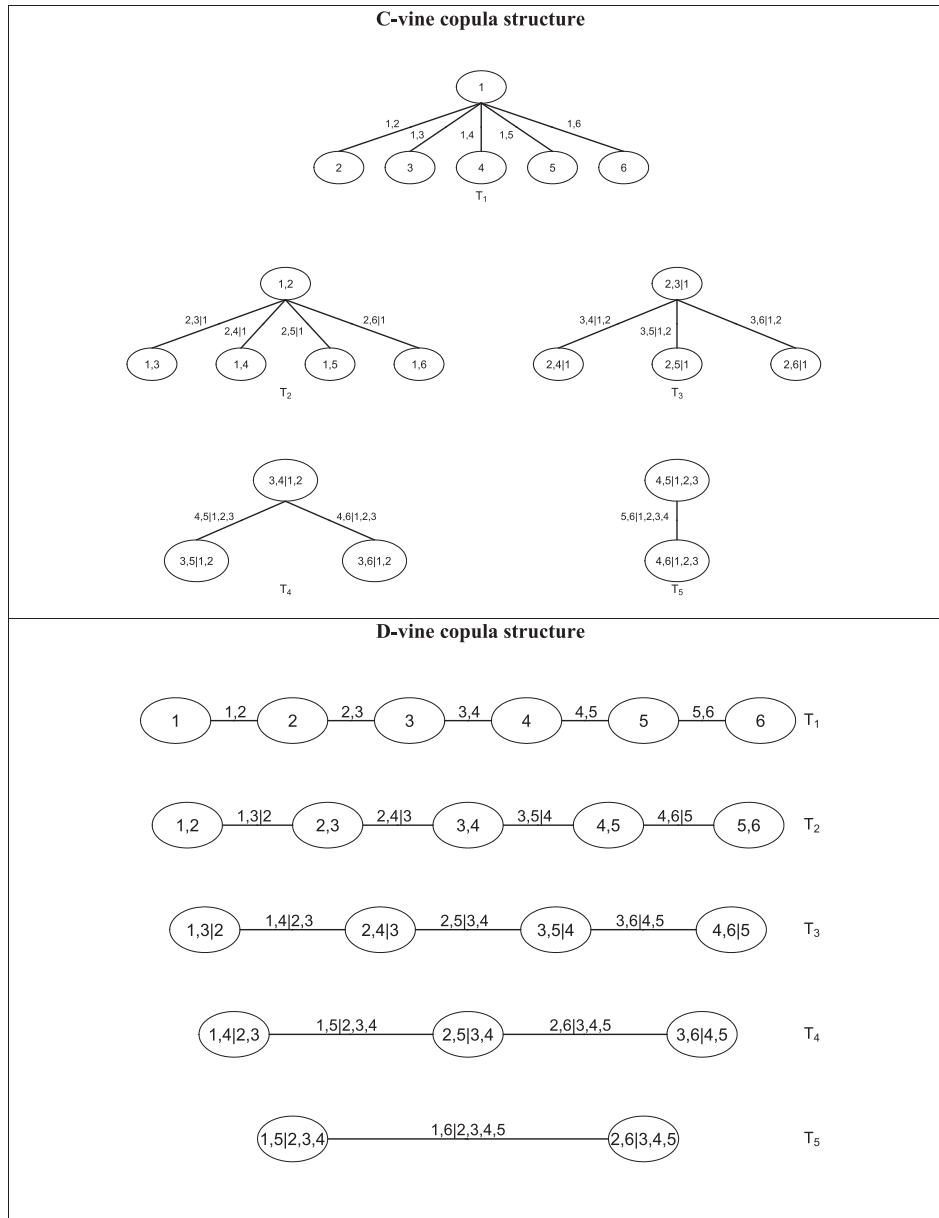

D-vine copula structure

Fig. 2. Six-dimensional C- (upper panel) and D-vine (lower panel) copula structures.

The procedure for fitting a joint distribution function using the C- or D-vine copula can be briefly summarized in four steps. The first step is to model the marginal distributions. For each price change series, we estimate its marginal distribution using an empirical distribution function and then transform the price change series into copula data (that is, a standard uniform variable).[8] The second step is to select an order of the variables for the C- or D-vine copula structures. For the C-vine copula

structure, we follow Czado et al. (2012) and select the variable that has the highest degree of association with all the other variables as the first variable. The degree of association is measured by summing the absolute values of pairwise Kendall's tau coefficients $\Theta_{\tau}^i = \sum_{j=1, i \neq j}^n |\tau_{i,j}|$ for each variable i. The variable that has the highest degree of association with the remainder of the variables is then selected as the second variables, and so on. For the D-vine copula structure, we follow Dißmann et al. (2013) and order the variables such that the sum of the absolute values of pairwise Kendall's tau coefficients $\Theta_{\tau} = \sum_{i=1}^{n-1} |\tau_{i,i+1}|$ is maximized. The third step is to choose a bivariate copula for each pair-copula. This study uses a sequential estimation approach proposed by Aas et al. (2009) with the Akaike Information Criterion (AIC) as a selection criterion. We consider 31 different parametric bivariate copulas.[9] Then, the

---

[8] It should be noted that the marginal distribution could also be estimated using a parametric estimation method. In this study, each marginal distribution is estimated nonparametrically in order to allow for the univariate asymmetry as well as to avoid the possible misspecification of parametric distributions (Charpentier et al., 2007). Similar to Bouyé and Salmon (2009), Power and Vedenov (2010), and Barbi and Romagnoli (2014), the marginal distributions are estimated using the unfiltered data. Other studies first apply a GARCH filter to the original data and then model the dependence structure of the filtered series (see, for example, Hsu et al., 2008; Lee, 2009; Sukcharoen et al., 2015 for studies on copula-based hedge ratios). An advantage of using the unfiltered data is that it allows us to avoid the first-stage estimation errors – the errors in the estimation of conditional mean and variance models, which could lead to the copula approach being inferior to the nonparametric approach that constructs the distribution functions of random variables directly from the unfiltered data.

[9] This is the maximal list of the R package: CDVine (Brechmann and Schepsmeier, 2013). The 31 bivariate copulas include Gaussian, Student's t, Clayton, Gumbel, Frank, Joe, BB1 (Clayton-Gumbel), BB6 (Joe-Gumbel), BB7 (Joe-Clayton), BB8 (Joe-Frank) copulas and the rotated versions (90, 180 and 270 degrees) of Clayton, Gumbel, Joe, BB1, BB6, BB7, and BB8 copulas.

final step is to estimate all copula parameters. The estimation is done sequentially starting from the first vine tree, where the maximum pseudo likelihood method described in Genest et al. (1995) is employed.

After obtaining all necessary parameters, we compute the four downside risk measures using a Monte Carlo simulation method. More specifically, the estimated vine copula densities are used to generate 10,000 draws of the six standard uniform variables, $\{u_{1,s}, u_{2,s}, \ldots, u_{6,s}\}_{s=1}^{10,000}$. For each variable i, these draws are converted to draws from the joint distribution of price changes using its inverse distribution function of the price change series. These simulated spot and futures price changes are then used to compute the refiner's hedged P&Ls in Eq. (3). For each hedging objective, the optimal hedge ratios are then derived by solving the minimization problems in Eq. (4) numerically using the Nelder-Mead direct search method (Nelder and Mead, 1965).

This study examines the usefulness of the C- and D-vine copula models in dealing with the downside risk in the refining industry based on their hedging effectiveness. For each hedging objective, the hedging effectiveness is measured as a percentage reduction in the downside risk of the hedged P&Ls relative to that of the unhedged P&Ls[10]:

$$\text{HE} = \left(1 - \frac{Risk(y_t(\boldsymbol{b}^*))}{Risk(y_t(0))}\right) \times 100 \qquad (18)$$

where $y_t(\boldsymbol{b}^*)$ is the hedged P&L, $y_t(0)$ is the unhedged P&L, and $Risk(\cdot)$ is SV, LPM, VaR, or ES, depending on the hedging objective. We also compare the hedging effectiveness of the vine copula models to that of the nonparametric method and three standard multivariate copula models: namely, the standard Gaussian, Student's t, and Clayton copula models.[11]

3. Data and preliminary analysis

We use weekly Wednesday closing spot and futures prices for West Texas Intermediate (WTI) crude oil, unleaded gasoline, and number 2 heating oil. In the rare cases where Wednesday prices are missing, Tuesday prices are taken instead.[12] All prices are obtained from the Datastream database, and converted into dollars per barrel. The price data span from December 31, 1986 to December 30, 2015, from which a sample of weekly changes in spot and futures prices are constructed.[13] To calculate the changes in the futures prices, the closing prices for the nearest-to-expiration futures contracts are used with the rollover occurring on Wednesday a week before the expiry of the contract.[14] At

the rollover date, care has been taken to ensure that the changes in futures prices are calculated using the same futures contract. Altogether, this results in a total of 1513 weekly observations for the changes in spot and futures prices.

Table 1 reports summary statistics (Panel A) and correlation matrix (Panel B) for the weekly changes in the spot and futures prices for the entire sample period. For each price change series, the mean is very small relative to its standard deviation. The changes in spot and futures prices of refined products (both gasoline and heating oil) are more volatile than those of raw material (crude oil), and for each commodity the price changes in the spot market is more volatile than the futures market. All price change series are slightly skewed and exhibit high excess kurtosis, suggesting that the price changes are not normally distributed. The significant Jarque-Bera test statistics for all the price change series confirm that the changes in spot and futures prices do not follow a normal distribution. The Augmented Dicky-Fuller (ADF) tests suggest that all the price change series are stationary. The correlation coefficients of at least 0.58 for all series pairs indicate that, more than half of the time, all the price series move in the same direction. The price changes in the spot and its corresponding futures markets are highly correlated (the correlation coefficients exceed 0.88) indicating that, in general, they have the same change trend.

Turning to the core of our empirical analysis, we evaluate the different hedging methods based on their out-of-sample hedging effectiveness. In the out-of-sample analysis, the following rolling window approach is followed.[15] First, we estimate the minimum-SV, minimum-LPM, minimum-VaR, and minimum-ES hedge ratios using the first 261 weekly observations. That is, our estimation window is approximately 5 years.[16] Next, the estimated optimal hedge ratios are used to construct the hedged P&Ls for the following 130 weeks (i.e., 2.5 years) for each hedging objective. Then, the estimation window is moved forward by

Table 1
Summary statistics and correlation analysis on weekly changes in spot and futures prices.

	ΔS^C	ΔS^G	ΔS^H	ΔF^C	ΔF^G	ΔF^H
Panel A: Summary statistics						
Mean	0.0123	0.0221	0.0141	−0.0045	0.1284	0.0196
Min	−14.5600	−44.3184	−14.7756	−14.4100	−13.8600	−14.7840
Max	14.1300	57.4560	18.7320	14.0800	28.9380	18.0180
SD	2.3879	3.6183	2.8382	2.3060	2.9175	2.6700
Skew	−0.1879	1.2066	0.0824	−0.3204	0.4437	−0.0134
Ex. Kurt.	5.1967	57.0524	6.0686	5.4245	8.3822	5.6270
J-B	1718.5*	206,140.0*	2332.6*	1888.6*	4495.1*	2004.2*
ADF	−27.2498*	−29.1573*	−28.1465*	−25.9375*	−26.6913*	−26.7788*
Panel B: Correlation matrix						
ΔS^C	1.0000	0.5849	0.7700	0.9775	0.7338	0.8315
ΔS^G		1.0000	0.6066	0.5995	0.8822	0.6443
ΔS^H			1.0000	0.7871	0.7210	0.9483
ΔF^C				1.0000	0.7496	0.8527
ΔF^G					1.0000	0.7819
ΔF^H						1.0000

Notes: Summary statistics (Panel A) and correlation matrix (Panel B) are presented for the weekly changes in the spot and futures prices for the period January 7, 1987 to December 30, 2015. The total number of observations is 1513 for each price change series. ΔS^C, ΔS^G, ΔS^H, ΔF^C, ΔF^G and ΔF^H denote the changes in crude oil spot, gasoline spot, heating oil spot, crude oil futures, gasoline futures, and heating oil futures prices, respectively. SD, Skew, and Ex. Kurt. represent sample standard deviation, skewness, and kurtosis, respectively. J-B is the Jarque-Bera test statistic, where * denotes the rejection of the null hypothesis of normality at the 1% significance level. ADF is the Augmented Dickey-Fuller test statistic, where * denotes the rejection of the null hypothesis that the respective price change series follows a unit root process at the 1% significance level.

---

[10] This is a variant of the measure of hedging effectiveness proposed by Ederington (1979).

[11] The nonparametric method adopted here is similar to the nonparametric approach by Harlow (1991), Rockafellar and Uryasev (2002), and Lien and Tse (2000). This approach is also known as the historical simulation method or the empirical distribution method.

[12] In the extremely rare cases where both Wednesday and Tuesday prices are missing, Monday prices are taken instead.

[13] Prices during the period of abnormal market conditions caused by Hurricane Katrina – from August 29, 2005 to September 9, 2005 – are removed from the estimation of optimal hedge ratios (but not from the evaluation of hedging effectiveness). During this period, the gasoline spot (futures) price increased by 57.46 (28.94) dollars per barrel in the first week and decreased by 44.32 (9.79) dollars per barrel in the second week. Given the abnormally large difference between the gasoline spot and futures price movements, including these data points in the estimation of optimal hedge ratios causes the optimal hedge ratio for gasoline to be negative (recommending a speculative position in the gasoline futures market – longing gasoline futures instead of shorting) regardless of the hedging model used. Because these abnormal price discrepancies between gasoline spot and futures markets were only temporary and the possibility of a Katrina-like event occurring in the near future is extremely low, taking a speculative position in the gasoline futures market in the following weeks would result in a huge loss. Therefore, excluding these data points when estimating optimal hedge ratios is justifiable. Note that Alexander et al. (2013) also remove these data points from the estimation of hedge ratios.

[14] Carchano and Pardo (2009) show that the choice of rollover date to construct the changes in the futures price series is not relevant. Also, the problem of thin market trading is of limited importance because for the commodities under consideration trading continues in high volumes right up to the futures expiration dates.

[15] Conlon and Cotter (2012, 2013) and Barbi and Romagnoli (2014) adopt a similar approach in their out-of-sample analyses. This rolling window approach allows us to account for the time variation in the distribution of price changes as well as to test the model's hedging effectiveness over a number of test windows.

[16] Alexander et al. (2013) and Barbi and Romagnoli (2014) also use a 5-year rolling window approach. Studies regarding the optimal length of the moving window for multiproduct hedging are still needed.

1 week, where the optimal hedge ratios and associated out-of-sample hedged P&Ls – the hedged P&Ls for the following 130 weeks – are recalculated.[17] This approach produces 1123 out-of-sample test windows. Finally, within each test window, the out-of-sample hedging effectiveness for each hedging objective is computed for all the hedging models. The mean and median hedging effectiveness are then calculated across the 1123 test windows.[18]

4. Empirical results

This section first presents evidence on the fit of the three standard multivariate copula models – the standard Gaussian copula (SGC), standard Student's t copula (SSC), and standard Clayton copula (SCC) models – and the two vine copula models – the C- and D-vine copula models. The section then proceeds to present our empirical findings for optimal crude oil, gasoline, and heating oil hedge ratios obtained using different hedging models (including the nonparametric (NP), SGC, SSC, SCC, C-vine copula and D-vine copula models). Then, comparisons of out-of-sample hedging effectiveness are made across different hedging models and hedging objectives. Finally, the out-of-sample performance of different hedging objectives for the best performing hedging model is assessed using various measures of hedging effectiveness.

4.1. Model fit

Table 2 provides some evidence on the fit of the five multivariate copula models: the SGC, SSC, SCC, C-vine copula, and D-vine copula models.[19] On average, the D-vine copula model yields the highest log-likelihood and lowest values of the AIC and Bayesian Information Criterion (BIC), whereas the SCC model provides the worst fit to the data.[20] The results are very consistent across all the 1123 estimation windows.[21]

The average number of parameters for each copula model is also listed in Table 2. The SCC model has only one parameter to characterize the overall dependence structure of the six random variables. It is very likely that this parameter restriction is a reason for the poor fit of the SCC model. The SGC model uses 15 pairwise correlation coefficients to capture the dependence structure of the random variables. However, it assumes no tail dependence, and could therefore underestimate the joint probability of extreme movements in all the petroleum prices. In addition to the 15 pairwise correlation coefficients, the SSC model adds one more parameter (a degree of freedom parameter) to characterize the tail dependence for all pairs of the random variables. However, using only one parameter to describe the overall tail dependence may be over-simplistic when dealing with more than two variables. These parameter restrictions are likely reasons for the superior fit of the vine copula models over the standard multivariate copula models.

Comparing between the two vine copula models, the superiority of the D-vine copula model may be explained by the difference in the way that the two models decompose the joint density function (specifically, the difference in the structure of the first tree). Referring to Fig. 2 (upper panel) and Eq. (16), the first tree of the C-vine copula model uses

Table 2
Average log-likelihood (LLH), Akaike Information Criterion (AIC), and Bayesian Information Criterion (BIC), and number of parameters for the multivariate copula models.

	LLH	AIC	BIC	Number of parameters
SGC	1169.54	− 2307.99	− 2254.52	15
SSC	1316.85	− 2600.68	− 2543.64	16
SCC	618.61	− 1234.78	− 1231.21	1
C-vine	1356.41	− 2666.33	− 2585.45	23
D-vine	1373.63	− 2700.73	− 2619.70	23

Notes: Each model is estimated using a rolling window approach with a window of approximately 5 years or 261 weeks. The total number of estimation windows is 1123 windows. SGC is the standard Gaussian copula model. SSC is the standard Student's t copula model. SCC is the standard Clayton copula model. C-vine is the canonical vine copula model. D-vine is the drawable vine copula model.

only one variable to link with the other five variables through different unconditional bivariate copulas. As a result, the first tree of the C-vine copula model captures the high dependence between the spot and its corresponding futures price changes in only one petroleum market. On the other hand, the first tree of the D-vine copula model permits a direct link between the spot and its corresponding futures price changes for all petroleum markets (see Fig. 2 (lower panel) and Eq. (17)). In other words, the variables in the first tree of the D-vine copula model for each estimation window can be ordered such that the spot and its corresponding futures variables are next (or linked) to each other. For example, the structure $S_t^H - F_t^H - F_t^C - S_t^C - F_t^G - S_t^G$ is selected for the first tree of the D-vine copula for our first estimation window. This feature is not allowed by the C-vine copula model and may be a reason why the D-vine copula model better fits the data than the C-vine copula model.

4.2. Minimum-downside risk hedge ratios

Table 3 reports the average minimum-SV, minimum-LPM, minimum-VaR, and minimum-ES hedge ratios (as well as their respective standard deviations) generated using different hedging models. On average, most hedging models (except the SCC model) recommend the hedge ratios of fairly similar magnitude for all hedging objectives. Depending on the hedging models and objectives, the average crude oil, gasoline, and heating oil hedge ratios are between 0.8 and 1.3. On the other hand, the SCC model yields the optimal gasoline and heating oil hedge ratios fairly close to 0 (no hedge), and the optimal crude oil hedge ratios slightly smaller than 0 (recommending a speculative position in the crude oil futures market – shorting crude oil futures instead of longing). The huge difference between the hedge ratios generated from the SCC model and the other copula models is due to the fact that the SCC model uses only one parameter to capture the dependence patterns across all the six markets.

Examining the standard deviations of the optimal hedge ratios, the heating oil hedge ratios are found to be much more volatile than the crude oil and gasoline hedge ratios. This corresponds to the relatively high level of excess kurtosis in the heating oil price changes, implying that the extreme price changes are observed more often in the heating oil market than in the other two markets. Thus, the optimal tail risk-minimizing hedge ratios for heating oil are more sensitive to the extreme price changes. In addition, the NP method generates the most volatile hedge ratios for all the three petroleum commodities (except for the case of minimum-SV hedge ratio for gasoline). This may be because the NP approach is very sensitive to new information from the data (especially when only 261 observations are used in the estimation of the nonparametric or empirical distribution). Further, the minimum-VaR and minimum-ES hedge ratios at the 99% confidence level are generally more dispersed than the other hedging objectives, which may be explained by the greater level of difficulty in estimating the extreme tails of the true distribution of the hedged P&Ls.

Table 3
Average optimal hedge ratios (with standard deviations in parentheses) of different hedging models and hedging objectives.

Model	Hedging objective							
	Semivariance	Lower partial	Value at risk (VaR)			Expected shortfall (ES)		
	(SV)	Moment (LPM)	90%	95%	99%	90%	95%	99%
Panel A: Crude oil hedge ratio								
NP	1.149 (0.225)	1.202 (0.251)	0.909 (0.253)	0.994 (0.300)	1.156 (0.433)	1.206 (0.219)	1.249 (0.300)	1.254 (0.408)
SGC	1.023 (0.133)	0.998 (0.159)	1.092 (0.118)	1.053 (0.127)	0.950 (0.186)	1.031 (0.131)	1.006 (0.152)	0.955 (0.223)
SSC	1.038 (0.093)	1.006 (0.114)	1.095 (0.112)	1.063 (0.112)	0.949 (0.187)	1.044 (0.092)	1.022 (0.101)	0.981 (0.141)
SCC	−0.097 (0.067)	−0.100 (0.079)	−0.071 (0.057)	−0.099 (0.067)	−0.059 (0.143)	−0.093 (0.055)	−0.094 (0.074)	−0.131 (0.126)
C-vine	1.078 (0.148)	1.047 (0.172)	1.105 (0.177)	1.094 (0.167)	1.012 (0.255)	1.086 (0.167)	1.069 (0.181)	1.005 (0.228)
D-vine	1.087 (0.113)	1.063 (0.128)	1.109 (0.157)	1.099 (0.145)	1.042 (0.219)	1.093 (0.134)	1.082 (0.138)	1.044 (0.184)
Panel B: Gasoline hedge ratio								
NP	1.155 (0.100)	1.190 (0.125)	0.970 (0.171)	1.112 (0.283)	1.054 (0.340)	1.171 (0.172)	1.196 (0.164)	1.231 (0.222)
SGC	1.077 (0.087)	1.078 (0.086)	1.063 (0.094)	1.055 (0.116)	1.054 (0.097)	1.052 (0.096)	1.049 (0.102)	1.067 (0.101)
SSC	1.106 (0.078)	1.092 (0.085)	1.088 (0.096)	1.090 (0.105)	1.083 (0.117)	1.081 (0.084)	1.081 (0.086)	1.061 (0.094)
SCC	0.053 (0.064)	0.053 (0.080)	0.020 (0.075)	0.052 (0.077)	−0.005 (0.181)	0.037 (0.075)	0.030 (0.106)	−0.011 (0.167)
C-vine	1.094 (0.100)	1.063 (0.105)	1.085 (0.130)	1.103 (0.117)	1.043 (0.128)	1.078 (0.108)	1.063 (0.110)	0.970 (0.150)
D-vine	1.138 (0.108)	1.131 (0.117)	1.116 (0.131)	1.140 (0.124)	1.135 (0.145)	1.126 (0.118)	1.124 (0.120)	1.085 (0.165)
Panel C: Heating oil hedge ratio								
NP	1.053 (0.585)	1.022 (0.738)	0.856 (0.601)	0.803 (0.650)	1.265 (0.815)	1.185 (0.549)	1.185 (0.651)	0.896 (1.023)
SGC	1.024 (0.301)	1.031 (0.340)	1.063 (0.264)	1.010 (0.332)	0.975 (0.405)	1.067 (0.300)	1.054 (0.327)	1.038 (0.430)
SSC	1.010 (0.249)	1.025 (0.293)	1.035 (0.247)	0.999 (0.275)	0.863 (0.418)	1.040 (0.239)	1.037 (0.253)	1.069 (0.336)
SCC	0.254 (0.141)	0.243 (0.152)	0.318 (0.181)	0.210 (0.161)	0.263 (0.259)	0.253 (0.138)	0.231 (0.139)	0.170 (0.173)
C-vine	1.156 (0.326)	1.216 (0.387)	1.081 (0.320)	1.095 (0.355)	1.152 (0.496)	1.188 (0.339)	1.224 (0.367)	1.339 (0.472)
D-vine	1.086 (0.244)	1.046 (0.298)	1.093 (0.302)	1.050 (0.295)	0.947 (0.439)	1.097 (0.265)	1.082 (0.286)	0.994 (0.393)

Notes: The optimal hedge ratios for different hedging objectives are estimated using a rolling window approach with a window of approximately 5 years or 261 weeks. The total number of estimation windows is 1123 windows. NP is the nonparametric method. SGC is the standard Gaussian copula model. SSC is the standard Student's t copula model. SCC is the standard Clayton copula model. C-vine is the canonical vine copula model. D-vine is the drawable vine copula model.

Table 4
Out-of-sample hedging effectiveness of different hedging models and hedging objectives.

Model	Hedging objective							
	SV	LPM	VaR reduction (%)			ES reduction (%)		
	Reduction (%)	Reduction (%))	90%	95%	99%	90%	95%	99%
NP	58.59 (61.91)	65.53 (72.66)	34.87 (35.08)	37.56 (37.90)	31.34 (32.70)	36.88 (38.84)	34.00 (35.29)	24.59 (22.61)
SGC	60.02 (61.69)	68.22 (76.34)	**41.61 (41.67)**	40.30 (40.90)	32.77 (36.98)	36.83 (36.81)	33.91 (34.52)	27.19 (29.91)
SSC	60.66 (63.36)	69.29 (76.07)	41.33 (41.35)	41.14 (41.93)	33.23 **(37.07)**	37.36 (37.64)	34.77 (36.21)	28.53 **(30.36)**
SCC	−3.95 (1.59)	−6.04 (1.63)	−4.01 (−2.42)	1.36 (1.58)	−0.98 (−0.07)	−0.35 (0.09)	−0.46 (−0.31)	−4.71 (−2.07)
C-vine	60.44 (63.87)	69.43 (75.46)	40.01 (39.97)	41.12 (42.66)	33.62 (35.14)	37.30 (38.57)	34.96 (35.67)	28.40 (30.76)
D-vine	**61.04 (64.39)**	**70.32 (76.66)**	39.87 (39.69)	**42.42 (43.51)**	**35.73** (37.02)	**38.14 (39.07)**	**36.08 (36.99)**	**30.52** (30.01)

Notes: The table reports the mean (median) out-of-sample hedging effectiveness for different hedging methods and hedging objectives. The mean and median hedging effectiveness are calculated across 1123 out-of-sample test windows. The best performing hedging method for each hedging objective is highlighted in bold type. NP is the nonparametric method. SGC is the standard Gaussian copula model. SSC is the standard Student's t copula model. SCC is the standard Clayton copula model. C-vine is the canonical vine copula model. D-vine is the drawable vine copula model. SV denotes Semivariance; LPM denotes Lower Partial Moment; VaR denotes Value at Risk; and ES denotes Expected Shortfall.

4.3. Out-of-sample hedging effectiveness

Table 4 presents the out-of-sample hedging effectiveness of the minimum-SV, minimum-LPM, minimum-VaR, and minimum-ES objectives for the six hedging models – the NP, SGC, SSC, SCC, C-vine copula, and D-vine copula models. For each hedging objective and model, the table gives the mean and median percentage reductions in the respective downside risk of the hedged P&Ls relative to the unhedged P&Ls. The mean and median values are calculated across the 1123 out-of-sample test windows. The best performing hedging model for each hedging objective is highlighted in bold type. Also, a paired *t*-test is performed to test the null hypothesis of equal out-of-sample hedging effectiveness between two hedging models.[22] The test results are reported in Tables A.1–A.4 in the Appendix.

4.3.1. Minimum-SV objective

Considering first the minimum-SV objective, all models (except the SCC model) produce, on average, at least 58% SV reductions. The D-vine

copula model is the most effective model, with a mean (median) SV reduction of 61.04% (64.39%). The SCC model performs extremely poorly with the mean SV reduction of −3.95% (i.e., increasing risk) and median SV reduction of 1.59%. Recall from Table 3, the SCC model recommends the gasoline and heating oil hedge ratios fairly close to 0, and thus fails to protect against adverse price movements in the gasoline and heating oil markets. In addition, it supports a speculative position in the crude oil futures market (i.e., the crude oil hedge ratios being <0), which could end up adding more risk to the unhedged position. In particular, this disappointing performance may be explained by the very poor fit of the SCC model (see Table 2).

Comparing with the widely used NP method, the D-vine copula model leads to a larger mean (median) SV reduction of about 2.45% (2.48%) points. It is evident from Fig. 3 (upper panel) that the D-vine copula model is superior to the NP method for most out-of-sample test windows (more specifically, about 66.61% of the cases). The paired *t*-test results also confirm that the mean hedging effectiveness of the D-vine copula model is significantly higher than that of the NP method (see Table A.1). The maximum improvement of the D-vine copula model over the NP model is 26.48% points for the March-2007-to-September-2009 test window, which covers the period of extreme fluctuations in crude oil prices.

---

[22] The test statistics are calculated using heteroskedasticity and autocorrelation consistent (HAC) standard errors.
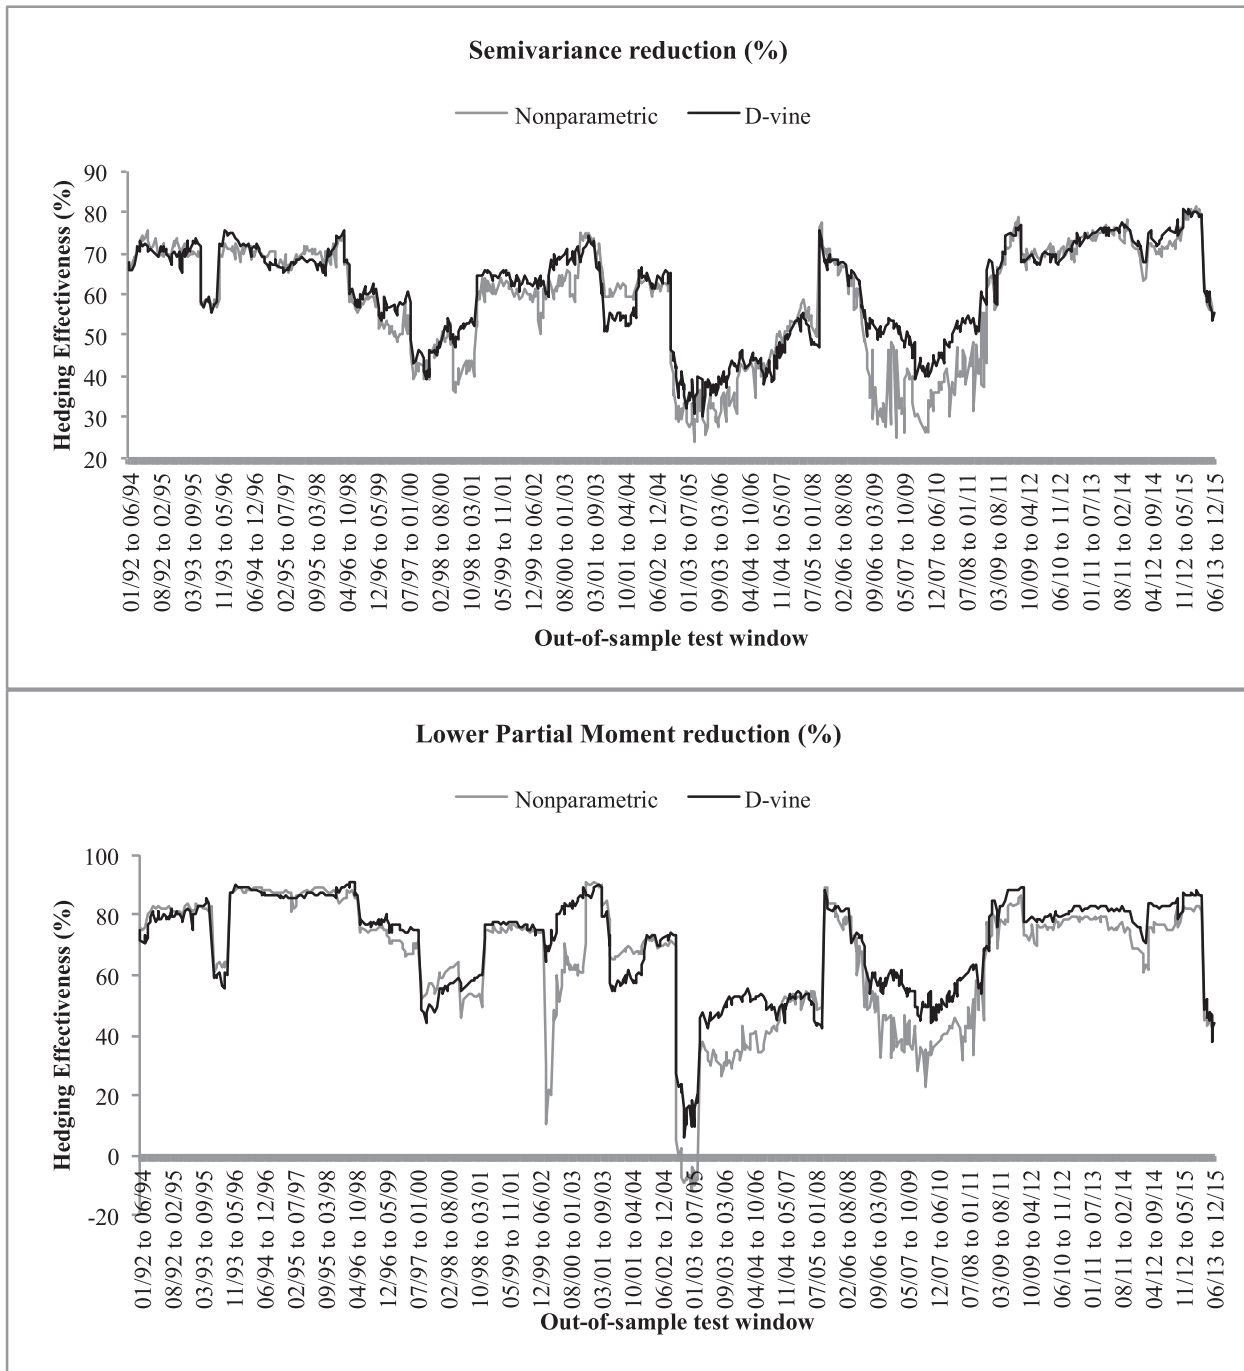
Fig. 3. Out-of-sample hedging effectiveness: percentage reductions in Semivariance (upper panel) and Lower Partial Moment (lower panel).

In fact, the D-vine copula model performs much better than the NP method for most test windows covering the years 2007 to 2010. On the other hand, the greatest improvement of the NP method over the D-vine copula model is only 8.99% points for the October-2001-to-April-2004 test window. Nevertheless, this suggests that the NP method may outperform the D-vine copula model when prices are relatively stable.

The D-vine copula model also produces better outcomes than the other copula models both in mean and median terms. The D-vine copula model clearly outperforms the SCC model. The mean (median) improvement of the D-vine copula model over the SGC, SSC and C-vine copula models ranges between 0.38% (0.52%) point and 1.02% (2.70%) points.[23]

Overall, under the minimum-SV framework, the D-vine copula model is on average able to significantly improve upon all the other models (see Table A.1). However, except for the case of the SCC model, the mean and median improvement offered by the D-vine copula model is only moderate. This is not totally unexpected because these models recommend the hedge ratios of fairly similar magnitude (see Table 3).

4.3.2. Minimum-LPM objective

We next consider the minimum-LPM objective. Similar to the minimum-SV objective, the D-vine copula model performs better than the other models both in mean and median terms (Table 4). Also, Table A.2 shows that the D-vine copula model has a significantly higher mean out-of-sample hedging effectiveness than all the other models. In particular, the D-vine copula model leads to a mean (median) LPM

---

[23] Results for each out-of-sample test window are not reported here but available upon request.

reduction of 70.32% (76.66%). Again, the SCC model performs the worst with the mean (median) hedging effectiveness of −6.04% (1.63%), confirming that using only one parameter is not enough to capture the dependence structure of the six-dimensional data.

Comparing with the NP method, the D-vine copula model leads to a 4.79% (4.00%) point increase in the mean (median) LPM reduction. Fig. 3 (lower panel) shows that the D-vine copula model offers higher levels of LPM reductions for most out-of-sample test windows. More specifically, the D-vine copula model produces greater LPM reductions than the NP method about 71.68% of the cases, with the maximum improvement of 59.89% points for the March-2002-to-August-2002 test window. It is also worth mentioning that the greatest improvement of the NP method over the D-vine copula model is only about 11.58% points. Besides, the NP method yields negative LPM reductions during the test windows November-2002-to-April-2005 to March-2003-to-September-2005 (a total of 19 out-of-sample test windows). Nevertheless, the mean (median) LPM improvement of the D-vine copula model over the SGC, SSC and C-vine copula models is quite modest, ranging between 0.89% (0.32%) and 2.10% (1.20%) points.

4.3.3. Minimum-VaR objective

For the minimum-VaR objective, all hedging models (except the SCC model) provide a VaR reduction of at least 30% both in mean and median terms (Table 4). Focusing on the mean hedging effectiveness, the D-vine copula model performs significantly better than all the other models at the 95% and 99% confidence levels (see Table A.3). When considering the median hedging effectiveness, the D-vine copula model performs the best only at the 95% confidence level. As expected, the SCC model performs the worst at all confidence levels. The hedging effectiveness of all hedging models, except the SCC model, is found to be lowest at the 99% confidence level, indicating that it is hard to hedge against a very extreme risk. In addition, as can be seen from Fig. 4, the hedging effectiveness for the minimum-VaR objective is relatively more volatile than for other hedging objectives. This is likely because VaR optimization is inherently more difficult than SV, LPM, and ES optimization (Gaivoronski and Pflug, 2004).

At the 90% confidence level, the mean and median reductions of VaR are greatest for the SGC model with a mean reduction of 41.61% and a median reduction of 41.67%. In particular, the SGC model leads to a higher mean (median) VaR reduction of 1.74% (1.98%) points relative to the D-vine copula model. At the 90% confidence level, the D-vine copula model also performs worse than the SSC and C-vine copula models. However, it is still able to improve upon the NP method with a larger mean (median) VaR reduction of about 5.00% (4.61%) points. It is evident from Fig. 4 (upper panel) that the D-vine copula model results in positive VaR reductions across all test windows, and clearly outperforms the NP method for most test windows. Panel A of Table A.3 also confirms that the D-vine copula model has, on average, a significantly higher out-of-sample hedging effectiveness than the NP method.

At the 95% confidence level, the D-vine copula model yields a mean (median) VaR reduction of 42.42% (43.51%), which is about 4.86% (5.61%) points higher than the NP method. Fig. 4 (middle panel) reveals that the D-vine copula model always yields positive VaR reduction, and that it offers significant improvements over the NP method in many out-of-sample test windows. Comparing with the SGC, SSC and C-vine copula models, the D-vine copula model leads to a larger mean (median) VaR reduction of at least 1.28% (0.85%) points.

At the 99% confidence level, the best performing hedging model in term of a mean VaR reduction is the D-vine copula model. On average, it offers a VaR reduction of 35.73%. The SSC model performs only slightly better than the D-vine copula model in term of a median VaR reduction (37.07% for the SSC model versus 37.02% for the D-vine copula model). Fig. 4 (lower panel) reveals that the hedging effectiveness of the D-vine copula model fluctuates greatly, and that the D-vine copula model yields negative VaR reductions for a few out-of-sample windows. Nevertheless, just as for the 90% and 95% confidence levels, the D-vine

copula model still performs better than the NP method with an increase in the mean (median) VaR reduction of 4.39% (4.32%) points. In addition, the negative VaR reductions are found in 35 test windows for the NP method but in only 7 test windows for the D-vine copula model. While statistically significant, the mean VaR improvement of the D-vine copula model over the SGC, SSC and C-vine copula models is quite modest, ranging between 2.11% and 2.96% points. However, we find that the SGC, SSC and C-vine copula models produce negative VaR reductions (i.e., increase the VaR of the unhedged position) at least 6 times more often than the D-vine copula model. Thus, the D-vine copula model is a safer choice for hedging the VaR of the refinery than the other models.

4.3.4. Minimum-ES objective

As can be seen from Table 4, in term of ES reduction, the hedging effectiveness of all hedging models is found to be largest at the lowest confidence (90% confidence level) and smallest at the largest confidence level (99% confidence level). In other words, the hedging effectiveness decreases as the confidence level increases. This indicates a greater difficulty in hedging a more extreme (tail) risk. Focusing on the mean hedging effectiveness, the D-vine copula model leads to the greatest ES reductions at all confidence levels. The paired t-test results in Table A.4 also suggest that the hedging effectiveness of the D-vine copula model is, on average, significantly higher than that of the other hedging models at all confidence levels. When we consider the median hedging effectiveness, the D-vine copula model performs the best for the 90% and 95% confidence levels, but not the 99% confidence level for which the SSC model is preferred. As before, the SCC model performs extremely poorly at all confidence levels.

As can be seen from Fig. 5, the D-vine copula model generally provides good hedging effectiveness at all the confidence levels. The mean (median) ES reductions offered by the D-vine copula model are 38.14% (39.07%), 36.08% (36.99%), and 30.52% (30.01%) for the 90%, 95%, and 99% confidence levels, respectively (Table 4). Unlike at the 90% and 95% confidence levels, it is evident from Fig. 5 (lower panel) that the D-vine copula model produces negative reductions in ES at the 99% confidence level for several out-of-sample test windows (more specifically, for a total of 29 test windows).

To find a possible reason for the occasional poor performance of the D-vine copula model, we investigate these 29 out-of-sample test windows more closely. Given the rolling window approach, these 29 test windows actually correspond to two periods of bad performance: (1) during the test windows (October-2002-to-March-2005) to (March-2003-to-September-2005) and (2) during the test windows (May-2013-to-October-2015) to (July-2013-to-December-2015). For the first period, the negative reductions in ES at the 99% confidence level are due to an additional loss to the unhedged P&Ls on March 30, 2005 (when the unhedged P&L has already fallen by 4.26 dollars per barrel). The extra loss on March 30, 2005 is particularly as a result of (1) the gasoline futures price moving in the opposite directions from the gasoline spot price, and (2) the heating oil futures price advancing more than the heating oil spot price.[24] For the second period, the negative ES reductions (at the 99% confidence level) occurs particularly because of a large magnitude of basis risk in the gasoline market on October 21, 2015, when the gasoline spot and futures prices move in the opposite directions. These two events suggest that the ability to hedge the extreme downside risk could decline considerably when the unhedged refining margin falls at the same time that the refining margin based on futures prices rises (presuming no speculation positions). In addition, it is worth noting that hedging may also increase the extreme tail risk if the refining margin based on futures prices rises (declines) more (less) than an increase (a decrease) in the unhedged refining margin. In other words, the occasional poor

---

[24] This is known as basis risk – risk that the changes in the futures prices deviate from the changes in the spot prices.
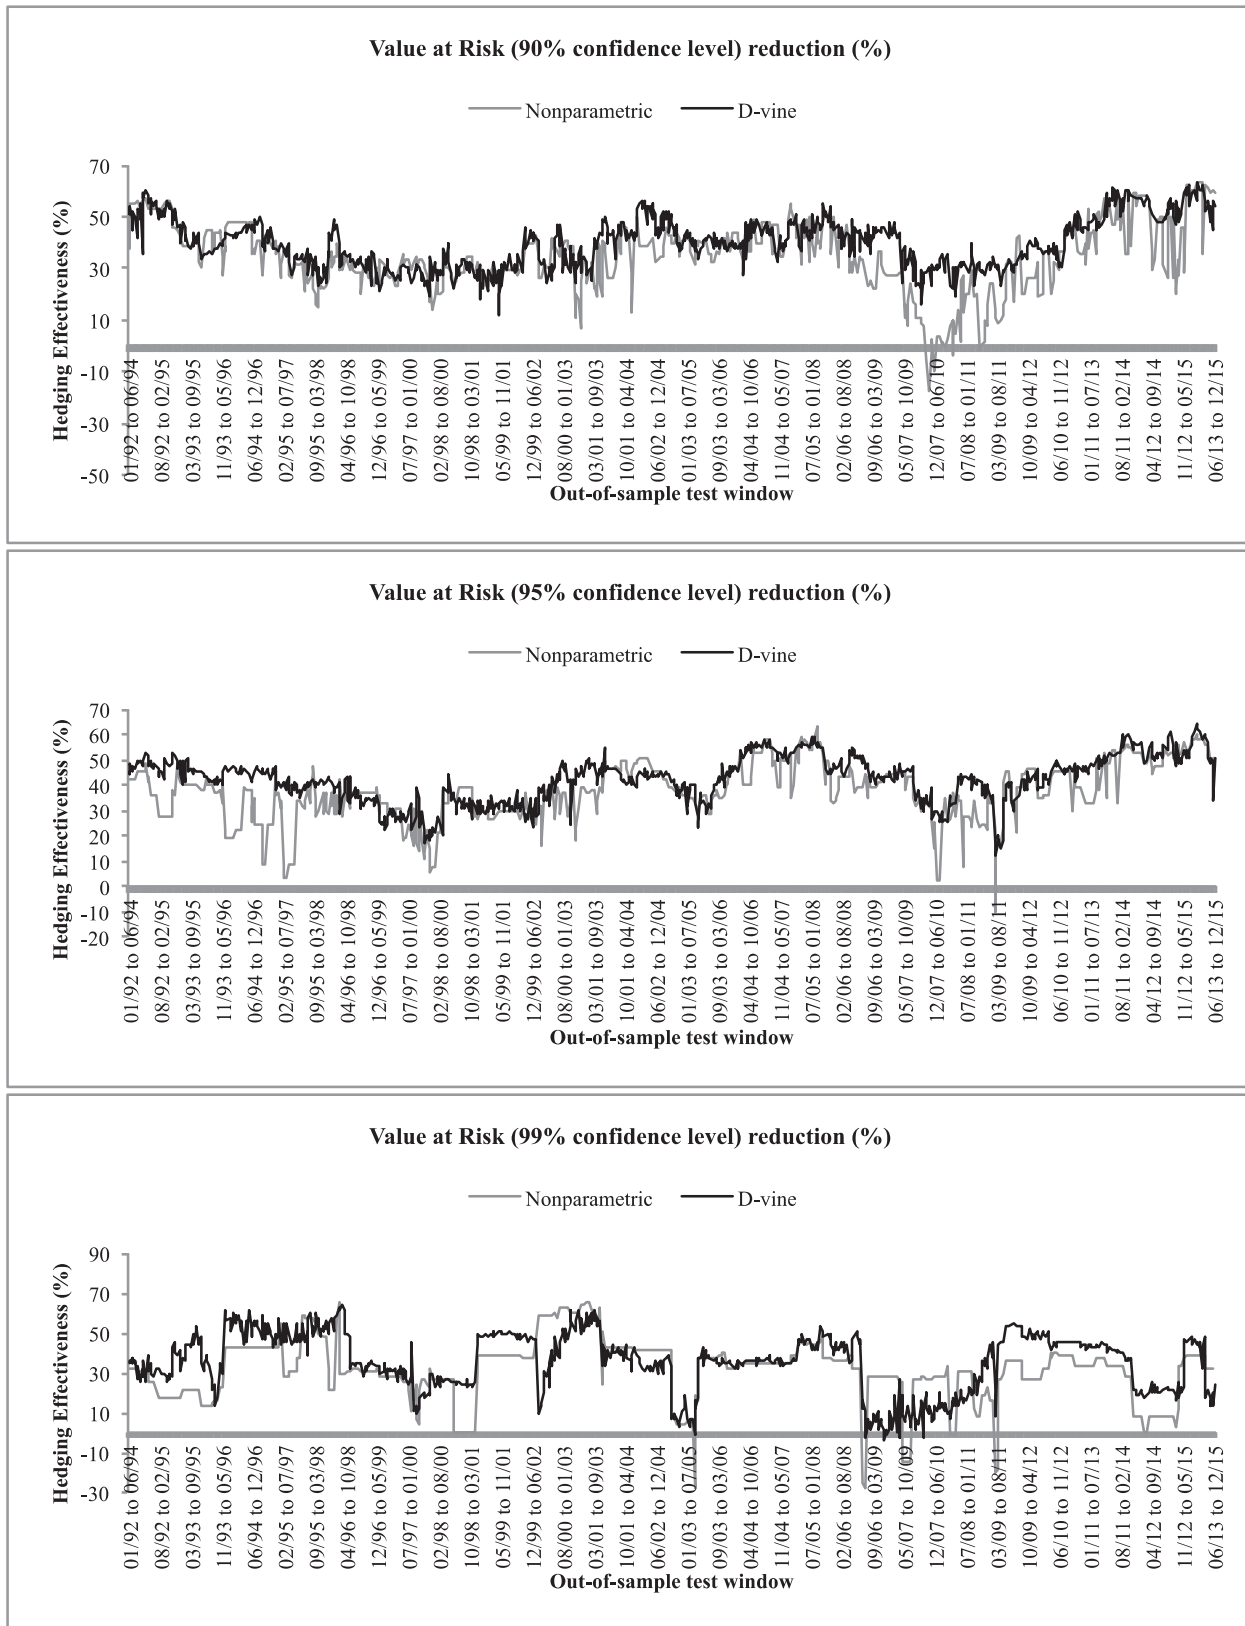
Fig. 4. Out-of-sample hedging effectiveness: percentage reductions in Value at Risk at the 90% (upper panel), 95% (middle panel), and 99% (lower panel) confidence levels.

performance of the D-vine copula model is likely explained by a sizable basis risk.

Despite its occasional poor performance, comparing with the NP method, the D-vine copula model yields larger mean (median) ES reductions of about 1.26% (0.23%), 2.08% (1.70%), and 5.93% (7.40%) points for the 90%, 95%, and 99% confidence levels, respectively. As expected, the D-vine copula model offers a larger improvement over the NP method as the confidence level becomes larger. This is because the

Fig. 5. Out-of-sample hedging effectiveness: percentage reductions in Expected Shortfall at the 90% (upper panel), 95% (middle panel), and 99% (lower panel) confidence levels.

NP method is based on the empirical distribution of the price changes and is therefore likely to provide poor estimates of the very extreme quantiles of the distribution (Pritsker, 2006). It is also evident from Fig. 5 that the D-vine copula model outperforms the NP method for most out-of-sample test windows at all confidence levels. In addition, the NP method produces negative ES reduction at the 99% confidence level much more often than the D-vine copula model (128 versus 29 test windows).

The D-vine copula model is also preferred to the SGC, SSC, and C-vine copula models at all confidence levels, except at the 99% confidence level when the median hedging effectiveness is considered. In this case, the SSC model produces a slightly greater reduction in the ES of about 0.35% point. Overall, the mean ES improvement of the D-vine copula model over these models is quite modest (though statistically significant). Nevertheless, these models produce poor hedging performance much more often than the D-vine copula model. Specifically, at the 99% confidence level, the negative ES reductions are found in 95, 52, and 43 test windows for the SGC, SSC, and C-vine copula models, respectively. Given this result, the D-vine copula model seems to be a better and safer choice than the other hedging models in managing the ES of the refinery.

4.4. Out-of-sample hedging performance of different hedging objectives across various measures of hedging effectiveness

Given the range of alternative minimum-downside risk hedging objectives available to refineries, it would be interesting to examine their performance across various measures of hedging effectiveness. In this section, we evaluate the out-of-sample hedging effectiveness of the different hedging objectives – minimum-SV, minimum-LPM, minimum-VaR (90%, 95%, and 99%), and minimum-ES (90%, 95%, and 99%) – using a variety of downside risk measures, including SV, LPM, VaR (90%, 95%, and 99%), and ES (90%, 95%, and 99%). Table 5 presents the mean hedging performance of the eight hedging objectives across the eight hedging effectiveness measures for the case of D-vine copula model, which is the best performing hedging model. The mean hedging effectiveness is calculated across the 1123 out-of-sample test windows. For each measure of hedging effectiveness, rankings of the eight hedging objectives are provided in the parentheses next to the mean hedging effectiveness. The best performing hedging objective for each hedging effectiveness measure is also highlighted in bold type.

Considering first the SV hedging effectiveness, 58%–61% of the SV of the unhedged position is removed for all hedging objectives. As expected, the minimum-SV objective performs best in term of reducing the SV of the unhedged P&Ls. In particular, it offers a mean SV reduction of 61.04%. The next-best hedging objective is the minimum-ES (90%) objective with the SV hedging effectiveness of 60.77%. It is worth noting that the two worst hedging objectives in this case are the minimum-VaR (99%) and minimum-ES (99%) objectives.

For the LPM hedging effectiveness, all hedging objectives can remove at least 68% of the LPM of the unhedged P&Ls. Unlike the previous case, we find that the minimum-LPM objective does not lead to greatest mean reduction of LPM out-of-sample. Instead, the best and the next-best hedging objectives are the minimum-ES (95%) and minimum-SV objectives. This is not totally unexpected because the hedge ratios for each hedging objective are derived using the in-sample data while the hedging effectiveness is calculated based on the out-of-sample data.

Nevertheless, similar to the case of the SV hedging effectiveness, we find that the minimum-VaR (99%) and minimum-ES (99%) objectives are the worst performing hedging objectives from the point of view of LPM reduction.

We next consider the VaR hedging effectiveness of the different hedging objectives. Depending on both hedging objectives and confidence levels, the VaR reduction ranges between 35% and 43%. At the 90% confidence level, we find that the minimum-VaR (90%) objective provides the greatest VaR (90%) reduction, followed by the minimum-SV objective. At the 95% confidence level, the two most effective hedging objectives are the minimum-SV and minimum-VaR (95%) objectives. At the 99% confidence level, the best and next-best hedging objectives are the minimum-ES (90%) and minimum-ES (95%) objectives. We also find that the minimum-VaR (99%) and minimum-ES (99%) objectives perform worst in reducing the VaR of the unhedged position at all confidence levels studied.

Finally, we examine the ES hedging effectiveness for each of the hedging objectives. Again, depending on hedging objectives and confidence levels, the ES reduction achieved by the various hedging objectives ranges between 30% and 38%. For both ES (90%) and ES (95%) hedging effectiveness, we find that the minimum-SV objective offers the greatest ES reduction, followed by the minimum-ES (90%) objective. Also, at both confidence levels, the minimum-VaR (99%) and minimum-ES (99%) objectives perform worst in reducing the ES of the unhedged P&Ls. For the ES (99%) objective, the best and next-best hedging objectives are the minimum-VaR (90%) and minimum-ES (90%) objectives, whereas the worst performing hedging objectives are the minimum-LPM and minimum-ES (99%) objectives.

Comparing the different hedging objectives across the various hedging effectiveness measures, the minimum-SV objective is found to be either the best or the next-best hedging objectives across the majority of hedging effectiveness. In addition, except for the case of the VaR (90%) hedging effectiveness, we find that the minimum-SV objective is able to reduce the downside risk by at least the same amount as the associated minimum-downside risk objective. For example, the minimum-SV objective is found to perform better than the minimum-LPM objective in reducing the LPM of the hedged position (70.45% versus 70.32% LPM reduction). These findings suggest that, regardless of the measures of downside risk the refineries wishing to reduce, the minimum-SV objective seems to be the best choice among the range of alternative hedging objectives available to them. Further, the minimum-VaR (99%) and minimum-ES (99%) objectives are found to perform relatively poorer than the other hedging objectives across the measures of hedging effectiveness considered. The relatively poor hedging performance is likely explained by the relatively small amount of data available for calculating the minimum-VaR (99%) and minimum-ES (99%) hedge ratios. These results suggest that, when choosing the hedging objectives, it is necessary for the refineries to account for the fact that the minimum-VaR and minimum-ES hedge ratios at the high confidence level (99%)

Table 5
Out-of-sample hedging performance of the different hedging objectives across various hedging effectiveness measures (D-vine copula model).

Hedging effectiveness	Hedging objective							
	Min-SV	Min-LPM	Min-VaR (90%)	Min-VaR (95%)	Min-VaR (99%)	Min-ES (90%)	Min-ES (95%)	Min-ES (99%)
SV reduction	**61.04% (1)**	60.73% (3)	60.04% (6)	60.66% (4)	58.69% (7)	60.77% (2)	60.65% (5)	58.63% (8)
LPM reduction	70.45% (2)	70.32% (4)	69.40% (6)	70.14% (5)	69.21% (7)	70.36% (3)	**70.58% (1)**	68.99% (8)
VaR (90%) reduction	39.76% (2)	38.54% (6)	**39.87% (1)**	39.53% (4)	36.18% (8)	39.53% (3)	38.73% (5)	36.40% (7)
VaR (95%) reduction	**42.72% (1)**	41.99% (4)	41.95% (5)	42.42% (2)	39.95% (7)	42.29% (3)	41.83% (6)	39.57% (8)
VaR (99%) reduction	36.55% (4)	36.01% (6)	36.57% (3)	36.51% (5)	35.73% (7)	**36.79% (1)**	36.72% (2)	35.22% (8)
ES (90%) reduction	**38.32% (1)**	37.81% (6)	37.83% (5)	38.03% (3)	36.30% (7)	38.14% (2)	37.84% (4)	36.15% (8)
ES (95%) reduction	**36.29% (1)**	35.88% (6)	35.93% (5)	36.13% (3)	34.77% (7)	36.25% (2)	36.08% (4)	34.67% (8)
ES (99%) reduction	31.13% (5)	30.66% (7)	**31.66% (1)**	31.43% (3)	31.04% (6)	31.50% (2)	31.43% (4)	30.52% (8)

Notes: The table reports the mean out-of-sample hedging performance of the eight hedging objectives across eight hedging effectiveness measures. The mean hedging effectiveness is calculated across 1123 out-of-sample test windows. The D-vine copula model is used to generate the minimum-SV, minimum-LPM, minimum-VaR, and minimum-ES hedge ratios. For each hedging effectiveness measure, rankings of the eight hedging objectives are reported in the parentheses next to the mean hedging effectiveness. The best performing hedging objective for each hedging effectiveness measure is also highlighted in bold type.

are associated with higher uncertainty and therefore more likely to perform worse than the other minimum-downside risk hedge ratios.

5. Conclusion

Oil refineries face the risk of losses that are associated with an increase in input prices (crude oil prices), a decrease in output prices (gasoline and/or heating oil prices), or a combination of both. In other words, they are exposed to downside price risk in multiple petroleum markets (including crude oil, gasoline and heating oil markets). The refineries may hedge against the risks of adverse input and output price movements using crude oil, gasoline, and heating oil futures. This paper proposes a multiproduct futures hedging model that minimizes the downside risk of the oil refineries, measured by Semivariance (SV), Lower Partial Moment (LPM), Value at Risk (VaR), or Expected Shortfall (ES). This is of special interest for the refineries that are particularly concerned about the negative impacts of unfavorable price movements in multiple petroleum markets.

The empirical analysis is based on a stylized problem of a typical U.S. oil refinery that converts 3 barrels of crude oil to 2 barrels of gasoline and 1 barrel of heating oil. The proposed hedging model constructs a joint distribution of six variables (spot and futures price changes in crude oil, gasoline, and heating oil markets) using a vine copula methodology, and determines the minimum-downside risk hedge ratios using a Monte Carlo optimization technique. The vine copula methodology, which is a relatively new class of multivariate copula approaches, is chosen because it allows us to capture important characteristics of petroleum price changes, including skewness and fat-tailedness in the marginal distributions of individual price change series as well as heterogeneous (tail) dependence patterns between different pairs of price changes. In this paper, two popular classes of vine copulas – the canonical (C-) and drawable (D-) vine copulas – are considered in the modeling of the dependence structure in petroleum spot and futures markets. We evaluate the suitability of the C– and D-vine copula models by examining their hedging effectiveness over 1123 out-of-sample test windows. In addition, we compare the out-of-sample hedging effectiveness of the vine copula models to that of several common alternative approaches, including the nonparametric (NP), standard Gaussian copula (SGC), standard Student's t (SSC), and standard Clayton copula (SCC) models.

The main findings are as follows. First, on average we find that both C– and D-vine copula models are able to effectively reduce the downside risk of the refinery, and that the D-vine copula model provides better out-of-sample hedging effectiveness than the C-vine copula model. The results are consistent across all the hedging objectives considered (namely, the minimum-SV, minimum-LPM, minimum-VaR, and minimum-ES objectives). The superiority of the D-vine copula model may be explained by its ability to directly capture the high dependence between the spot and its corresponding futures price changes in all petroleum markets, which is a feature that is not allowed by the C-vine copula model. Depending on the hedging objective, the mean (median) downside risk reductions offered by the D-vine copula model between 30.52% (30.01%) to 70.32% (76.66%). Second, for the minimum-VaR (99% confidence level) and minimum-ES (99% confidence level) objectives, the D-vine copula model yields negative downside risk reduction (that is, increases downside risk of the unhedged position) for few out-of-sample test windows (more specifically, for <30 out of 1123 test windows). We find that the occasional poor performance of the D-vine copula model is likely due to a sizable basis risk (or the risk that futures prices do not move in line with the underlying spot prices). However, the D-vine copula model produces poor hedging performance much less often than the other alternative hedging models.

Third, the D-vine copula model is on average preferred to the widely used NP method regardless of which hedging objective is considered. The superiority of the D-vine copula model over the NP method is generally seen across numerous out-of-sample test windows, which signals

the relevance of explicit modeling of the extreme price dependence. Finally, the D-vine copula model on average leads to greater downside-risk reductions than the SGC, SSC, SCC, and C-vine copula models. As expected, the improvement over the SCC model, which uses only one parameter to capture the dependence structure of six variables, is enormous. However, the improvement over the SGC, SSC, and C-vine copula models is quite modest. Nevertheless, we find that these models (as well as the NP method) produce poor hedging performance for a much greater number of out-of-sample test windows than the D-vine copula model. Given these results, the D-vine copula model seems to be a good and safe hedging model for the refinery that wants to minimize its downside risk. Moreover, comparing the performance of the different hedging objectives across the various hedging effectiveness, we find that the minimum-SV objective is the best choice among the range of alternative hedging objectives available to the refineries.

As indicated above, our analysis might be especially useful for petroleum (as well as non-petroleum) producers who seek to reduce the risks of adverse price movements in input and output markets. In addition, the findings reported in this paper provide additional evidence that there is a benefit from modeling the joint distribution (more specifically, the dependence structure) more realistically.

Acknowledgements

Appendix A. Supplementary data

Supplementary data to this article can be found online at http://dx.doi.org/10.1016/j.eneco.2017.07.012.

References

Aas, K., Czado, C., Frigessi, A., Bakken, H., 2009. Pair-copula constructions of multiple dependence. Insur. Math. Econ. 44 (2), 182–198.

Alexander, C., Prokopczuk, M., Sumawong, A., 2013. The (de)merits of minimum-variance hedging: application to the crack spread. Energy Econ. 36, 698–707.

Allen, D.E., Ashraf, M.A., McAleer, M., Powell, R.J., Singh, A.K., 2013. Financial dependence analysis: applications of vine copulas. Statistica Neerlandica 67 (4), 403–435.

Aloui, R., Aïssa, M.S.B., Hammoudeh, S., Nguyen, D.K., 2014. Dependence and extreme dependence of crude oil and natural gas prices with applications to risk management. Energy Econ. 42, 332–342.

Awudu, I., Wilson, W., Dahl, B., 2016. Hedging strategy for ethanol processing with copula distributions. Energy Econ. 57, 59–65.

Barbi, M., Romagnoli, S., 2014. A copula-based quantile risk measure approach to estimate the optimal hedge ratio. J. Futur. Mark. 34 (7), 658–675.

Bedford, T., Cooke, R.M., 2001. Probability density decomposition for conditionally dependent random variables modeled by vines. Ann. Math. Artif. Intell. 32 (1), 245–268.

Bekiros, S., Hernandez, J.A., Hammoudeh, S., Nguyen, D.K., 2015. Multivariate dependence risk and portfolio optimization: an application to mining stock portfolios. Resour. Policy 46, 1–11.

Bouyé, E., Salmon, M., 2009. Dynamic copula quantile regressions and tail area dynamic dependence in forex markets. Eur. J. Financ. 15 (7–8), 721–750.

Brechmann, E., Czado, C., Paterlini, S., 2014. Flexible dependence modeling of operational risk losses and its impact on total capital requirements. J. Bank. Financ. 40, 271–285.

Brechmann, E.C., Schepsmeier, U., 2013. Modeling dependence with C- and D-vine copulas: the R package CDVine. J. Stat. Softw. 52 (3), 1–27.

Cao, Z., Harris, R.D.F., Shen, J., 2010. Hedging and value at risk: a semi-parametric approach. J. Futur. Mark. 30 (8), 780–794.

Carchano, Óscar, Pardo, Ángel, 2009. Rolling over stock index futures contracts. J. Futur. Mark. 29 (7), 684–694.

Chang, C.-L., McAleer, M., Tansuchat, R., 2010. Analyzing and forecasting volatility spillovers, asymmetries and hedging in major oil markets. Energy Econ. 32 (6), 1445–1455.

Charpentier, A., Fermanian, J.-D., O. S., 2007. The estimation of copulas: theory and practice. In: Risk, J. (Ed.), Copulas: From Theory to Application in Finance. Risk Books, London, pp. 35–60.

Chen, S., Wilson, W., Larsen, R., Dahl, B., 2016. Risk management for grain processors and "copulas". Can. J. Agric. Econ. 64, 365–382.

Conlon, T., Cotter, J., 2012. An empirical analysis of dynamic multiscale hedging using wavelet decomposition. J. Futur. Mark. 32 (3), 272–299.

Conlon, T., Cotter, J., 2013. Downside risk and the energy Hedger's horizon. Energy Econ. 36, 371–379.

Czado, C., 2010. Pair-copula constructions of multivariate copulas. In: Jaworski, P., Durante, F., Härdle, W., Rychlik, T. (Eds.), Copula Theory and Its Applications. Springer, New York, pp. 93–109.

Czado, C., Schepsmeier, U., Min, A., 2012. Maximum likelihood estimation of mixed C-vines with application to exchange rates. Stat. Model. 12 (3), 229–255.

Demirer, R., Lien, D., 2003. Downside risk for short and long hedgers. Int. Rev. Econ. Financ. 12 (1), 25–44.

Dißmann, J., Brechmann, E.C., Czado, C., Kurowicka, D., 2013. Selecting and estimating regular vine copulae and application to financial returns. Comput. Stat. Data Anal. 59, 52–69.

Ederington, L.H., 1979. The hedging performance of the new futures markets. J. Financ. 34 (1), 157–170.

Fishburn, P.C., 1977. Mean-risk analysis with risk associated with below-target returns. Am. Econ. Rev. 67 (2), 116–126.

Gaivoronski, A., Pflug, G., 2004. Value-at-risk in portfolio optimization: properties and computational approach. J. Risk 7 (2), 1–31.

Genest, C., Kilani, G., Rivest, L.-P., 1995. A semiparametric estimation procedure of dependence parameters in multivariate families of distributions. Biometrika 82 (3), 543–552.

Grégoire, V., Genest, C., Gendron, M., 2008. Using copulas to model price dependence in energy markets. Energy Risk 5 (5), 58–64.

Haigh, M.S., Holt, M.T., 2002. Crack spread hedging: accounting for time-varying volatility spillovers in the energy futures markets. J. Appl. Econ. 17 (3), 269–289.

Hammoudeh, S., Li, H., Jeon, B., 2003. Causality and volatility spillovers among petroleum prices of Wti, gasoline and heating oil in different locations. N. Am. J. Econ. Financ. 14 (1), 89–114.

Harlow, W.V., 1991. Asset allocation in a downside-risk framework. Financ. Anal. J. 47 (5), 28–40.

Harris, R.D.F., Shen, J., 2006. Hedging and Value at Risk. J. Futur. Mark. 26 (4), 369–390.

Hsu, C.-C., Tseng, C.-P., Wang, Y.-H., 2008. Dynamic hedging with futures: a copula-based GARCH model. J. Futur. Mark. 28 (11), 1095–1116.

Ji, Q., Fan, Y., 2011. A dynamic hedging approach for refineries in multiproduct oil markets. Energy 36 (2), 881–887.

Joe, H., 1996. Families of M-Variate Distributions with Given Margins and M (M-1)/2 Bivariate Dependence Parameters. Lecture Notes-Monograph Series 28 pp. 120–141.

Joe, H., 1997. Multivariate Models and Multivariate Dependence Concepts. CRC Press, New York.

Kurowicka, D., Cooke, R.M., 2005. Distribution-free continuous Bayesian belief nets. In: Wilson, A., Limnios, N., Keller-McNulty, S., Armijo, Y. (Eds.), Modern Statistical and Mathematical Methods in Reliability. World Scientific, Singapore, pp. 309–323.

Kurowicka, D., Cooke, R.M., 2006. Uncertainty Analysis with High Dimensional Dependence Modelling. John Wiley & Sons, Sussex.

Lee, H.-T., 2009. A copula-based regime-switching GARCH model for optimal futures hedging. J. Futur. Mark. 29 (10), 946–972.

Lien, D., Tse, Y.K., 1998. Hedging time-varying downside risk. J. Futur. Mark. 18 (6), 705–722.

Lien, D., Tse, Y.K., 2000. Hedging downside risk with futures contracts. Appl. Financ. Econ. 10 (2), 163–170.

Liu, P., Vedenov, D., Power, G.J., 2017. Is hedging the crack spread no longer all It's cracked up to be? Energy Econ. 63, 31–40.

Low, R.K.Y., Alcock, J., Faff, R., Brailsford, T., 2013. Canonical vine copulas in the context of modern portfolio management: are they worth it? J. Bank. Financ. 37 (8), 3085–3099.

McNeil, A.J., Frey, R., 2000. Estimation of tail-related risk measures for heteroscedastic financial time series: an extreme value approach. J. Empir. Financ. 7 (3–4), 271–300.

Nelder, J.A., Mead, R., 1965. A simplex method for function minimization. Comput. J. 7 (4), 308–313.

Nelsen, R.B., 2006. An Introduction to Copulas. 2nd ed. Springer, New York.

Power, G.J., Vedenov, D., 2010. Dealing with downside risk in a multi-commodity setting: a case for a "Texas hedge"? J. Futur. Mark. 30 (3), 290–304.

Pritsker, M., 2006. The hidden dangers of historical simulation. J. Bank. Financ. 30 (2), 561–582.

Riccetti, L., 2013. A copula–GARCH model for macro asset allocation of a portfolio with commodities. Empir. Econ. 44 (3), 1315–1336.

Rockafellar, R.T., Uryasev, S., 2002. Conditional value-at-risk for general loss distributions. J. Bank. Financ. 26 (7), 1443–1471.

Roy, A.D., 1952. Safety first and the holding of assets. Econometrica 20 (3), 431–449.

Serra, T., Gil, J.M., 2012. Biodiesel as a motor fuel price stabilization mechanism. Energy Policy 50, 689–698.

Sklar, A., 1959. Fonctions de Répartition à n Dimensions et Leurs Marges. Publications de l' Institut Statistique de l'Université de Paris 8 pp. 229–231.

Sukcharoen, K., Choi, H., Leatham, D.J., 2015. Optimal gasoline hedging strategies using futures contracts and exchange-traded funds. Appl. Econ. 47 (32), 3482–3498.

Unser, M., 2000. Lower partial moments as measures of perceived risk: an experimental study. J. Econ. Psychol. 21 (3), 253–280.

Veld, C., Veld-Merkoulova, Y.V., 2008. The risk perceptions of individual investors. J. Econ. Psychol. 29 (2), 226–252.

Weiß, G.N.F., Supper, H., 2013. Forecasting liquidity-adjusted intraday value-at-risk with vine copulas. J. Bank. Financ. 37 (9), 3334–3350.

Zhang, D., 2014. Vine copulas and applications to the European Union sovereign debt analysis. Int. Rev. Financ. Anal. 36, 46–56.

Zhang, B., Wei, Y., Yu, J., Lai, X., Peng, Z., 2014. Forecasting var and ES of stock index portfolio: a vine copula method. Phys. A Stat. Mech. Appl. 416, 112–124.

Zimmer, D.M., 2015. Analyzing comovements in housing prices using vine copulas. Econ. Inq. 53 (2), 1156–1169.