

Tilburg University

## Multivariate Option Pricing Using Dynamic Copula Models

van den Goorbergh, R.W.J.; Genest, C.; Werker, B.J.M.

*Publication date:*  
2003

[Link to publication in Tilburg University Research Portal](#)

*Citation for published version (APA):*

van den Goorbergh, R. W. J., Genest, C., & Werker, B. J. M. (2003). *Multivariate Option Pricing Using Dynamic Copula Models*. (CentER Discussion Paper; Vol. 2003-122). Finance.

### General rights

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the public portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the public portal

### Take down policy

If you believe that this document breaches copyright please contact us providing details, and we will remove access to the work immediately and investigate your claim.

Center



Discussion Paper

No. 2003–122

**MULTIVARIATE OPTION PRICING USING DYNAMIC COPULA  
MODELS**

By R.W.J. van den Goorbergh, C. Genest, B.J.M. Werker

December 2003

ISSN 0924-7815

# Multivariate Option Pricing Using Dynamic Copula Models

Rob W. J. van den Goorbergh\*      Christian Genest<sup>†</sup>  
Bas J. M. Werker<sup>‡</sup>

This version: September 30, 2003

## Abstract

This paper examines the behavior of multivariate option prices in the presence of association between the underlying assets. Parametric families of copulas offering various alternatives to the normal dependence structure are used to model this association, which is explicitly assumed to vary over time as a function of the volatilities of the assets. These dynamic copula models are applied to better-of-two-markets and worse-of-two-markets options on the S&P500 and Nasdaq indexes. Results show that option prices implied by dynamic copula models differ substantially from prices implied by models that fix the dependence between the underlyings, particularly in times of high volatilities. Furthermore, the normal copula produces option prices that differ significantly from non-normal copula prices, irrespective of initial volatility levels. Within the class of non-normal copula families considered, option prices are robust with respect to the copula choice.

*JEL classification:* G13.

*Keywords:* Time-varying dependence, best-of-two-markets options, non-normality.

---

\*Department of Finance and CentER for Economic Research, Tilburg University, PO Box 90153, 5000 LE Tilburg, The Netherlands. Corresponding author. E-mail: [r.w.j.vdngoerbergh@uvt.nl](mailto:r.w.j.vdngoerbergh@uvt.nl), Phone: +31 (0)13 466 8222, Fax: +31 (0)13 466 2875.

<sup>†</sup>Département de mathématiques et de statistique, Université Laval, Québec, Canada G1K 7P4.

<sup>‡</sup>Department of Finance, Department of Econometrics and Operations Research, and CentER for Economic Research, Tilburg University, Tilburg, The Netherlands.

# 1 Introduction

In today's economy, multivariate (or rainbow) options are viewed as excellent tools for hedging the risk of multiple assets. These options, which are written on two or more underlying securities or indexes, usually take the form of calls (or puts) that give the right to buy (or sell) the best or worst performer of a number of underlying assets. Other examples include forward contracts whose payoff is equal to that of the best or worst performer of its underlyings, and spread options on the difference between the prices of two assets.

One of the key determinants in the valuation of multivariate options is the dependence between the underlying assets. Consider for instance a bivariate call-on-max option, namely a contract that gives the holder the right to purchase the more valuable of two underlying assets for a pre-specified strike price. Intuitively, the value of such an option should be smaller if the underlyings tend to move together than when they move in opposite directions. More generally, the dependence between the underlyings could change over time. Accounting for time variation in the dependence structure between assets should prove helpful in providing a more realistic valuation of multivariate options.

Over the years, various generalizations of the Black–Scholes (1973) Brownian motion framework have been used to model multivariate option prices. Examples include Margrabe (1978), Stulz (1982), Johnson (1987), Reiner (1992), and Shimko (1994). In these papers, the dependence between assets is modelled by their correlation. However, unless asset returns are well represented by a multivariate normal distribution, correlation is often an unsatisfactory measure of dependence; see, for instance, Embrechts, McNeil, and Straumann (2002). Furthermore, it is a stylized fact of financial markets that correlations observed under ordinary market conditions differ substantially from correlations observed in hectic periods. In particular, asset prices have a greater tendency to move together in bad states of the economy than in quiet periods; see, for instance, Boyer, Gibson, and Loretan (1999) and Patton (2002a, 2002b) and references therein. These “correlation breakdowns,” associated with economic downturns, suggest a dynamic model of the dependence structure of asset returns.

In this paper, the relation between multivariate option prices and the dependence structure of the underlying financial assets is modelled dynamically through copulas. A copula is a multivariate distribution function each of whose marginals is uniform on the unit interval. It has been known since the work of Sklar (1959) that any multivariate continuous distribution function can be uniquely factored into its marginals and a copula. Thus while

correlation measures dependence through a single number, the dependence between multiple assets is fully captured by the copula. From a practical point of view, the advantage of the copula-based approach to modelling is that appropriate marginal distributions for the components of a multivariate system can be selected by any desired method, and then linked through a copula or family of copulas suitably chosen to represent the dependence prevailing between the components.

The use of copulas to price multivariate options is not new. For example, in Rosenberg (1999), univariate options data are used to estimate marginal risk-neutral densities, which are linked with a Plackett copula to obtain a bivariate risk-neutral density from which bivariate claims are valued. This semiparametric procedure uses a particular identifying assumption on the risk-neutral correlation to fix the copula parameter. Cherubini and Luciano (2002) extend Rosenberg's work by considering other families of copulas. In Rosenberg (2003), a risk-neutral bivariate distribution is estimated from nonparametric estimates of the marginal distributions and a nonparametric estimate of the copula.

An innovating feature of the present paper, however, is that, contrary to earlier works on multivariate option pricing, the dependence structure of the underlying assets is not treated as fixed, but rather as possibly varying over time. Taking into account this time variation is important because it may influence option prices. This paper proposes a model for the time variation of the dependence structure, in which a parametric copula is specified whose dependence parameter is allowed to change with the volatilities of the underlying assets. A distinct advantage of the parametric approach is that while the model may be misspecified, the robustness of the conclusions can easily be verified by repeating the analysis for as many different copula families as desired.

A similar dynamic-copula approach has already been used in the foreign exchange market literature by Patton (2002a), who found time variation to be significant in a copula model for asymmetric dependence between two exchange rates where the dependence parameter followed a ARMA-type process. While Patton's goal was to study the effect of asymmetric dependence on portfolio returns, the objective of the present paper is very different. The main focus here is on the effect of time variation in the underlying dependence structure on the price of multivariate options.

In the empirical study presented herein, multivariate options on two important American equity index returns are considered: the S&P500 and the Nasdaq. An analysis of the results suggests that allowing for time variation in the dependence structure of the underlyings produces substantially dif-

ferent option prices than under constant dependence, particularly in times of increased volatility. Moreover, option prices implied by a normal dynamic dependence structure differ significantly from option prices implied by non-normal dynamic dependence structures. These findings suggest that unless the dependence between the S&P500 and Nasdaq stock indexes is well described by a normal copula, alternative copula families should be considered. Option prices turned out to be robust among the alternative—i.e., non-normal—copula models considered in this study.

The remainder of this paper is organized as follows. Section 2 describes the payoff structure of better-of-two-markets and worse-of-two-markets claims, and explains in detail the proposed dynamic-dependence option valuation scheme. The empirical results are presented in Section 3, and conclusions are given in Section 4.

## 2 Option valuation with time-varying dependence

Multivariate options come in a wide variety of payoff schemes. The most commonly traded options of this kind are basket options on a portfolio of assets, such as index options. Other examples include spread options, some of which are traded on commodity exchanges (see, for example, Rosenberg (1998)), or dual-strike and multivariate-digital options.

This paper concentrates on European-type options on the best (worst) performer of several assets, sometimes referred to as outperformance (underperformance) options. As these are typically traded over the counter, price data are not available. Therefore, valuation models cannot be tested empirically. However, a robustness study comparing models with different assumptions remains feasible, and this is the objective pursued herein. While the study described in the sequel is restricted to options on better-and worse-of-two-markets claims, the technique is sufficiently general to analyze the aforementioned alternative multivariate options as well, and may thus be of wider interest.

One can distinguish four types of better-of-two-markets or worse-of-two-markets claims: call options on the better performer, put options on the worse performer, call options on the worse performer, and put options on the better performer. These may be referred to as call-on-max, put-on-min, call-on-min, and put-on-max options, respectively. Their payoffs at maturity are:

$$\begin{aligned}
\text{call on max} & : \max\{\max(R_1, R_2) - E, 0\}, \\
\text{put on min} & : \max\{E - \min(R_1, R_2), 0\}, \\
\text{call on min} & : \max\{\min(R_1, R_2) - E, 0\}, \\
\text{put on max} & : \max\{E - \max(R_1, R_2), 0\},
\end{aligned}$$

where  $R_i$  is the return at maturity on index  $i \in \{1, 2\}$ , and  $E$  denotes the exercise price of the option.

The proposed scheme for valuating these options is as follows. First, each of the two objective marginal distributions of the daily index returns underlying a type of claim is modelled, and their risk-neutral counterparts are derived. Next, a parametric family of copulas is chosen to fix the joint risk-neutral distribution of the index returns. The fair value of the option is then determined by taking the discounted expected value of the option's payoff under the risk-neutral distribution.

The specification chosen for the objective marginal distributions is from Duan (1995). It is general enough to capture volatility clustering, a stylized fact of equity returns for which there is overwhelming empirical evidence at the daily frequency, while still providing a relatively easy transformation to risk-neutral distributions. Each of the objective marginal distributions of the index returns is modelled by a GARCH(1,1) process. It is repeated here for the sake of completeness; see Bollerslev (1986). For  $i \in \{1, 2\}$ ,

$$\begin{aligned}
R_{i,t+1} & = \mu_i + \eta_{i,t+1}, \\
\eta_{i,t+1} | \mathcal{I}_t & \sim \mathcal{N}(0, h_{i,t}), \\
h_{i,t+1} & = \omega_i + \beta_i h_{i,t} + \alpha_i \eta_{i,t+1}^2,
\end{aligned}$$

where  $\omega_i > 0$ ,  $\beta_i > 0$ , and  $\alpha_i > 0$ . Here,  $\mathcal{I}_t$  denotes the information available at time  $t$ . However, it must be stressed that, in the light of Sklar's theorem, in principle *any* choice for the marginal distributions is consistent with the copula approach. The vast collection of alternatives that have been used by other authors to model univariate index return distributions includes (variants of) continuous-time geometric Brownian motion of Black and Scholes (1973), and the discrete-time binomial model of Cox, Ross, and Rubinstein (1979). Again, the GARCH specification that is employed here is appealing as it allows for an easy change of measure in addition to being able to capture volatility clustering. In particular, Duan (1995) shows that, under certain conditions, the change of measure comes down to a change in the drift.

An alternative, nonparametric approach is to use univariate option price data to obtain arbitrage-free estimates of the marginal risk-neutral densities,

as in Ait-Sahalia and Lo (1998). This route is taken by Rosenberg (2003). Clearly, an advantage of this approach is that it does not impose restrictions on the asset return processes or on the functional form of the risk-neutral densities. However, this flexibility comes at the cost of imprecise estimates, especially if the distributions are time-varying.

The second step in the proposed valuation scheme is to fix the joint risk-neutral distribution of the index returns by choosing a copula. A set of well-known one-parameter copula families is considered for this purpose. They are the Frank, Gumbel–Hougaard, Plackett, Galambos, and normal families. Their cumulative distribution functions are given in Appendix A. For all of these copulas, there is a one-to-one relation between the dependence parameter—denoted  $\theta$ —and Kendall’s nonparametric measure of association. For any copula  $C_\theta$ , Kendall’s tau is related to  $\theta$  in the following way:

$$\tau(\theta) = 4EC_\theta(U, V) - 1, \quad (1)$$

where  $(U, V)$  is distributed as  $C_\theta$ , and  $E$  denotes the expectation operator with respect to  $U$  and  $V$ . Appendix B displays closed-form formulas for the population value of Kendall’s tau for some of the copula models under consideration.

This relation suggests a natural way to estimate the copula. An estimate of  $\theta$  is readily obtained by computing the sample version of tau on a (sub)sample of paired index-return observations, inverting Relation (1), and plugging in the sample tau.<sup>1</sup> This method-of-moment type procedure yields a rank-based estimate of the association parameter which is consistent, under the assumption that the selected family of copulas describes accurately the dependence structure of the equity indexes. Other methods could be used without fundamentally altering this approach, e.g., inversion of Spearman’s rho, or the maximum pseudo-likelihood method.

The proposed technique assumes that the objective and risk-neutral copulas are identical. Rosenberg (2003) makes this assumption as well. If multivariate option price data were available, this assumption could be tested or the appropriate risk-neutral copula could be estimated. Only data on prices of multivariate claims would reveal information about the risk-neutral dependence structure. Information about the risk-neutral depen-

---

<sup>1</sup>The sample version of Kendall’s tau is defined as follows. Let  $\{(X_1, Y_1), \dots, (X_n, Y_n)\}$  be a random sample of  $n$  observations from a vector  $(X, Y)$  of continuous random variables. Two distinct pairs  $(X_i, Y_i)$  and  $(X_j, Y_j)$  are said to be concordant if  $(X_i - X_j)(Y_i - Y_j) > 0$ , and discordant if  $(X_i - X_j)(Y_i - Y_j) < 0$ . Kendall’s tau for the sample is then defined as  $t = (c - d)/(c + d)$ , where  $c$  denotes the number of concordant pairs, and  $d$  is the number of discordant pairs.

dence structure can never be extracted from univariate option prices—which *are* available—as these only bear relevance to the risk-neutral marginal processes, and not to the joint risk-neutral process. Identification of the multivariate density requires knowledge of both the marginal densities and the dependence function that links them together.

Time variation in the copula is modelled by allowing the parameter of dependence to evolve through time according to a particular equation. The forcing variables in this equation are the conditional volatilities of the underlying assets. These are also the forcing variables that are typically chosen to model time-varying correlations; see, e.g., the BEKK model introduced by Engle and Kroner (1995). Additional motivation is provided by the evidence on correlation breakdowns, which suggests that financial markets exhibit high dependence in periods of high volatility. Patton (2002a) proposes an ARMA-type process linking the dependence parameter to absolute differences in return innovations, which is another way to capture the same idea.

To be more specific, let  $\tau_t$  be Kendall’s measure of association at time  $t$ , and let  $h_{i,t}$  be the objective conditional variance estimate at time  $t$  of underlying index return  $i \in \{1, 2\}$  implied by Duan’s GARCH option pricing model. It is assumed that

$$\tau_t = \gamma(h_{1,t}, h_{2,t}) \tag{2}$$

for some function  $\gamma(\cdot, \cdot)$  to be specified later. This conditional measure of association governs the degree of dependence for the risk-neutral copula under consideration.

The proposed valuation scheme is implemented using Monte Carlo simulations. Pairs of random variates are drawn from the copula implied by the estimated conditional risk-neutral measure of association, which are then transformed to return innovations using Duan’s GARCH model. Subsequently, the payoffs implied by these innovations are averaged and discounted at the risk-free rate. The result then constitutes the fair value of the option. Algorithms for random variable generation from the non-normal copulas are given in Genest and MacKay (1986), Genest (1987), Ghoudi, Khoudraji, and Rivest (1998), and Nelsen (1999). For the normal copula, a straightforward Cholesky decomposition may be used.

### 3 Pricing options on two equity indexes

The dynamic-dependence valuation scheme outlined in Section 2 is applied to better-of-two-markets and worse-of-two-markets options on the S&P500 and Nasdaq indexes. A sample consisting of pairs of daily returns on the S&P500 and Nasdaq from January 1, 1993 to August 30, 2002 was obtained from Datastream. The sample size is  $T = 2422$ . The maximum likelihood estimates of the GARCH parameters for the marginal index return processes may be found in Table I. The values for  $\alpha$  and  $\beta$  nearly add up to one. These estimates are in line with previously reported values.

Figure 1 depicts the time series of the estimated standardized GARCH innovations  $(\eta_{1,t+1}/\sqrt{h_{1,t}}, \eta_{2,t+1}/\sqrt{h_{2,t}})$  for the last 250 trading days in the sample. (For clarity, the picture is restricted to a subsample; other episodes show a similar pattern.) Note that outliers typically occur simultaneously and in the same direction. This positive dependence between the two series is even more apparent from Figure 2, which displays the support set of the empirical copula of the standardized return innovations. This scatter plot consists of the observed pairs of ranks (divided by  $T + 1$ ) for the estimated standardized GARCH innovations of the two markets. Under regularity conditions, the empirical copula function converges to the true (here, objective) copula function; see Van der Vaart and Wellner (1996). Notice the pronounced positive dependence, particularly in the tails. The sample version of Kendall's tau for the entire sample amounts to 0.60, confirming positive dependence. Figure 3 gives an impression of how this dependence measure of the standardized return innovations evolves over time. It shows rolling-window estimates of Kendall's tau using window sizes of two months, i.e., Kendall's tau at day  $t$  is computed using the 20 trading days prior to day  $t$ , day  $t$  itself, and the 20 trading days after day  $t$ . While the estimates show considerable variation, a slightly upward trend over the sample period is discernable.

The time variation in the copula is governed by Equation (2). It models the dependence measure as a function of the conditional volatilities of the index returns. The following specification of this function is proposed:

$$\gamma(h_1, h_2) = \gamma_0 + \gamma_1 \log \max(h_1, h_2). \quad (3)$$

To motivate this specification, recall that the evidence on correlation breakdowns suggests that increased dependence occurs in hectic periods. Hence, theory predicts a positive value of  $\gamma_1$ . The maximum operator reflects that hectic periods in either market may cause dependence to go up. Since volatility in both markets is highly dependent, the actual specification is likely not

to affect the results in the present section too much. The parameters  $\gamma_0$  and  $\gamma_1$  were estimated by regressing the rolling-window estimates of Kendall's tau on the estimated log maximum conditional volatility. This is illustrated in Figure 4. The slope coefficient,  $\gamma_1$ , was estimated at 0.063; positive, as expected. The estimated dependence measure implied by these parameter estimates,

$$\gamma(h_{1,t}, h_{2,t}) = \gamma_0 + \gamma_1 \log \max(h_{1,t}, h_{2,t}),$$

was then used to fix the conditional risk-neutral copula at time  $t$ . Return innovations were sampled from this conditional copula to compute the price of the option. In total, the Monte Carlo study was based on 100,000 replications, leading to simulation errors in the order of magnitude of 1 basis point for one-month maturity claims.

Clearly, the option price depends on the initial levels of volatility of the underlyings. Prices for three levels of initial volatility were computed: low, medium, and high volatility, where medium volatility is defined as the estimated unconditional variance  $\omega/(1 - \beta - \alpha)$ , and low and high volatility are one-fourth of and four times this amount, respectively. Furthermore, different maturities were considered, ranging from one day to one month (i.e., 20 trading days). The strike price was set at levels between .98 and 1.02. Finally, the risk-free rate was assumed to be 4 percent per annum.

The results show that allowing for time varying dependence leads to different option prices than under static dependence, in particular in times of high volatility. This is illustrated in Figure 5 which displays, for various copula parametrizations, the price (measured in basis points) of a one-month put-on-max option as a function of the exercise price implied by dynamic dependence, and compares it to the option price under three levels of static dependence: low, medium, and high static dependence. The medium level of dependence is equal to the average measure of dependence found in the sample, 0.60; the low and high levels are 0.40 and 0.80, respectively. Note that a static model for the dependence structure, which uses the sample measure of dependence of 0.60, underestimates the option price generated by the dynamic model considerably for all copula parametrizations and over the entire range of strike prices considered. The difference is significant since the 95% confidence intervals of the price estimates do not overlap. In the interest of clarity, confidence intervals are not displayed here, but available from the authors upon request. Note that the prices implied by dynamic copulas are between the high and the medium static-dependence prices, suggesting that the dynamic model implies a dependence that is on average stronger than in the medium static-dependence case. Interestingly, price differences between

the dynamic and static model vanish as initial volatilities are at a medium level; see Figure 6. The same holds for low initial volatilities (not shown), again, across a broad range of copula families and strike prices.

It is also interesting to compare option prices produced by different dynamic copula families. It turns out that prices implied by the normal copula deviate substantially from prices implied by the other copula families. Outside the normal class, the copula choice appears to be irrelevant. This suggests that unless the dependence between index returns can be described by a normal model, alternative specifications should be considered. These findings are illustrated in Figures 7 and 8 which depict dynamic-dependence one-month call-on-max and put-on-min option prices respectively, as a function of their strike under medium initial volatilities. The prices implied by the normal copula are significantly lower than the prices implied by the other copulas across the whole range of strike prices. The effect is there at other maturities as well. The difference between normal and non-normal prices is also found for high and low initial volatility levels. The differences are less significant for call-on-min and put-on-max options.

## 4 Conclusions

This paper studies the relation between multivariate options prices and the dependence structure of the underlying assets. A copula-based model was proposed for the valuation of claims on multiple assets. A novel feature of the proposed model is that, contrary to earlier works on multivariate option pricing, the dependence structure is not taken as fixed, but rather as potentially varying with time. The time variation in the dependence structure was modelled using various parametric copulas by letting the copula parameter depend on the conditional volatilities of the underlyings.

This dynamic copula model was applied to better- and worse-of-two-markets options on the S&P500 and Nasdaq indexes for a variety of copula parametrizations. Option prices implied by the dynamic model turned out to differ substantially from prices implied by a model that fixes the dependence between the underlying indexes, especially in high-volatility market conditions. Hence, the application suggests that time variation in the dependence between the S&P500 and the Nasdaq is important for the price of options on these indexes. A comparison of option prices computed from different copula families shows that the normal family produces prices that differ significantly from the ones implied by the non-normal alternatives. These findings suggests that if the dependence between the index returns

is not well represented by a normal copula, alternative copulas need to be considered. The empirical relevance of such alternatives is apparent given the evidence of non-normality in financial markets.

## A One-Parameter Copula Families

The table below displays several one-parameter copula families.

Frank	$C_\theta(u, v) = \frac{1}{\theta} \log \left\{ 1 + \frac{(e^{\theta u} - 1)(e^{\theta v} - 1)}{e^\theta - 1} \right\}$
Gumbel–Hougaard	$C_\theta(u, v) = \exp \left\{ - \left(  \log u ^\theta +  \log v ^\theta \right)^{\frac{1}{\theta}} \right\}$
Plackett	$C_\theta(u, v) = \frac{1 + (\theta - 1)(u + v) - \sqrt{[1 + (\theta - 1)(u + v)]^2 - 4uv\theta(\theta - 1)}}{2(\theta - 1)}$
Galambos	$C_\theta(u, v) = uv \exp \left\{ \left(  \log u ^\theta +  \log v ^\theta \right)^{-\frac{1}{\theta}} \right\}$
Normal	$C_\theta(u, v) = N_\theta(\Phi^{-1}(u), \Phi^{-1}(v))$

Note:  $\Phi$  is the standard (univariate) normal distribution function, and  $N_\theta$  denotes the standard bivariate normal distribution function with correlation coefficient  $\theta$ .

## B Kendall's tau

The table below provides expressions—closed-form if available—of the relation between Kendall's tau and the dependence parameter for the copula families considered in Appendix A.

Frank	$\tau(\theta) = 1 - 4 \{D_1(-\theta) - 1\} / \theta$
Gumbel-Hougaard	$\tau(\theta) = 1 - 1/\theta$
Plackett	$\tau(\theta) = 4 \int_0^1 \int_0^1 C_\theta(u, v) dC_\theta(u, v) - 1$
Galambos	$\tau(\theta) = \frac{\theta+1}{\theta} \int_0^1 \left( \frac{1}{t^{1/\theta}} + \frac{1}{(1-t)^{1/\theta}} - 1 \right)^{-1} dt$
Normal	$\tau(\theta) = \frac{2}{\pi} \arcsin \theta$

Note:  $D_1$  denote the first-order Debye function,  $D_1(-\theta) = \frac{1}{\theta} \int_0^\theta \frac{t}{e^t - 1} dt + \frac{\theta}{2}$ .

## References

- Ait-Sahalia, Yacine, and Andrew W. Lo, 1998, Nonparametric Estimation of State-Price Densities Implicit in Financial Asset Prices, *Journal of Finance* 53, 499–547.
- Black, Fisher, and Myron S. Scholes, 1973, The Pricing of Options and Corporate Liabilities, *Journal of Political Economy* 81, 637–654.
- Bollerslev, Tim, 1986, Generalized Autoregressive Conditional Heteroskedasticity, *Journal of Econometrics* 31, 307–327.
- Boyer, Brian H., Michael S. Gibson, and Mico Loretan, 1999, Pitfalls in Tests for Changes in Correlations, Working paper, International Finance Discussion Papers, No. 597, Board of Governors of the Federal Reserve System.
- Cherubini, Umberto, and Elisa Luciano, 2002, Multivariate Option Pricing with Copulas, *ICER Applied Mathematics Working Paper Series, WP no. 5*.
- Cox, John C., Stephen A. Ross, and Mark Rubinstein, 1979, Option Pricing: A Simplified Approach, *Journal of Financial Economics* 7, 229–263.
- Duan, Jin-Chuan, 1995, The Garch Option Pricing Model, *Mathematical Finance* 5, 13–32.
- Embrechts, Paul, Alexander J. McNeil, and Daniel Straumann, 2002, Correlation and dependence in risk management: properties and pitfalls, in M. A. H. Dempster, eds.: *Risk Management: Value at Risk and Beyond* (Cambridge University Press, Cambridge, England ).
- Engle, Robert F., and Kenneth F. Kroner, 1995, Multivariate Simultaneous Generalized ARCH, *Econometric Theory* pp. 122–150.
- Genest, Christian, 1987, Frank’s Family of Bivariate Distributions, *Biometrika* 74, 549–555.
- Genest, Christian, and R. Jock MacKay, 1986, Copules archimédiennes et familles de lois bidimensionnelles dont les marges sont données, *The Canadian Journal of Statistics* 14, 145–159.
- Ghoudi, K., A. Khoudraji, and L.-P. Rivest, 1998, Propriétés Statistiques des Copules de Valeurs Extrêmes Bidimensionnelles, *The Canadian Journal of Statistics* 26, 187–197.

- Johnson, Herb, 1987, Options on the Maximum or the Minimum of Several Assets, *Journal of Financial and Quantitative Analysis* 22, 277–283.
- Margrabe, William, 1978, The Value of an Option to Exchange One Asset for Another, *Journal of Finance* 33, 177–186.
- Nelsen, Roger B., 1999, *An Introduction to Copulas. Lecture Notes in Statistics no 139.* (Springer-Verlag, New York).
- Patton, Andrew J., 2002a, Modelling Time-Varying Exchange Rate Dependence Using the Conditional Copula, *Working Paper*.
- Patton, Andrew J., 2002b, Skewness, Asymmetric Dependence, and Portfolios, *Working Paper*.
- Reiner, Eric, 1992, Quanto Mechanics, in *From Black-Scholes to Black Holes: New Frontiers in Options* (RISK Books, London ).
- Rosenberg, Joshua V., 1998, Pricing Multivariate Contingent Claims using Estimated Risk-Neutral Density Functions, *Journal of International Money and Finance* 17, 229–247.
- Rosenberg, Joshua V., 1999, Semiparametric Pricing of Multivariate Contingent Claims, *Stern School of Business, Working Paper S-99-35*.
- Rosenberg, Joshua V., 2003, Nonparametric Pricing of Multivariate Contingent Claims, *Journal of Derivatives* 10.
- Shimko, David C., 1994, Options On Futures Spreads: Hedging, Speculation, and Valuation, *Journal of Futures Markets* 14, 183–213.
- Sklar, A., 1959, Fonctions de répartition à  $n$  dimensions et leurs marges, *Publ. Inst. Statist. Univ. Paris* 8, 229–231.
- Stulz, René M., 1982, Options on the Minimum or the Maximum of Two Risky Assets: Analysis and Applications, *Journal of Financial Economics* 10, 161–185.
- Van der Vaart, Aad W., and Jon A. Wellner, 1996, *Weak Convergence and Empirical Processes.* (Springer-Verlag, New York).

Table I: Maximum likelihood estimates of the GARCH parameters for the marginal index return processes. Figures in brackets are robust quasi-maximum likelihood standard errors.

Parameter	S&P500		Nasdaq	
$\mu \times 10^2$	0.0674	(0.0168)	0.0812	(0.0246)
$\omega \times 10^5$	0.0680	(0.0398)	0.1895	(0.0987)
$\beta$	0.9258	(0.0220)	0.8906	(0.0309)
$\alpha$	0.0680	(0.0198)	0.1015	(0.0288)

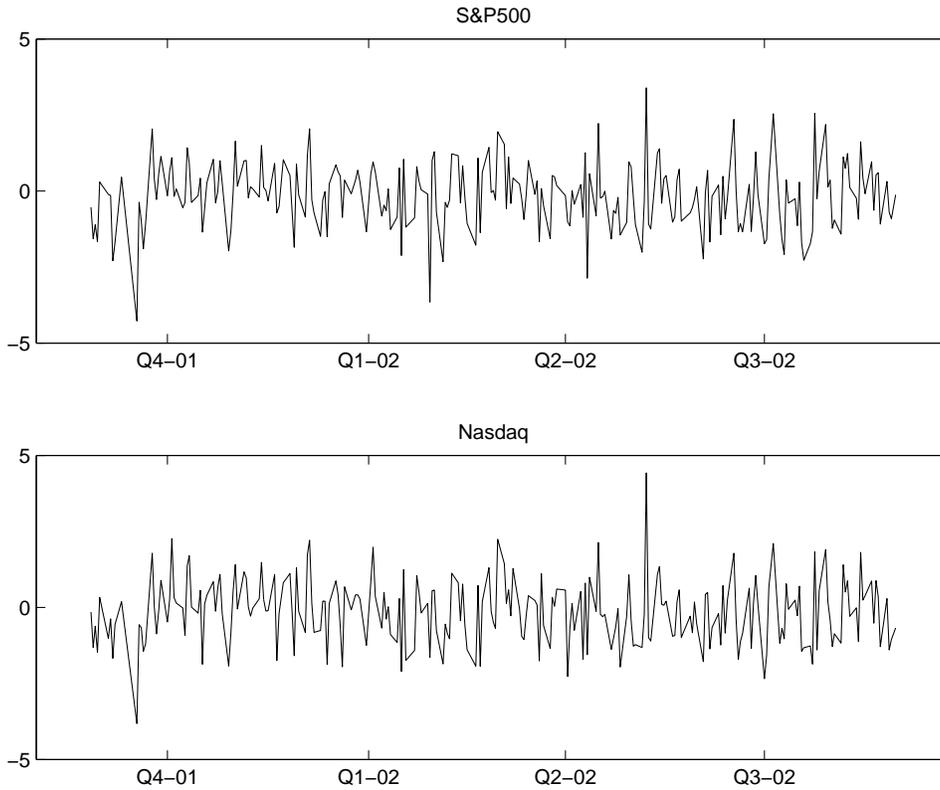


Figure 1: Daily standardized GARCH innovations for S&P500 and Nasdaq for the last 250 trading days in the sample.

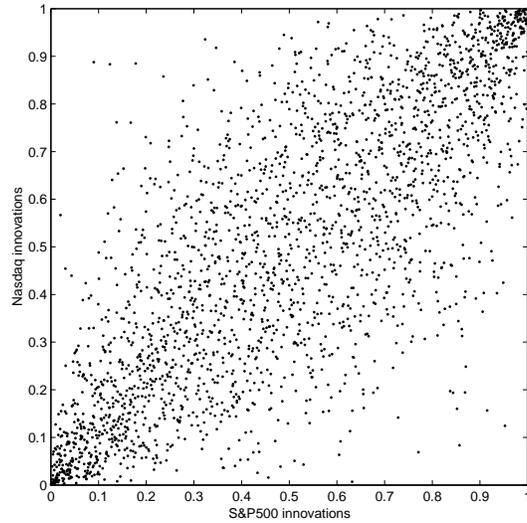


Figure 2: Support set of the empirical copula of the standardized GARCH innovations.

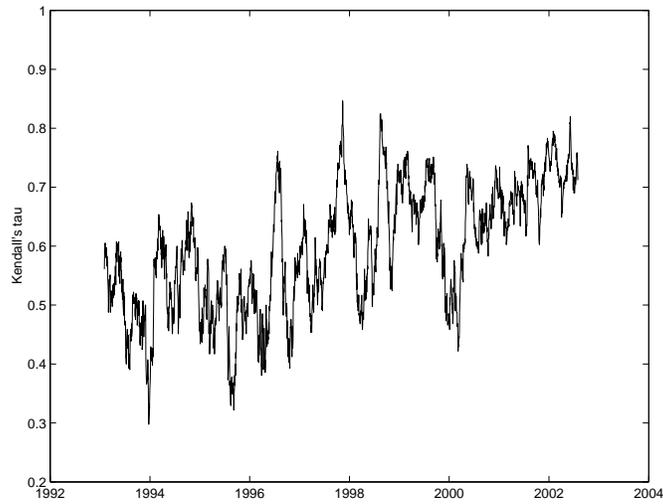


Figure 3: Rolling-window estimates of Kendall's tau for the standardized return innovations using a window size of 41 trading days.

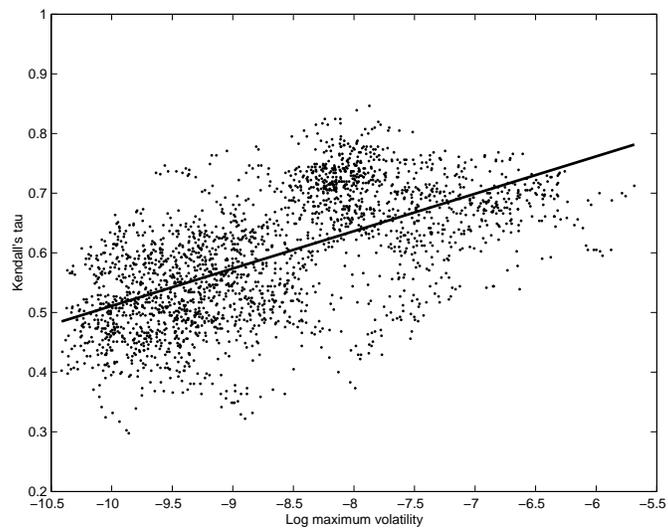


Figure 4: Regression of rolling-window estimates of Kendall's tau for the standardized return innovations on the logarithm of the maximum return volatility.

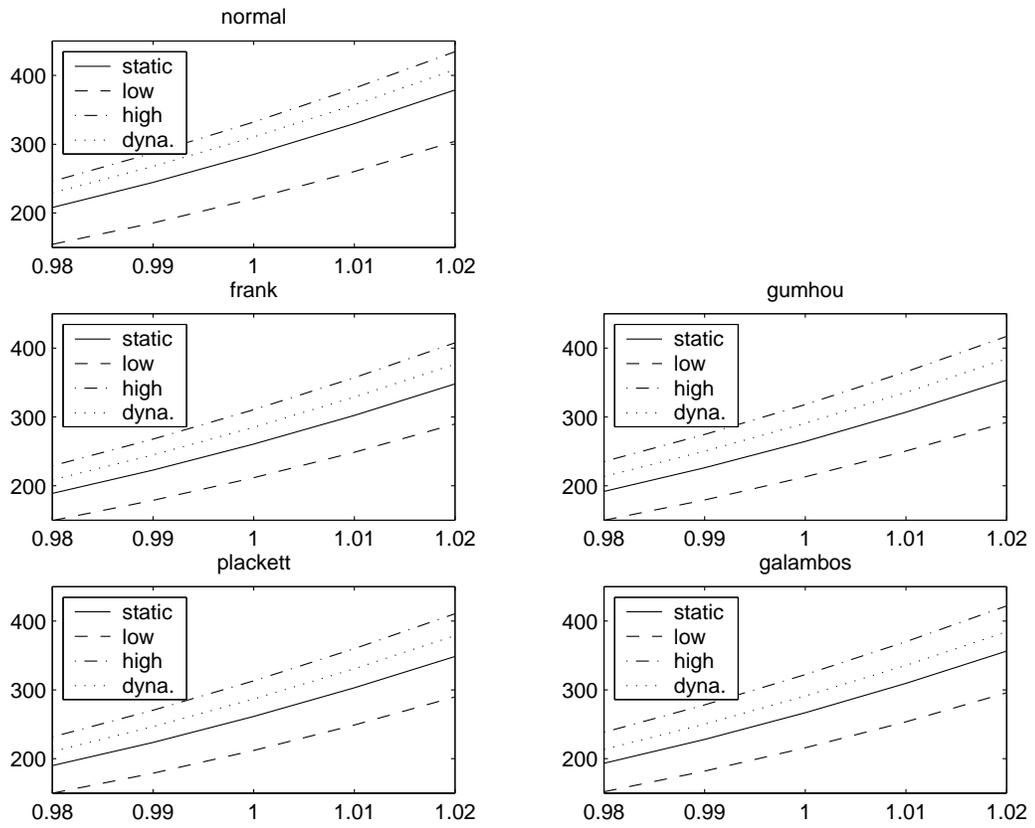


Figure 5: One-month maturity put-on-max prices as a function of the strike under high initial volatilities for dynamic dependence and for low, medium, and high static dependence for various copulas.

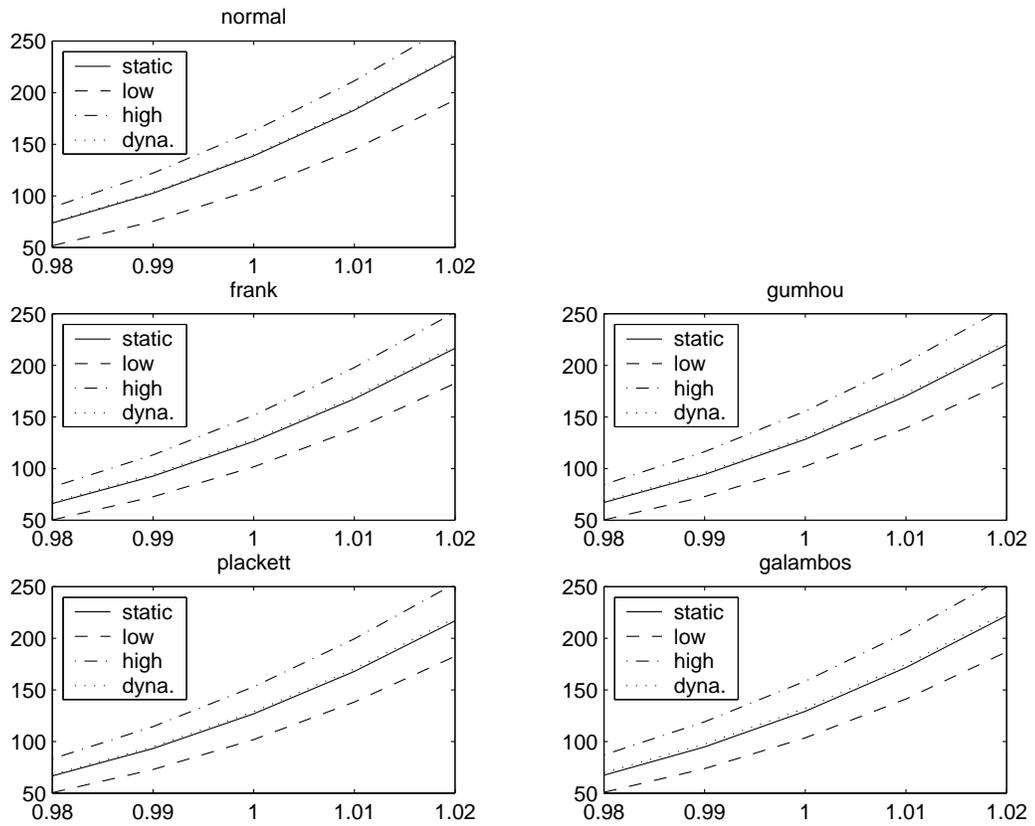


Figure 6: One-month maturity put-on-max prices as a function of the strike under medium initial volatilities for dynamic dependence and for low, medium, and high static dependence for various copulas.

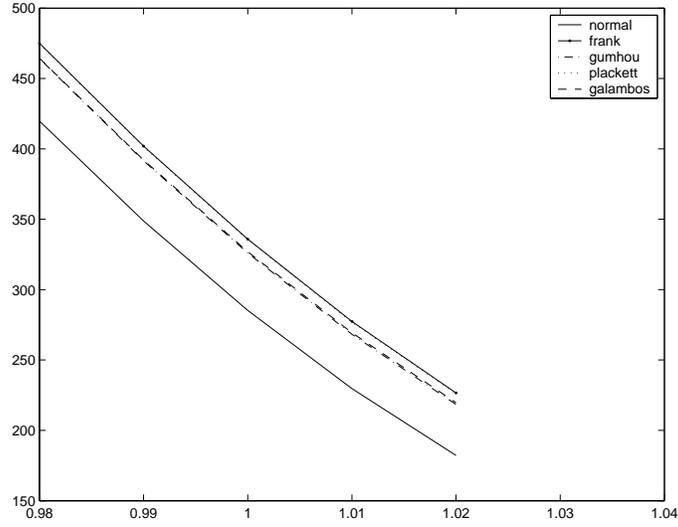


Figure 7: One-month call-on-max prices as a function of the strike under dynamic dependence and medium initial volatilities for various copula models.

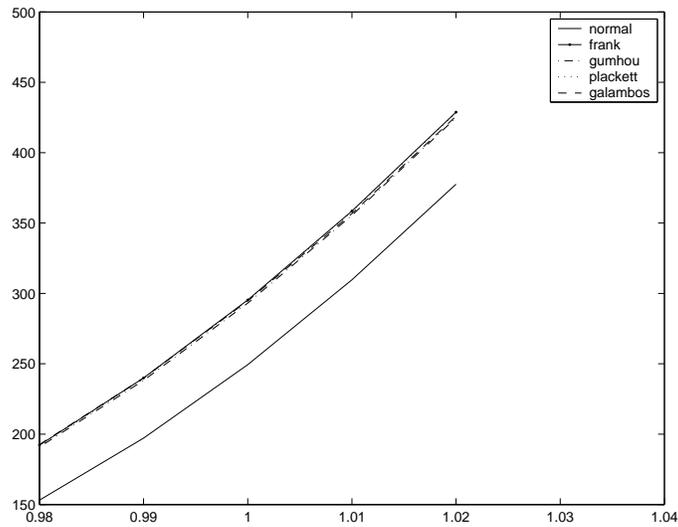


Figure 8: One-month put-on-min prices as a function of the strike under dynamic dependence and medium initial volatilities for various copula models.