EVALUATING DENSITY FORECASTS VIA THE COPULA APPROACH

by

Xiaohong Chen and Yanqin Fan

Working Paper No. 02-W25R

October 2002
Revised September 2003

DEPARTMENT OF ECONOMICS
VANDERBILT UNIVERSITY
NASHVILLE, TN 37235

www.vanderbilt.edu/econ

Evaluating Density Forecasts via the Copula Approach*

Xiaohong Chen
Department of Economics
New York University
New York, NY 10003
Email: xiaohong.chen@nyu.edu

Yanqin Fan
Department of Economics
Vanderbilt University
Nashville, TN 37235-1819
Email: yanqin.fan@vanderbilt.edu

First version: October 2002
This version: September 2003

Abstract

In this paper, we develop a general approach for constructing simple tests for the correct density forecasts, or equivalently, for i.i.d. uniformity of appropriately transformed random variables. It is based on nesting a series of i.i.d. uniform random variables into a class of copula-based stationary Markov processes. As such, it can be used to test for i.i.d. uniformity against alternative processes that exhibit a wide variety of marginal properties and temporal dependence properties, including skewed and fat-tailed marginal distributions, asymmetric dependence, and positive tail dependence. In addition, we develop tests for the dependence structure of the forecasting model that are robust to possible misspecification of the marginal distribution.

*We thank the editor and an anonymous referee for valuable comments.

1 Introduction

How to evaluate density forecasts of parametric dynamic models is very important for risk management in finance and insurance. Recently, a number of authors including Diebold, Gunther, and Tay (1998), Diebold, Tay, and Wallis (1999), Diebold, Hahn, and Tay (1999), Clements and Smith (2000), and Elerian, Chib, and Shephard (2001) have applied and extended Rosenblatt's (1952) transformation to evaluating density forecasts.[1]

Specifically, let $\{Y_t\}_{t=1}^n$ denote a time series and Ω_t represent the information set at time t (not including Y_t). Let $F_t(\cdot|\Omega_t)$ be the forecast of the distribution of Y_t given the information Ω_t. Diebold, Gunther, and Tay (1998) and Bai (2003) show that the transformed variables $U_t = F_t(Y_t|\Omega_t)$, $t = 1, \ldots, n$, are i.i.d. uniform over (0,1) if and only if the forecasts $\{F_t(\cdot|\Omega_t)\}$ are correct. Based on this result, Diebold, Gunther, and Tay (1998) propose to evaluate density forecasts by checking the uniformity and serial independence of the U_t's graphically via the histogram of $\{U_t\}$ and the correlograms of $\{(U_t - \bar{U})^i\}$, where $i = 1, 2, 3, 4$ and $\bar{U}$ is the sample mean of the U_t's. Via a simulation study, they demonstrate that the proposed procedures reveal useful information about possible deviations of the forecast model from the correct model. Bai (2003), among others, proposes a consistent test for uniformity of the U_t's based on the empirical distribution function under the assumption that the U_t's are i.i.d.. Hong and Li (2002) develop a joint test for uniformity and serial independence by comparing a kernel estimate of the joint density function of U_{t-j} and U_t with the uniform density on the unit square and report superior Monte Carlo performance of their test over existing tests. Thompson (2002) provides theoretical justification for the graphical procedures used in papers such as Diebold, Gunther, and Tay (1998) by describing a family of specification tests for uniformity and serial independence based on the empirical distribution function and/or the sample periodogram.

Berkowitz (2001) argues that consistent nonparametric tests typically require the availability of large data sets to achieve accurate inference. Instead he advocates simple parametric tests. Noting the potential difficulties in nesting the null hypothesis of i.i.d. uniformity in a more general parametric class of models, Berkowitz (2001) first transforms U_t to $Z_t = \Phi^{-1}(U_t)$ in which $\Phi(\cdot)$ is the standard normal distribution function, and then imposes a linear autoregressive (AR) structure on the Z_t's. Under the linear AR structure for $\{Z_t\}$, testing i.i.d. uniformity of the U_t's and testing serial independence of the U_t's are equivalent to testing certain hypotheses on the specific values of the parameters in the AR model, which can be carried out by likelihood ratio (LR) tests.

The main objective of this paper is to introduce a general approach for constructing parametric tests for the correct density forecasts, or equivalently, for i.i.d. uniformity of $\{U_t\}$, that are valid for a wide variety of alternative processes, such as those with skewed and fat-tailed marginal

[1] Rosenblatt's (1952) transformation allows us to transform a multivariate random vector into a vector of i.i.d. uniform scalar random variables. Recently, this transformation has also become popular in testing parametric time series models. For instance, Bai (2003), Corradi and Swanson (2001), and Inoue (1999) have applied it to testing the parametric specification of conditional distributions of dynamic models, while Hong and Li (2002) and Thompson (2002) have applied it in testing diffusion models.

distributions, asymmetric dependence, and positive tail dependence. It is based on nesting a series of i.i.d. uniform random variables into a class of copula-based stationary Markov processes. In addition, we develop tests for serial independence of $\{U_t\}$ or the dependence structure of the forecasting model that are robust to possible misspecification of the marginal distribution of U_t.

A copula is simply a multivariate probability distribution function with uniform marginals. The importance of copulas in multivariate modeling is justified by the Sklar's (1959) theorem. Loosely speaking, it states that any given multivariate distribution can be expressed as the copula function evaluated at its marginal distribution functions and any given copula function when evaluated at any marginal distributions becomes a multivariate distribution. Hence the information in the joint distribution is decomposed into those in the marginal distributions and that in the copula function. Consequently a copula allows one to model the marginal distributions and the dependence structure of a multivariate random variable separately. Moreover, the copula measure of dependence is invariant to any increasing transformations. See Joe (1997) and Nelsen (1999) for detailed discussions on the theory and examples of copulas.

Because of their flexibility in modeling the distribution of multivariate random variables, copulas have gained popularity in the finance and insurance community over the last few years. We refer readers to Embrechts, McNeil and Straumann (1999), and Patton (2001) for extensive lists of references in the finance and insurance literatures. While most existing papers in the finance and insurance literatures use copulas to model the contemporaneous dependence between two time series, a number of papers have employed copulas to model temporal dependence of a time series: Joe (1997) proposes a class of stationary first order Markov processes based on copulas; Darsow, Nguyen, and Olsen (1992) provide a necessary and sufficient condition for a copula-based process to be Markovian; Bouyé, Gaussel, and Salmon (2002) use copulas to model stationary nonlinear autoregressive dependence of time series and apply the models to financial returns and transactions based forex data; Gagliardini and Gourieroux (2002) suggest a class of proportional hazard duration time series models and the corresponding copula functions; Chen and Fan (2002) establish the temporal dependence properties of a class of copula-based stationary time series models and the asymptotic properties of a semiparametric estimator of the copula function. The simulation results in Chen and Fan (2002) reveal the richness of both the marginal and dependence patterns that can be generated by copula-based time series models.

In this paper, we make use of the recent developments in copula-based time series modelling to develop tools for evaluating density forecasts. In particular, we establish tests for correct density forecasts, or i.i.d. uniformity of the U_t's, against the alternative that $\{U_t\}$ is a stationary copula-based Markov process. Since a copula is itself a multivariate distribution with uniform marginals, it provides a natural way to model the temporal dependence structure of the transformed variables $\{U_t\}$. As we will show in the next section, the linear AR structure imposed on $\{Z_t\}$ in Berkowitz (2001) implies that the process $\{U_t\}$ is a copula-based Markov process with the Gaussian copula and a specific parametric marginal distribution. It is well known that the dependence structure of time series characterized by the Gaussian copula is symmetric regardless of the marginal distribution. As a result, the tests in Berkowitz (2001) do not have power against alternative processes of $\{U_t\}$ that exhibit asymmetric dependence. By choosing a non-Gaussian copula such as the Frank copula, the Clayton copula, among many others, the tests developed in this paper have power against processes exhibiting complicated nonlinear asymmetric dependence. On the other hand, by choosing fat-tailed distributions such as the Student's t distribution, our tests have power against time series with extreme values. In addition, by leaving the marginal distribution unspecified, we develop tests for the dependence structure of the forecasting model or the serial independence of $\{U_t\}$ that are robust to the choice of the marginal distribution of U_t. Another advantage of the copula approach in our context is that copulas and hence copula-based measures of dependence are invariant to strictly increasing transformations of random variables. This ensures that in cases where U_t is an increasing transformation of the innovation of the forecast model, the dependence structure of the innovations of the model is the same as that of $\{U_t\}$ as modeled by copulas. This information may then be used to update the forecast model in case it is being rejected.

The rest of this paper is organized as follows. In Section 2, we introduce the class of copula-based stationary time series models for $\{U_t\}$ that nest the null hypothesis of i.i.d. uniformity. Section 3 first develops tests for the null hypothesis of correct density forecasts or the i.i.d. uniformity of $\{U_t\}$ for the class of copula-based stationary time series models whose marginal distribution belongs to a parametric family of distributions and then tests for the null hypothesis of the correct specification of the dependence structure of the forecasting model or the serial independence of $\{U_t\}$ for the class of semiparametric models in which the marginal distribution is unspecified. The last section concludes. In this paper, it is understood that the null hypothesis of correct density forecasts is equivalent to that of i.i.d. uniformity of $\{U_t\}$, and the null hypothesis of the correct specification of the dependence structure of the forecasting model is equivalent to that of the serial independence of $\{U_t\}$.

2 The Class of Time Series Models for $\{U_t\}$

Let $\{U_t : t = 1, ..., n\}$ be the transformed variables defined in the Introduction. Throughout the rest of this paper, we will work with the following assumption:

Assumption 1: $\{U_t : t = 1, ..., n\}$ *is a sample of a stationary first-order Markov process generated from* $(G^*(\cdot), C(\cdot, \cdot; \alpha^*))$*, where* $G^*(\cdot)$ *is the true invariant distribution which is absolutely continuous with respect to Lebesgue measure on real line;* $C(\cdot, \cdot; \alpha^*)$ *is the true parametric copula for* (U_{t-1}, U_t) *up to unknown value* α^* *and is absolutely continuous with respect to Lebesgue measure on* $[0, 1]^2$.

Under **Assumption 1**, the null hypothesis of i.i.d. uniform U_t's is equivalent to the joint hypothesis that $G^*(u)$ is the uniform distribution and $C(v_1, v_2; \alpha^*) = v_1 v_2$ is the independence copula. By choosing $G^*(\cdot)$ and $C(\cdot, \cdot; \alpha^*)$ separately, one can design tests for i.i.d. uniformity of the U_t's that have power against alternative processes that exhibit a variety of marginal and dependence properties such as skewed, fat-tailed marginal distributions and nonlinear, asymmetric dependence, among others.

The following example shows that the approach taken in Berkowitz (2001) is a special case of our general copula approach in which the copula function is the Gaussian copula and the invariant distribution G^* takes a specific form.

Example 1: Let the copula $C(\cdot, \cdot; \alpha^*)$ be the Gaussian copula:

$$C(v_1, v_2; \alpha^*) = \Phi_{\alpha^*}(\Phi^{-1}(v_1), \Phi^{-1}(v_2)),$$

where $\Phi_{\alpha^*}(\cdot, \cdot)$ is the distribution function of the bivariate normal distribution with means zero, variances 1, and correlation coefficient α^*. Then the process $\{\Phi^{-1}(G^*(U_t))\}$ is a Gaussian process and hence

$$\Phi^{-1}(G^*(U_t)) = \alpha^* \Phi^{-1}(G^*(U_{t-1})) + \epsilon_t, \tag{2.1}$$

where $\epsilon_t \sim N(0, 1 - \alpha^{*2})$. If the marginal distribution $G^*(\cdot)$ is of the form:

$$G^*(u) = \Phi(\frac{\Phi^{-1}(u) - \mu}{\sigma/\sqrt{1 - \alpha^{*2}}}), \tag{2.2}$$

then (2.1) reduces to the linear AR(1) model for $\{Z_t = \Phi^{-1}(U_t)\}$ adopted in Berkowitz (2001):

$$Z_t - \mu = \rho(Z_{t-1} - \mu) + \eta_t, \quad \text{where } \eta_t \text{ is i.i.d. } N(0, \sigma^2).$$

By allowing $G^*(\cdot)$ to be a different distribution than (2.2), one can design tests for i.i.d. uniformity that have power against processes characterized by the Gaussian copula, but different marginal distributions. Coupled with fat-tailed distributions such as the Student's t distribution, time series models with the Gaussian copula can produce extremely large and small values. However no clustering of such extremely large or small values occurs as the Gaussian copula does not have tail dependence. Moreover, it is well known that the dependence structure of time series characterized by the Gaussian copula is symmetric regardless of its marginal distribution. To develop tests that have power against processes exhibiting complicated nonlinear asymmetric dependence and clusters of large and/or small values, non-Gaussian copulas must be used. A wide variety of non-Gaussian copulas is available, see Joe (1997) and Nelsen (1999) for properties of specific copulas. Chen and Fan (2002) provide plots of time series generated from the Joe-Clayton copula coupled with normal and student's t distributions. These plots reveal the richness of both the marginal and temporal dependence structure that can be generated by copula-based time series models.

The copula approach facilitates the incorporation of additional information into the model. For example, suppose one tests the null hypothesis of i.i.d. uniformity sequentially by testing the hypothesis of uniformity first and then that of serial independence. In cases where the uniformity of U_t is not rejected, one may wish to take this into account when testing for serial independence. This can be done easily by choosing an appropriate copula function as the joint distribution of U_{t-1}

and U_t and then testing if the copula function is the independence copula. On the other hand, if one wishes to test the serial independence first, tests that are robust to the choice of the marginal distribution would be desirable. Consider, for example, the Gaussian copula model. Klaassen and Wellner (2001) show that in the Gaussian copula model, $|\alpha^*|$ equals the maximum correlation coefficient of U_{t-1} and U_t. Hence U_{t-1} and U_t are independent if and only if $\alpha^* = 0$. Robust tests for serial independence of $\{U_t\}$ can thus be constructed from consistent estimators of α^*. Since (2.1) implies

$$\alpha^* = \text{corr}(\Phi^{-1}(G^*(U_{t-1})), \Phi^{-1}(G^*(U_t))), \tag{2.3}$$

one gets $\alpha^* = \text{corr}(U_{t-1}, U_t)$ if and only if $G^*(\cdot) = \Phi(\cdot)$, which is impossible due to the $[0,1]$ support of U_t. As a result, in the Gaussian copula model, the correlogram of $\{U_t\}$ can never capture all the dependence structure of $\{U_t\}$. Instead the correlogram of $\{\Phi^{-1}(G^*(U_t))\}$ does.

One important property of copulas that makes them useful in evaluating density forecasts is their invariance to strictly increasing transformations of random variables (see Theorem 2.4.3 in Nelsen (1999). This may be useful in updating the forecast model in case it is being rejected, as shown in the following example.

Example 2: Consider Vasicek's (1977) model:

$$dY_t = \kappa(\beta - Y_t)dt + \sigma dW_t, \tag{2.4}$$

where $\{W_t\}$ is the standard Brownian motion, β is the long run mean, and κ is the speed of mean reversion to β. The conditional distribution of Y_t given Ω_t is $N(\beta + (Y_{t-1} - \beta)e^{-\kappa}, V_E(1 - e^{-2\kappa}))$ in which $V_E = \sigma^2/(2\kappa)$. Hence

$$U_t = \Phi(\frac{Y_t - [\beta + (Y_{t-1} - \beta)e^{-\kappa}]}{\sqrt{V_E(1 - e^{-2\kappa})}}). \tag{2.5}$$

Let $\epsilon_t = Y_t - [\beta + (Y_{t-1} - \beta)e^{-\kappa}]$. Since $\Phi(\cdot)$ is a strictly increasing function, the copula associated with the joint distribution of ϵ_{t-1} and ϵ_t is the same as that associated with the joint distribution of U_{t-1} and U_t. The dependence structure of the process $\{\epsilon_t\}$ may then be incorporated into the original model for $\{Y_t\}$. For instance, if the copula of (U_{t-1}, U_t) is the Gaussian copula and the marginal distribution of U_t is uniform, then the joint distribution of ϵ_{t-1} and ϵ_t is

$$H(u_1, u_2) = C(\Phi(\frac{u_1}{\sqrt{V_E(1 - e^{-2\kappa})}}), \frac{u_2}{\sqrt{V_E(1 - e^{-2\kappa})}}); \alpha) = \Phi_\alpha(\frac{u_1}{\sqrt{V_E(1 - e^{-2\kappa})}}, \frac{u_2}{\sqrt{V_E(1 - e^{-2\kappa})}}).$$

The joint distribution of ϵ_{t-1} and ϵ_t implies that $\{\epsilon_t\}$ is a linear AR(1) process which in turn implies that the process $\{Y_t\}$ is a linear AR(2) process.

Diebold, Gunther, and Tay (1998) suggest that one should examine the correlograms of $\{(U_t - \bar{U})\}^i$ for $i = 1, 2, 3, 4$; they will reveal dependence through the conditional mean, conditional variance, conditional skewness, or conditional kurtosis, see also Thompson (2002). One potential drawback of the Pearson's correlation coefficient is that it is not invariant to strictly increasing transformations. In the above example, this means that the serial correlation coefficient of $\{U_t\}$ is not the same as that of $\{\epsilon_t\}$. Copula-based dependence measures such as Kendall's tau and Spearman's rho, see Joe (1997) and Nelsen (1999) are known to be invariant to strictly increasing transformations of random variables and hence may be used as well.

3 Evaluating Density Forecasts

In this section, we develop two classes of tests: one tests the null hypothesis of correct density forecasts, or equivalently, of the i.i.d. uniformity of the U_t's; the other tests the null of correct specification of the dependence structure of the forecasting model, or the serial independence of $\{U_t\}$, regardless of the specification of the marginal of U_t.

3.1 Tests for the Correct Density Forecasts

Let $G(u) = G(u;\beta)$ be a parametric marginal distribution function such that there exist parameter values β^* and β_0 satisfying $G^*(u) = G(u;\beta^*)$ and $G(u;\beta_0) = u$ for any $u \in [0,1]$. In this case, the true joint distribution of (U_{t-1}, U_t) is of a parametric form:

$$H^*(u_1, u_2) \equiv C(G(u_1;\beta^*), G(u_2, \beta^*); \alpha^*). \tag{3.1}$$

Under the given parametric copula model for the dependence structure in $\{U_t\}$, independence corresponds to $\alpha^* = \alpha_0$ for a specific value α_0 (i.e. $C(v_1, v_2; \alpha_0) = v_1 v_2$). For instance, if the copula is Gaussian, then $\alpha_0 = 0$. Hence, the null hypothesis of i.i.d. uniformity of $\{U_t\}$ is equivalent to

$$H_0 : \ \beta^* = \beta_0, \alpha^* = \alpha_0.$$

The log-likelihood function is given by

$$L(\alpha, \beta) = \frac{1}{n}\sum_{t=1}^{n} \log g(U_t;\beta) + \frac{1}{n}\sum_{t=2}^{n} \log c(G(U_{t-1};\beta), G(U_t;\beta);\alpha), \tag{3.2}$$

where $g(u;\beta)$ is the density function of the marginal distribution $G(u;\beta)$ and $c(v_1, v_2;\alpha)$ is the density function of the copula $C(v_1, v_2; \alpha)$. Noting that α and β may have common elements such as in **Example 1**, we let τ_0 and τ^* respectively be the vectors of the distinct elements of $(\alpha_0', \beta_0')'$ and $(\alpha^{*'}, \beta^{*'})'$. In addition, let $\hat{\tau} = \arg\max_\tau L(\alpha, \beta)$ denote the MLE of τ^*. Then any hypothesis regarding the value of the parameter τ^* can be tested by the LR test. In particular, the LR test for $\tau^* = \tau_0$ is test for the correct density forecasts and it extends the one in Berkowitz (2001) in two directions: First, the copula function is not necessarily Gaussian; Second, the marginal distribution G^* can be of any specific form as long as it includes the uniform distribution as a special case.

To implement the above tests, one needs to choose a class of parametric distributions $G(u;\beta)$ such that $G(u, \beta_0) = u$ for some β_0. One important class of such distributions is given by

$$g(u;\beta) = C(\beta)\exp\{\sum_{i=1}^{k} \beta_i \pi_i(u)\}, \ 0 < u < 1,$$

where k is a positive integer, $\beta = (\beta_1, \ldots, \beta_k)'$, and $\{\pi_i(u)\}$ is the set of orthonormal polynomials on the uniform distribution (the normalized Legendre polynomials) and $C(\beta)$ is a constant depending on β, introduced to ensure that the probability density function integrates to one. Obviously, $g(u; \beta) = u$ when $\beta = 0$.

3.2 Robust Tests for the Dependence Structure of the Forecasting Model

When the marginal distribution G^* is completely unknown, the true joint distribution of U_{t-1} and U_t is of a semiparametric form: $H^*(u_1, u_2) = C(G^*(u_1), G^*(u_2); \alpha^*)$. The null hypothesis of correct specification of the dependence structure of the forecasting model, or equivalently the serial independence of $\{U_t\}$, is equivalent to

$$H_0' : \alpha^* = \alpha_0.$$

Tests for H_0' based on the semiparametric model are thus robust to the choice of the marginal distribution G^*.

In this case, we estimate α^* by a pseudo-MLE $\tilde{\alpha}$ defined as

$$\tilde{\alpha} = \text{argmax}_\alpha \tilde{L}(\alpha), \quad \tilde{L}(\alpha) = \frac{1}{n} \sum_{t=2}^{n} \log c(G_n(U_{t-1}), G_n(U_t); \alpha), \tag{3.3}$$

where $G_n(\cdot)$ is the rescaled empirical distribution function defined as

$$G_n(u) = \frac{1}{n+1} \sum_{t=1}^{n} I\{U_t \le u\}. \tag{3.4}$$

The estimator $\tilde{\alpha}$ is introduced in Genest, Ghoudi, and Rivest (1995) for the case where n i.i.d. observations are available from a bivariate distribution. Chen and Fan (2002) establish conditions under which $\tilde{\alpha}$ is $\sqrt{n}$-consistent and asymptotically normally distributed for the class of copula-based time series models specified in Assumption 1.

Let $\mathcal{A} \subset R^d$ be the parameter space and assume $\alpha^* \in int(\mathcal{A})$. In addition, let $l(v_1, v_2, \alpha) = \log c(v_1, v_2, \alpha)$, $l_\alpha(v_1, v_2, \alpha) \equiv \frac{\partial l(v_1, v_2, \alpha)}{\partial \alpha}$ and $l_{\alpha,j}(v_1, v_2; \alpha) \equiv \frac{\partial^2 l(v_1, v_2; \alpha)}{\partial v_j \partial \alpha}$ for $j = 1, 2$. Assuming the conditions in Chen and Fan (2002) are satisfied, the following corollary of their Proposition 3.3 will be used to construct our tests.

Corollary 3.1 *Under Assumption 1 and conditions for Proposition 3.3 in Chen and Fan (2002), if H_0' holds, then $\sqrt{n}(\tilde{\alpha} - \alpha^*) \to N(0, B^{-1})$, where $B = E[l_\alpha(V_1, V_2, \alpha^*) l_\alpha'(V_1, V_2, \alpha^*)]$ in which $V_t = G^*(U_t)$.*

Proof: Proposition 3.3 in Chen and Fan (2002) states that $\sqrt{n}(\tilde{\alpha} - \alpha^*) \to N(0, B^{-1} \Sigma B^{-1})$ in distribution, where $\Sigma = \lim_{n\to\infty} Var(\sqrt{n} A_n^*)$ in which

$$A_n^* \equiv \frac{1}{n-1} \sum_{t=2}^{n} [l_\alpha(V_{t-1}, V_t, \alpha^*) + W_1(V_{t-1}) + W_2(V_t)]$$

$$W_1(V_{t-1}) \equiv \int_0^1 \int_0^1 [I\{V_{t-1} \leq v_1\} - v_1] l_{\alpha,1}(v_1, v_2; \alpha^*) c(v_1, v_2; \alpha^*) dv_1 dv_2$$
$$W_2(V_t) \equiv \int_0^1 \int_0^1 [I\{V_t \leq v_2\} - v_2] l_{\alpha,2}(v_1, v_2; \alpha^*) c(v_1, v_2; \alpha^*) dv_1 dv_2.$$

Under Assumption 1, $\{U_t\}$ is an independent process if and only if the true dependence parameter α^* is such that $C(v_1, v_2; \alpha^*) = v_1 v_2$ and hence $c(v_1, v_2; \alpha^*) = 1$. By making use of this, one can easily verify that $E[l_{\alpha,1}(v_1, V_2, \alpha^*)] = 0$ which yields $W_1(V_{t-1}) = 0 = W_2(V_t)$ under independence. Hence under independence,

$$\Sigma = Var\{l_\alpha(V_1, V_2, \alpha^*)\} + 2Cov\{l_\alpha(V_1, V_2, \alpha^*), l_\alpha(V_2, V_3, \alpha^*)\} = Var\{l_\alpha(V_1, V_2, \alpha^*)\},$$

where the first equality follows from the independence of the V_t's and the second one follows from the fact that under independence of the V_t's, $E[l_\alpha(V_1, V_2, \alpha^*)|V_i] = 0$ for $i = 1, 2$.

□

This corollary extends the asymptotic efficiency result of the α_0 estimate based on i.i.d. observations in Genest, Ghoudi, and Rivest (1995) to the case of time series observations $\{U_{t-1}, U_t\}$.

Assuming α_0 lies in the interior of the parameter space, we will develop two asymptotic tests for H_0': a pseudo Wald test and a pseudo LR test.

To construct the pseudo Wald test, we need to estimate the asymptotic variance of $\tilde{\alpha}$ under H_0'. If the copula is Gaussian, then under H_0', $\alpha^* = \alpha_0 = 0$. It follows from the Gaussian copula example in Chen and Fan (2002) that under H_0', $\sqrt{n}\tilde{\alpha} \sim N(0, 1)$. Following Klaassen and Wellner (2001), one can also show that the estimator $\tilde{\alpha}$ is asymptotically equivalent to the normal scores correlation coefficient, see (2.3). This suggests a simple consistent test for H_0' when the copula is Gaussian: First form $\tilde{V}_t^* = \Phi^{-1}(G_n(U_t))$; Second, estimate α by the OLS estimator in the regression of $\tilde{V}_{t-1}^*$ on $\tilde{V}_t^*$; Finally reject H_0' at 5% significance level if the estimator falls outside of the interval $[-2n^{-1/2}, 2n^{1/2}]$.

For a specific non-Gaussian copula, it may be possible to obtain an explicit expression for the asymptotic variance of $\tilde{\alpha}$ by Corollary 3.1. However, noting that the asymptotic variance of $\tilde{\alpha}$ does not depend on the marginal distribution G^*, it can thus be estimated easily by simulation. More specifically, recall under H_0', $B = E[l_\alpha(V_1, V_2, \alpha_0) l_\alpha'(V_1, V_2, \alpha_0)]$, where V_1 and V_2 are independent uniform random variables on the interval $(0, 1)$. This suggests that one consistent estimator of B under H_0' should be

$$\hat{B} = \frac{2}{N(N-1)} \sum \sum_{1 \leq i < j \leq N} l_\alpha(V_i, V_j, \alpha_0) l_\alpha'(V_i, V_j, \alpha_0), \tag{3.5}$$

where $\{V_i\}_{i=1}^N$ is a large number (N) of i.i.d. uniform random variables on the interval $(0, 1)$. In summary,

Proposition 3.2 *Let* $\widetilde{W} = (\tilde{\alpha} - \alpha_0)' \hat{B} (\tilde{\alpha} - \alpha_0)$ *and conditions for Corollary 3.1 hold. Then under* H_0', $n\widetilde{W} \to \chi^2_{[d]}$ *in distribution.*

Liang and Self (1996) show that in general the Pseudo-LR test does not follow an asymptotic χ^2 distribution due to the inefficiency of the pseudo-likelihood estimator. However, as shown in Corollary 3.1, under H_0', the pseudo-likelihood estimator $\tilde{\alpha}$ is asymptotically efficient. This in turn verifies the following result.

Proposition 3.3 *Let conditions for Corollary 3.1 hold. Then under H_0', $2n[\tilde{L}(\tilde{\alpha}) - \tilde{L}(\alpha_0)] \to \chi^2_{[d]}$ in distribution.*

4 Concluding Remarks

In this paper, we demonstrate that the copula approach to evaluating density forecasts provides flexibility in designing tests that have power against a wide range of alternative processes. In addition it allows us to construct simple tests for the correct specification of the dependence structure of the forecasting model that are robust to misspecification of the marginal distribution. Together with the test in Bai (2003), such tests will be useful in detecting the failure of the forecasting model; its dependence structure or the marginal distribution, when it is rejected.

This paper makes two simplifying assumptions: the forecast model is completely known, and α_0 (the value leads to the independence copula) is in the interior of the parameter space. These are restrictive assumptions and are adopted in this paper in order not to complicate the exposition of the main idea. We are currently working on relaxing these assumptions.

References

[1] Bai, J. (2003), "Testing Parametric Conditional Distributions of Dynamic Models," *Review of Economics and Statistics*, forthcoming.

[2] Berkowitz, J. (2001), "Testing Density Forecasts, With Applications to Risk Management," *Journal of Business & Economic Statistics* 19(4), 465-474.

[3] Bouyé, E. , N. Gaussel, and M. Salmon (2002), "Investigating Dynamic Dependence Using Copulae," Manuscript, Financial Econometrics Research Center.

[4] Chen, X. and Y. Fan (2002), "Semiparametric Estimation of Copula-Based Time Series Models," Manuscript.

[5] Clements, M. P. and J. Smith (2000), "Evaluating the Forecast Densities of Linear and Nonlinear Models: Applications to Output Growth and Unemployment," *Journal of Forecasting* 19, 255-276.

[6] Corradi, V. and N. R. Swanson (2001), "Bootstrap Specification Tests With Dependent Observations and Parameter Estimation Error," Working paper.

[7] Darsow, W., B. Nguyen, and E. Olsen (1992), "Copulas and Markov Processes," *Illinois Journal of Mathematics* 36, 600-642.

[8] Diebold, F. X. , T. Gunther, and A. S. Tay (1998), “Evaluating Density Forecasts, with Applications to Financial Risk Management,” *International Economic Review* 39, 863-883.

[9] Diebold, F. X. , J. Hahn, and A. S. Tay (1999), “Multivariate Density Forecast Evaluation and Calibration in Financial Rist Management: High-Frequency Returns on Foreign Exchange,” *Review of Economics and Statistics* 81, 661-673

[10] Diebold, F. X. , A. S. Tay, and K. Wallis (1999), “Evaluating Density Forecasts of Inflation: the Survey of Professional Forecasts”, in R. Engle and H. White (eds), *Festschrift in Honor of C.W.J. Granger*, Oxford: Oxford University Press.

[11] Elerian, O. , S. Chib, and N. Shephard (2001), “Likelihood Inference for Discretely Observed Non-linear Diffusions,” *Econometrica* 69, 959-993.

[12] Embrechts, P. , A. McNeil, and D. Straumann (1999), “Correlation and Dependence Properties in Risk Management: Properties and Pitfalls,” in M. Dempster, ed., *Risk Management: Value at Risk and Beyond*, Cambridge University Press.

[13] Gagliardini, P. and C. Gourieroux (2002): “Duration Time Series Models with Proportional Hazard”, working paper, CREST and University of Toronto.

[14] Genest, C. , K. Ghoudi, and L.-P. Rivest, “A Semiparametric Estimation Procedure of Dependence Parameters in Multivariate Families of Distributions,” *Biometrika* 82 (3), 543-552.

[15] Hong, Y. and H. Li (2002), “Nonparametric Specification Testing for Continuous-Time Models with Application to Spot Interest Rates,” Cornell University.

[16] Inoue, A. (1999), “A Conditional Goodness-of-fit Test for Time Series,” Working paper.

[17] Joe, H. (1997), *Multivariate Models and Dependence Concepts*, Chapman & Hall/CRC.

[18] Klaassen, C. and J. Wellner (2001), “Efficient Estimation in the Bivariate Normal Copula Model: Normal Margins are Least-Favorable,” *Bernoulli* 3, 55-77.

[19] Liang, K.-Y. and S. G. Self (1996), “On the Asymptotic Behavior of the Pseudolikelihood Ratio Test Statistic,” *J. R. Statist. Soc.* B 58, 785-796.

[20] Nelsen, R. B. (1999), *An Introduction to Copulas*, Springer.

[21] Patton, A. (2001), “Modeling Time-Varying Exchange Rate Dependence Using the Conditional Copula,” Manuscript, UCSD.

[22] Rosenblatt, M. (1952), “Remarks on a Multivariate Transformation,” *Annals of Mathematical Statistics* 23, 470-472.

[23] Sklar, A. (1959), “Fonctions de r’epartition ’a n dimensionset leurs marges,”*Publ. Inst. Statis. Univ. Paris* 8, 229-231.

[24] Thompson, S. B. (2002), “Evaluating the Goodness of Fit of Conditional Distributions, with an Application to Affine Term Structure Models,” Harvard University.

[25] Vasicek, O. (1977), “An Equilibrium Characterization of the Term Structure,” *Journal of Financial Economics* 5, 177-188.