

Statistics for Risk Management: Frontier Issues

Statistica per la Gestione del Rischio: Questioni di Frontiera

Umberto Cherubini
Università di Bologna
cherubin@polyhedron.it

Riassunto: Vengono passati in rassegna gli argomenti di frontiera della ricerca nel campo della valutazione del rischio. Il vasto coinvolgimento degli intermediari finanziari nei mercati dei prodotti derivati ha posto al centro dell'attenzione le tecniche di analisi e simulazione di distribuzioni non normali. L'utilizzo di prodotti non lineari consente infatti la costruzione dell'intera distribuzione dei rendimenti, anziché la determinazione dei soli primi due momenti come nelle strategie tradizionali. Vengono passate in rassegna le tecniche proposte e/o utilizzate a livello univariato e multivariato per l'analisi delle distribuzioni dei rendimenti. A livello univariato, le tecniche sono articolate in una forte compressione dei dati e nella scelta di un modello statistico da calibrare, sulla base di informazione storica e/o implicita nei prezzi dei derivati. A livello multivariato, l'impiego delle funzioni di copula permette l'analisi separata di distribuzioni marginali e struttura di dipendenza.

Keywords: Risk Management, Historical Simulation, Copula Functions

1. Financial Derivatives Management: the Challenge to Statistics

Starting with the 90s, the problem of risk evaluation and risk management has become one of the core topics of discussion among academics, practitioners and regulators. This focus has been mainly motivated by the radical change that the intermediation business has undergone over the last decades. Nowadays, the banking business has progressively shifted its focus from the intermediation of funds to the intermediation of risks. The main engine behind this change is represented by the growing involvement of financial intermediaries in the derivative markets (forward, futures and options, with more and more sophisticated features). These are not only used for the management of proprietary portfolios but also to build new financial products for the intermediation business, such as in the so called structured finance market. The most recent innovations in the derivatives market have also changed the traditional intermediation business by enabling the banks to modify their credit risk burden by transferring all or part of it to the market or other institutions: it is all the more common practice for banks to resort to tools such as securitisation and credit derivatives to buy and sell "protection" on credit portfolios. These innovations have also created new kinds of financial intermediaries which specialise in the use of derivatives for the "yield enhancement" of investment strategies: the main example is given by the so called hedge funds. While it is far beyond this analysis to give an account of the use of derivatives in finance, it may be useful to give a synthetic assessment of how this business calls for a deep involvement of statistical techniques and resources. Without getting into the details as to how, for example, a hedge fund works, we may give a synthetic view of its operation by stating

that it is using derivatives to design a proper probability distribution for its returns. So, if one would like disentangle the role of financial engineers and the statisticians in a financial institution, he could say that the former price derivatives and use them to design a specific probability distribution, while the latter provide a representation of the distribution, well beyond the standard classical assumption of normality. It must be clear that non-normality of the distribution is no longer an empirical issue, but it is an investment instrument used in the market. The world of finance is then seeking, particularly from statistics, tools to cope with non-normal distributions: open questions, that we address in this paper, refer to estimation and simulation strategies, particularly for the evaluation of risk, and the issue of representing dependence between non-normal variables. The worst nightmare of any trader is the joint movement, in front of changing market conditions, of the wide variety of options held in his portfolio. Since such positions, as their payoffs, are intrinsically non-linear, linear correlation does not seem to be a well suited statistics to represent the comovement among their values. Recent literature in mathematical finance spots in copula functions theory a promising technique that statistics can provide to overcome, or at least to address, these problems. The reason for this interest rests on the fact that copula functions provide a useful opportunity to separate the problem of multivariate distribution representation in the specification of the marginals and the dependence structure.

The interest in probability distributions is then the common ground of work of statisticians, financial engineers and risk managers. Because of the use derivatives, the task has grown more complex, well beyond the interest in mean and variance of the earlier times, when only linear products were used in the investment strategies and the assumption of gaussian distribution of asset returns was taken for granted, at least as a good approximation.

As for risk management, the practice of derivatives trading has lead to the construction of new tools for the evaluation and measurement of risk. Since the beginning of the 90s, a new measure, called Value-at-Risk, has gathered momentum and is nowadays the most widely used tool. Technically speaking, Value-at-Risk is simply a percentile of the conditional distribution of returns over a given period. In other words, assuming one is bound to hold a portfolio for a given period, Value-at-Risk gives a measure of how much value he can lose over that time at a given probability level. The length of time one is assumed to hold the portfolio is known as the “unwinding period” and reflects the idea that one needs some time to unbundle his position and get out from the market: so, the unwinding period is meant to account for the limited liquidity of the market, i.e. the difficulty to find counterparties to which one could sell the position. In this respect, this measure is directly derived from the derivatives market, and in particular those for futures, where brokers require customers to deposit the cash equivalent of a percentage of the value of the position (known as margin) to have time to close the position is the market moves against it. In technical terms, Value-at-risk (VaR) is defined as

$$Prob(V(\tau) - V(t) \leq - VaR) \equiv \alpha \quad (1)$$

where V is the value of the position, $\tau - t$ is the unwinding period and α is the probability level.

The Value-at-Risk concept has recently been criticised on the grounds that it may, at least in principle, lead to questionable results. More explicitly, it can happen that the Value-at-Risk figure for a portfolio is greater than the sum of the Value-at-Risk figures

of the portfolio components, while one would indeed expect the former to be weakly smaller than the latter because of the diversification effect. In technical terms, Value-at-Risk is not a sub-additive measure, a condition which features among the requirements (axioms) that describe a “coherent” risk measure, according to the definition in the paper by Artzner, Delbaen and Elber (1997). A large literature has been growing to provide new measures able to overcome this flaw: among these, the most serious challenge to Value-at-Risk seems to be represented by the so-called Expected Shortfall (ES) measure, which could be defined as

$$ES \equiv E[V(\tau) - V(t) | V(\tau) - V(t) \leq -VaR] \quad (2)$$

The expected shortfall measure represents the expected loss, conditional on the loss being greater than the percentile representing the Value-at-Risk. It is immediate to see that this measure is more accurate than Value-at-Risk, because it directly refers to the loss of value: in fact, we may think of two portfolios with the same Value-at-Risk but different expected shortfall figures, in which case one would naturally choose the lower one. On top of that, unlike Value-at-Risk, expected shortfall may be proved to be sub-additive and may be used as a constraint in a portfolio problem optimisation without making the problem itself untractable, as it happens with Value-at-Risk constraints. Despite these advantages, however, expected shortfall has not yet substituted Value-at-Risk as the most widely used measure of risk, and it still plays a largely secondary role in the risk management practice. There are many reasons for this, well beyond the natural sluggishness of technology changes. While getting into the details of discussion would involve a large digression in the field of financial mathematics, which is beyond our scope here, it could be useful to notice that the advantages of the expected shortfall measure would not be substantial in a market in which the returns and losses are normally distributed: indeed, in this case, as in the more general case of elliptic distributions, the Value-at-Risk measure itself may be proved to be sub-additive, like the expected loss. Under gaussian returns, then, both measures are found to give the same risk-management directions. It is not the same for more complex distributions, that can be induced on the returns and losses by the presence of “optionality” in the portfolio: it is the case of the structured finance products, but also of any product with credit risk. It is intuitive, also to a reader new to the very basic concepts of finance, that the presence of credit risk, which means a possibility to lose whole or part of the capital invested, affects the distribution of returns by making it substantially skewed toward losses. The growing attention to credit risk provides this effect with a pervasive relevance throughout most of the operations of the intermediaries, and it will possibly lead to further confrontation between the Value-at-Risk figure and new-born challengers, such as expected shortfall. Again, however, the theory of probability distributions, beyond the standard gaussian one, will represent the main topic of the future.

2. Fat tails in one dimension: simulation techniques

The problem of non-gaussian distribution of returns has posed interesting questions to statisticians and quantitative analysts in general. One of the most frequently discussed issue in banks is the question of leptocurtosis, i.e. the “fat tails” problem. On top of that,

statistics has also spread to the traders' jargon, where is now common to hear terms like "trading the skew". As far as risk management applications are concerned, we can distinguish three basic issues. The first is how to reduce the dimension of the problem in order to address the question of non-normality with more powerful tools. The second issue has to do with the kind of information and the data to be used for the construction of the risk measure. The last crucial topic is which statistical model to use in substitution for the normal one. Across all of these choices, the work of a risk manager may be seen more as an art than a science, and the success of setting up an effective risk management unit is linked to the way in which these choices are blended.

Data compression techniques. As far as the need to reduce the dimension of the analysis is concerned, we may distinguish three different strategies. The first, which is widely used in historical simulation applications, is to use the actual portfolio of a desk or an intermediary institution as the tool to reduce the dimensionality of the problem. In plain words, the portfolio is re-evaluated using a sample of data drawn from past history and the time series of the returns that is obtained is then evaluated using advanced techniques that are available for uni-variate problems. The second choice is to select a limited number of relevant risk factors, such as the stock market or the foreign exchange, with respect to which a matrix of sensitivities, or factor loadings, is periodically estimated. The simulation is then performed on the time series of these risk factors and the effect on the value of the positions is reconstructed using the factor loadings. The third way to go is to resort to the traditional data compression techniques provided by statistics, such as principal components and factor analysis. Obviously, which is the right method to use crucially depends on the application: on one hand, using the actual portfolio as a data compression tool enables to use the most advanced techniques to analyse the distribution of returns, at the cost of re-running the whole procedure any time that the portfolio is reallocated; on the other hand, using more standard statistical data compression techniques could separate the time of specification of the model from that of reallocation of the portfolio, at the cost of simplifying assumptions concerning the dependence structure among the risk factors. If finally one uses observed risk factors in the analysis, he could extract more information concerning the sources of risk, to the great advantage of the hedging decision making process. So, choosing a specific data compression strategy means striking the right balance between all of these pros and cons.

Implicit and Historical Information. The second issue has to do with the kind of information one wants to use in risk management. This kind of issue is quite peculiar to financial applications, and it requires some familiarity with quantitative finance techniques to be grasped. To keep things simple, it suffices to say that in finance we use two different probability measures: one, which we call "objective", refers to the distribution of returns that can be observed and estimated from time series of market data, while the other, which we call "risk-neutral measure" (or also equivalent martingale measure), emerges directly from the specific pricing technique that rules out arbitrage opportunities and that is the basic principle of modern finance applications. Loosely speaking, under the "objective" measure the expected return of any asset in the economy may be decomposed in a risk-free part and a risk-premium. More technically, we may write

$$E[r_i] = r_f + \lambda' \sigma_i \quad (3)$$

where r_f is the risk free rate and $\lambda'\sigma_i$ is the inner product of the vector of market prices of the risk factors (λ) and the sensitivities of any asset i to them (σ_i). The fact that no arbitrage opportunities is available on the market is warranted by the fact that the factors market prices of risk are the same across all of the $i = 1, 2, \dots, N$ assets traded in the economy. Under the risk-neutral probability measure Q we have instead

$$E_Q[r_i] = r_f \quad (4)$$

and all of the risky assets yield the risk-free rate. The latter probability space is widely used in finance to price derivatives. This brief digression leads us straight to our problem. If we are interested in estimating the probability distribution of returns, we may either choose a time series of prices of the asset, or a cross-section of prices of derivatives written on it. To be more precise, there are two questions involved here: the first is whether risk should be evaluated under the objective or the risk-adjusted measure, the second is which information to use to calibrate the chosen probability distribution. As for the former question, a detailed discussion would call for a deep understanding of the financial issues of arbitrage free pricing and firm evaluation, so we do not address it here: it suffices to say that the vast majority (say the world except this author) follows the common wisdom of using the risk-adjusted measure for pricing and the objective probability for risk management. As for the latter question, the choice is between the analysis of the time series of returns and the analysis of the current derivatives quotes. Again, any choice has pros and cons. Using time series analysis enables to cover a large number of markets, but it may be flawed by the limited ability to account for structural breaks, which is exactly what risk managers are more interested in, as well as abrupt changes of the sample space itself, like in the so-called peso problem. On the opposite, using implied information from derivatives overcome most of these flaws, providing a forward looking estimate consistent with market quotes: however, this comes at the cost of restrictive assumptions that have to be imposed to back-out the objective probability from the risk-neutral one, a procedure that is by no means easy to set up. In fact, extracting the objective probability from the risk-neutral one calls for the specification of the dynamics of market prices of risk (λ), which pulls back the problem to a time series dimension. Moreover, obtaining reliable implied information from derivatives prices is possible only for very few well developed markets, and even in these cases the information obtained may carry substantial “noise” due to liquidity problems with the market and the corresponding significance of the prices used for calibration. Once again, setting up the right system is a matter of choice and there is no doubt that any contribution on new techniques to compare and possibly combine historical and implied information would have a relevant fall out on the risk management practice.

New Models for the Distribution of Returns. The third issue is particularly interesting for statisticians, and has to do with the specification and estimation of the distributions of returns. In principle, we may think of three different approaches to the problem. The first is to try to specify a new distribution for the returns, which could be able to represent the probability law better than the gaussian one. The second is to try to simulate the distribution using historical data. The third is to rely on worst case scenarios for the distribution.

As for the need to specify a new distribution, one could for example resort to a Student-t distribution, as proposed by Wilson (1993) in a model with estimation risk in which the

agents are not assumed to know exactly the volatility and correlation of assets, but only have a prior on their distribution: if the prior is inverted-gamma and the conditional distribution of the returns is gaussian, Student-t distribution emerges as the right choice. Another choice would be to work with flexible functional forms. An example is the so called generalised Pearson family (Cobb, Coppstein and Chen, 1983). Zellner and Highfield (1988) show that the probability distributions in this family are obtained by choosing the maximal entropy density under constraints on the first k moments and provide an algorithm for computation, based on an expansion in polynomial terms. It is important to notice that this choice of expansion in polynomials around a key density also corresponds to a pricing model for options, due to Jarrow and Rudd (1982), that could be used to calibrate the probability on market quotes rather than on time series data. If the latter route is followed instead, Gallant and Tauchen (1989) seminonparametric approach provide a natural and suitable extension allowing for more complex key densities, such as that of an ARCH model.

This is not the only flexible functional choice provided and used. Another opportunity is to resort to the so-called Johnson (1949) family of distributions, and in particular to the subset of distribution obtained by translating the log-normal distribution. A FORTRAN algorithm for the calibration of this model is provided in Hill et al. (1976) (see the software included in Cherubini and Della Lunga, 2001 for a translation of the code to VBA). The idea of working with a family of distributions including the log-normal one is particularly useful for asset prices, as we know that this corresponds to a geometric brownian motion dynamics of the price, that is used in the most famous option pricing model, Black and Scholes (1973). Of course the log-normal model for prices corresponds to normally distributed returns, but a possible adjustment that can be used to account for deviations from normality is translation of the log-normal distribution, as proposed by Johnson half a century ago. This pricing technique is used in many option pricing softwares around, and this of course would give the choice of Johnson family an advantage. Moreover, along this way, recent developments in the pricing literature could be exploited to develop the risk-management counterpart: for example, Brigo and Mercurio (2002) resort to a mixture of translated log-normal distributions to fit option prices, also recovering the corresponding continuous time dynamics.

If one does not want to commit himself to the choice of any particular functional form, even among the flexible ones, he can resort to historical simulation. Under the simplest possible application, the historical simulation structure consists of a histogram of the portfolio value, simulated using past history as the set of possible scenarios. Early empirical analyses, mainly conducted in central banks (see Mahoney, 1995 and Jackson et al. 1995) found that even this simple procedure improves substantially the accuracy of the measure over the assumption of normal returns. Of course, this techniques also run into some problems. The first is dependence among the observations across time: if returns are not i.i.d. the straight application of historical simulation can lead to a biased measure of risk. To solve this problem, Barone-Adesi et al. (1998a,b) propose to filter data by using a GARCH model and resort to bootstrap techniques both to forecast future volatility and to simulate the conditional distribution over the unwinding period. As an example, say we use a GARCH(1,1) model to filter the returns, so that they are assumed to be generated by the model

$$\begin{aligned} r_t &= \varepsilon_t \\ \sigma_t^2 &= \omega + \alpha_1 \varepsilon_{t-1}^2 + \beta_1 \sigma_{t-1}^2 \end{aligned} \quad (5)$$

Based on this model, the time series of residuals ε_t and estimated volatilities σ_t are used to construct the standardised residuals $e_t \equiv \varepsilon_t / \sigma_t$. A bootstrap sampling is then applied to standardised residuals. If the unwinding period consists of n days, n standardised residuals are drawn from the time series for each iteration. The first is used to simulate the next period shock $z_{t+1} = e_1 \hat{\sigma}_{t+1}$, where $\hat{\sigma}_{t+1}$ is the volatility forecast, and to forecast volatility at time $t + 2$, and so on until all of the n standardised residuals are used. The process is repeated for a large number of iterations and the distribution of values of the portfolio over n days is finally tabulated. Nowadays, this algorithm represents the most advanced example of historical simulation application. Nevertheless, it still is subject to some flaws. The first is that the chance of getting i.i.d. residuals is due to the particular model specification for the returns dynamics. The second problem, which is common to all of the historical simulation applications, is that the set of possible scenarios is limited by what has happened over the past: this makes the model, and historical simulation in general, hardly applicable to periods of sudden structural changes and to emerging markets that are endowed with a short time series of data. This problem is common to both historical simulation and stress testing techniques based on market events: in both cases the shortcoming could be overcome by resorting to Montecarlo techniques or machine learning algorithms (genetic algorithms and the like), aimed at looking for the worst possible scenarios over a set of states possibly wider than those proposed by history.

The issue of looking for the worst possible scenario leads us to the third set of options. A possible idea is to look at a large set of possible distributions, and to concentrate on the behavior of the tails. A stream of literature in statistics, known as extreme value theory, has investigated this problem. For most distributions the density of the minimum values attained are represented by

$$g(z) = \frac{1}{\alpha} \exp\left(-\frac{z-\beta}{\alpha}\right) \exp\left[-\exp\left(-\frac{z-\beta}{\alpha}\right)\right] \quad (6)$$

$$f(z) = \frac{1}{\alpha\tau} \left(\frac{z-\beta}{\alpha}\right)^{\frac{1}{\tau}-1} \exp\left[-\left(\frac{z-\beta}{\alpha}\right)^{\frac{1}{\tau}}\right] \quad (7)$$

where $g(z)$ is known as Gumbel and $f(z)$ is the Fréchet distribution, while parameters β and α represent location and dispersion. It can be noticed that the two distributions differ because of the presence of a parameter τ which represents “fat-tails”. Longin (2000) proposes to calibrate these distributions on the time series of portfolio returns. In this way, the Value-at-Risk figure is calibrated using information on the tails only.

As a final kind of approach, one could consider to assume a set of probability measures, and to apply worst case analysis to the choice of the distributions themselves. This approach is indeed well suited to address a second dimension of risk management, which has to do with market liquidity. To take the simplest example, consider a very illiquid market of options, quoting a very large bid-ask spread for implied volatilities and option prices. Say you want to extract the probability distribution from such prices: which price would you consider? Of course, bid and ask quotations would lead you to

pretty different assessments of the distribution. Given the portfolio, the idea is to look for the probability distribution that would give you the lowest value. Cherubini and Della Lunga (2001b) apply this technique to the evaluation of credit risk, showing that this approach can provide a reasonable solution to the underperformance of structural credit risk models, i.e. option based models for the evaluation of corporate securities. In fact, because credit risk can be seen, the reader not familiar to corporate finance would take our word for it, as a hidden option written on the asset balance sheet of a firm, credit risk applications provide the ideal example of extremely illiquid markets.

We have seen that at the univariate level there is a large choice of techniques that can be applied to analyze and calibrate the distribution. Very often, however, we are interested in the measure of the dependence structure in a multivariate non-normal setting. This problem gets us to the use of copula functions, which represents the very frontier of current research in risk management.

3. The Multivariate Setting: Copula Functions

It is not always a good idea, or it is not always feasible, to compute a single time series for the whole portfolio of a financial institution. For example, Value-at-Risk is computed for each of the different units of the intermediary, and it should be used to allocate the overall capital among them. In this case, we are obviously interested in the dependence structure among the different businesses: if for example all of the businesses are perfectly positively dependent, adverse market conditions will have a serious impact on the institution as a whole, while if some of the businesses are negatively dependent, they could serve as a hedge for some of the others, so that losses on one of them could be balanced against by gains in the others. The same question arises if we want to manage a portfolio of loans, or corporate bonds: again, it matters a lot whether we lend to a set of borrowers whose default probabilities are independent rather than perfectly dependent.

From the financial theory point of view, these are the same old diversification and hedging arguments, which are very easy to handle in a gaussian setting. The technical question is that they are not so easy to handle in the case in which the returns of the different portfolios are not normally distributed. In the example above, this is exactly the case if the desks involved in our capital allocation problem make substantial use of options in their investment strategies. This is again the case if we have a portfolio of credit risks, that from the financial point of view is equivalent to a portfolio of options.

The problems above have caused the research in risk management to discover a new tool, which is borrowed from statistics, and that has the uneasy name of “copula” functions. From a mathematical point of view, a copula function ($C(u,v)$ in the simple bivariate setting) is defined by three requirements. Both the domain of the arguments, u and v , and the image of the copula must be defined in the closed unit interval. It must be $C(0,v) = C(u,0) = 0$, $C(1,v) = v$ and $C(u,1) = u$: we say that the copula function is grounded. Finally, it must be 2-increasing: if $u_1 \leq u_2$ and $v_1 \leq v_2$ then $C(u_2,v_2) - C(u_2,v_1) - C(u_1,v_2) + C(u_1,v_1) \geq 0$. The interest for statistics is based on a result known as Sklar theorem. This result states that given a joint distribution function $F(x,y)$ and marginal distributions $F_1(x)$ and $F_2(y)$ it is always possible to define a copula function $C(u,v)$ such that $F(x,y) = C(F_1(x), F_2(y))$. The opposite result also holds: given a copula function

$C(u,v)$ and two univariate distributions $F_1(x)$ and $F_2(y)$, it is possible to construct a joint distribution by setting $C(F_1(x), F_2(y))$.

The appeal of the method is immediately clear: using copula function representations enables one to address the issue of specification of marginal distributions and the joint distribution separately. It is clear that, once the marginal distributions have been specified, the choice of the copula function contributes to the construction of the joint distribution by specifying the dependence structure. The simplest example of copula that we may think of corresponds to the independence case, and it is the product copula $C_{ind}(F_1(x), F_2(y)) = F_1(x) F_2(y)$. What about the cases of perfect positive and negative dependence? We may exploit a result due to Fréchet, concerning the bounds of joint probability distributions to obtain

$$\max[F_1(x) + F_2(y) - 1, 0] \leq C(F_1(x), F_2(y)) \leq \min[F_1(x), F_2(y)] \quad (8)$$

Copula functions are then required to stay between a lower bound $C_{min}(F_1(x), F_2(y)) \equiv \max[F_1(x) + F_2(y) - 1, 0]$ corresponding to perfect negative dependence and an upper bound $C_{max}(F_1(x), F_2(y)) \equiv \min[F_1(x), F_2(y)]$ corresponding to perfectly positive dependence. So, the choice of copula is directly linked to the specification of the dependence structure of the variables. The simplest way to build a copula is to consider a mixture of the three cases considered, setting

$$C(F_1, F_2) = \beta C_{min}(F_1, F_2) + (1 - \beta - \alpha) C_{ind}(F_1, F_2) + \alpha C_{max}(F_1, F_2) \quad (9)$$

with $\beta, \alpha \geq 0$ and $\beta + \alpha \leq 1$. This is known as the Fréchet family of copulas. Notice that by setting one of the two parameter to zero or their sum to one we encompass the extremal and the product copula, a propriety that is called comprehensiveness. It is intuitive to guess that for other values of the parameters we may represent imperfect dependence. Indeed, copula functions can be directly related to measures of dependence.

Before looking at this relationship, it is useful to digress briefly on the dependence measure that is extensively used in the market and in the applied work in finance. The measure always used to represent comovements between the markets is the linear correlation coefficient. Indeed, using this measure may run into problems if the marginal distributions are allowed to take very general shapes. While the reader is referred to Embrechts et al. (1999) for details, to get an idea of the problem consider that the correlation coefficient between the returns of our two desks may be written as

$$\rho = \frac{\text{cov}(r_1, r_2)}{\sigma_1 \sigma_2} = \frac{1}{\sigma_1 \sigma_2} \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} [F(x, y) - F_1(x)F_2(y)] dx dy \quad (10)$$

If the returns from the two desks are perfectly dependent, the correlation can be calculated by substituting the upper Fréchet bound in the equation: the problem is that it may well be the case that we obtain a correlation figure lower than 1. So, under arbitrary shapes for the marginal distributions, corresponding to non linear relationships between the variables, the correlation measure may not span the whole range from -1 to $+1$. It is to be expected that hedge funds and derivatives desks will sooner or later have to switch to non-parametric measures of dependence, such as Spearman's ρ_s and Kendall's τ .

Unlike linear correlation, these measures do not depend on marginal distributions, just like copula functions don't. It is not surprising, then, that these measures are directly related to copula functions. We have in fact

$$\tau = 4 \int_0^1 \int_0^1 C(u, v) dC(u, v) - 1 \quad (11)$$

for Kendall's τ and

$$\rho_s = 12 \int_0^1 \int_0^1 C(u, v) du dv - 3 \quad (12)$$

for Spearman's ρ_s . It may be checked that by substituting the extremal copulas into (11) and (12) we obtain -1 and $+1$ in the cases of perfect negative and positive dependence respectively. We may also establish a relationship between copula parameters and these non-parametric statistics. In the simple example of the Fréchet family presented above, we have in fact

$$\tau = \frac{(\beta - \alpha)(\beta + \alpha + 2)}{3} \quad (13)$$

$$\rho_s = \alpha - \beta \quad (14)$$

This suggests a first strategy for estimation, by selecting the parameters to fit the non parametric statistics. Another result that is useful for the estimation is that the joint density turns out to be equal to the product of the cross-derivative of the copula times the marginal densities, which is very useful to perform maximum likelihood estimation. Finally, if one would like to calibrate the model according to the behavior in the tails, the copula function could be used to calibrate a tail dependence index.

This new technique of copula function looks like one of the most promising lines of research in risk management and in the modelling of financial time series in general. The obvious application in risk management is to provide a view of the joint probability of loss, beyond the capital allocated on the basis of Value-at-Risk, of different desks or lines of business of the same financial institution as studied in Cherubini and Luciano (2001), or on the probability distribution of the sum of losses, as in Luciano and Marena (2001). Other interesting applications are in the field of credit risk management and credit derivatives (Li, 2000 and Schonbucher and Schubert, 2001). New results are obtained in the field of multivariate option pricing and credit risk in derivatives transactions (Cherubini and Luciano 2001, 2002). As for the specification of the dynamics of financial time series, the use of copula has contributed to more flexible models for the dynamics of dependence across time series, as in the application to foreign exchange provided by Patton (2001), with a conditional copula model and marginal GARCH dynamics, as well as a more general setting of dependence across time, as in the application provided by Bouyé et al. (2001).

References

- Artzner P., Delbaen F. and Eber J.M. (1999) Coherent measures of risk, *Mathematical Finance*, 9, 3, 203-228.
- Barone-Adesi G., Bourgoin F. and Giannopoulos K. (1998a) Don't Look Back, *Risk*, 11, 100-104.
- Barone-Adesi G., Giannopoulos K. and Vosper L. (1998b) VaR Without Correlations for Non-Linear Portfolios, *Journal of Futures Markets*, 19, 583-602
- Black F. and Scholes, M. (1973) The pricing of options and corporate liabilities, *Journal of Political Economy*, 81, 637-654.
- Brigo D. and Mercurio F. (2002) Lognormal mixture dynamics and calibration to market volatility smiles, *International Journal of Theoretical and Applied Finance*, forthcoming.
- Bouyé E., Gaussel N. and Salmon M. (2000) Investigating dynamic dependence using copulas, *working paper*, Financial Econometrics Research Center, City University Business School.
- Cherubini U. and Della Lunga G. (2001a) *Il Rischio Finanziario*, McGraw Hill, Milan
- Cherubini U. and Della Lunga G. (2001b) Liquidity and credit risk, *Applied Mathematical Finance*, 8, 2, 79-95.
- Cherubini U. and Luciano E. (2001) Value-at-Risk trade-off and capital allocation, *Economic Notes*, 30, 2.
- Cherubini U. and Luciano E. (2002) Bivariate option pricing with copulas, *Applied Mathematical Finance*, forthcoming.
- Cherubini U. and Luciano E. (2002) Pricing vulnerable options with copulas, *Icer working paper*, Turin.
- Cobb L., Koppstein P. and Chen N.H. (1983) Estimation and moment recursion for multimodal distributions of the exponential family, *Journal of the American Statistical Association*, 78, 38, 124-130
- Embrechts P., McNeil, A. and Straumann, D. (1999) Correlation and Dependence in Risk Management: Properties and Pitfalls, *ETHZ working paper*.
- Gallant A.R. and Tauchen G. (1989) Semiparametric estimation of conditionally constrained heterogeneous processes: asset pricing applications, *Econometrica*, 57, 1091-1120.
- Jackwerth J. and Rubinstein M. (1996) Recovering probability distributions from option prices, *Journal of Finance*, 51, 1611-1631.
- Jackson P., Maude D. and Perraudin W. (1995) Value-at-Risk techniques: an empirical study, *Risk Measurement and Systemic Risks, Proceedings of a joint central bank research conference*.
- Jarrow R. and Rudd A. (1982) Approximate options valuations for arbitrary stochastic processes, *Journal of Financial Economics*, 10, 347-369.
- Johnson N. L. (1949) Systems of frequency curves generated by methods of translation, *Biometrika*, 36, 149-176.
- Hill I. D., Hill R. and Holder R. L. (1976) Fitting Johnson curves by moments in *Applied Statistics Algorithms*, Griffiths P., Hill I. D. (Eds), Ellis Horwood, Chichester.
- Li D.X. (2000) On default correlation: a copula function approach, *Journal of Fixed Income*, march, 43-54.

- Longin F. M. (2000) From Value-at-Risk to stress testing: the extreme value approach, *Journal of Banking & Finance*, 24, 1097-1130.
- Luciano E. and Marena M. (2001) Value at Risk bounds for portfolios of non-normal returns, *working paper*.
- Manhoney J.M. (1995) Empirical based versus model based approaches to Value-at-Risk : an examination of foreign exchange and global equity portfolios, *Risk Measurement and Systemic Risks, Proceedings of a joint central bank research conference*.
- Nelsen R.B. (1999) *An Introduction to Copulas*, Springer-Verlag, New York.
- Patton A. (2001) Modelling time varying exchange rate dependence using the conditional copula, *working paper*, University of California at San Diego
- Schonbucher P.J. and Schubert D. (2001) Copula dependent default risk in intensity models, *working paper*, University of Bonn.
- Wilson T. (1993) Infinite wisdom, *Risk*, 6, 6, 37-45.
- Zellner A. and Highfield R.H. (1988) Calculations of maximum entropy distributions and approximations of marginal posterior distributions, *Journal of Econometrics*, 37, 2, 195-210.