Working Paper no. 334

Using copulas to construct bivariate foreign exchange distributions with an application to the sterling exchange rate index

Matthew Hurd, Mark Salmon and Christoph Schleicher

November 2007

Bank of England

Using copulas to construct bivariate foreign exchange distributions with an application to the sterling exchange rate index

*Matthew Hurd**
*Mark Salmon***
and
Christoph Schleicher§

Working Paper no. 334

* Foreign Exchange Division, Bank of England.
Email: matthewhurd@cantab.net.
** Financial Econometrics Research Centre, Warwick Business School.
Email: mark.salmon@wbs.ac.uk.
§ Monetary Instruments and Markets Division, Bank of England.
Email: christoph@iam.ubc.ca (corresponding author).

The views expressed in this paper are those of the authors, and not necessarily those of the Bank of England. We would like to thank Michael Bennett, Andrew Patton, Alessio Sancetta and Petra Wikström, as well as seminar participants at the Bank of England, the CEF 2005 conference in Washington, the GFC 2005 conference in Dublin, the FFM 2006 conference in Aix-en-Provence, the University of Essex, the London School of Economics, and the University of York, for useful comments and discussion. This paper was finalised on 1 February 2007.

The Bank of England's working paper series is externally refereed.

Information on the Bank's working paper series can be found at www.bankofengland.co.uk/publications/workingpapers/index.htm.

Publications Group, Bank of England, Threadneedle Street, London, EC2R 8AH;
telephone +44 (0)20 7601 4030, fax +44 (0)20 7601 3298, email mapublications@bankofengland.co.uk.

ISSN 1749-9135 (on-line)

Contents

Abstract

We model the joint risk-neutral distribution of the euro-sterling and the dollar-sterling exchange rates using option-implied marginal distributions that are connected via a copula function that satisfies the triangular no-arbitrage condition. We then derive a univariate distribution for a simplified sterling effective exchange rate index. Our results indicate that standard parametric copula functions, such as the commonly used Normal and Frank copulas, fail to capture the degree of asymmetry observed in the data. We overcome this problem by using a non-parametric dependence function in the form of a Bernstein copula which is shown to produce a very close fit. We further give an example of how our approach can be used to price currency index options.

Key words: Option-implied distributions, effective exchange rate indices, copula functions, option pricing.

JEL classification: C16, C39, G13.

Summary

Option contracts give the right, but not the obligation, to buy or sell a financial asset at a predetermined price, known as the 'strike price'. As such, the value of an option depends on the likelihood that its holder will exercise this right. Therefore, option prices contain information about the probability that market participants attach to different outcomes of future asset prices. A common way to summarise this information is by estimating probability distributions of future asset prices implied by option prices.

Option contracts are traded for a wide range of currencies. Nevertheless, for the purpose of monetary policy, the effective exchange rate index (ERI) is more relevant than individual bilateral exchange rates. The main purpose of this paper is to develop a method that estimates option-implied distributions for effective exchange rates.

The sterling effective exchange rate is a weighted average of a large number of sterling bilateral exchange rates. However, it can be reasonably approximated by a function of only two exchange rates – the sterling-euro and the sterling-dollar bilaterals. The distribution of the sterling ERI can then be modelled as a function of the *joint distribution* of the sterling-euro and sterling-dollar exchange rates. The joint distribution describes the probability of all possible joint outcomes of the two exchange rates.

We observe option prices on the sterling-euro and the sterling-dollar exchange rates and use them to compute individual probability distributions for the two exchange rates. In order to link these two individual distributions to a joint distribution we make use of a so-called copula. Copulas are functions that link probabilities of individual events ('*it will be cloudy tomorrow*', '*it will rain tomorrow*') to those of outcomes of joint events ('*it will be cloudy and it will rain tomorrow*'). In the context of this paper, they join two one-dimensional distributions to create one two-dimensional distribution. Copulas are useful, because they provide a very general description of dependence patterns.

A potential problem is the fact that there are a very large (in fact infinite) number of copulas to choose from. We overcome this problem by imposing a no-arbitrage condition between the joint distribution of the sterling-euro and sterling-dollar exchange rates on the one hand, and the

univariate distribution of the euro-dollar cross-rate on the other hand. Observed option prices need to satisfy this condition, because otherwise they would present an opportunity to make a profit in excess of the risk-free rate without taking on any risk. Our no-arbitrage condition is a generalisation of the standard triangular no-arbitrage condition between any two spot exchange rates and their cross-rate. This narrows the choice of the copula function to those that are consistent with no-arbitrage.

In an empirical application we derive option-implied distributions of the sterling ERI on a daily basis between 2000 and 2005. We show that the distribution has seen considerable variation during this time. We also show how we can compute distributions of the sterling ERI that are conditional on movements in the euro-dollar exchange rate. This allows us to gauge the sensitivity of the sterling ERI to changes in the cross-rate of the United Kingdom's main trading partners.

In a second application we show that our method can be simply modified to calculate prices for options on exchange rate indices. We show that, contrary to standard models, the copula-based model generates a smile effect: options with strike prices that are further away from the current level of the effective exchange rate are relatively more expensive.

1 Introduction

This paper describes a new approach for constructing bivariate risk-neutral distributions implied by option prices. We consider FX options and model the dependence between a set of three currencies using a copula function that satisfies a triangular no-arbitrage condition such that the bivariate distribution of any two bilaterals is consistent with the univariate distribution of the third currency pair. This allows us to derive a risk-neutral distribution of the simplified sterling exchange rate index (ERI) under a consistent numeraire measure.

There is a very limited literature on the derivation of multivariate risk-neutral distributions and existing approaches are either based on fitting the implied correlation coefficient (Bikos (2000) and Taylor and Wang (2004)), individual option contracts (Bennett and Kennedy (2004)), or historical realised data (Rosenberg (2003)) in order to determine the dependence pattern. In contrast our method is novel in the sense that we exploit all available information in a set of (option-implied) density functions. We find that standard parametric copula functions (such as Gaussian, Frank or Gumbel) along with the perturbed normal method of Bennett and Kennedy (2004) are not able to provide a good fit to the data. This is especially true when one or more of the marginal distributions is heavily skewed. We overcome this problem by fitting a non-parametric Bernstein copula which can approximate any possible shape for the joint density.

The ERI is a key statistic for monetary policy purposes and hence also for market participants in assessing potential interest rate movements. Using our method we can produce theoretically and empirically consistent distributions for the ERI along with interval probabilities and standard summary statistics such as synthetic implied volatilities and risk-reversals. Our empirical results indicate that since 2002 the standard deviation of the sterling ERI has been relatively low (6% on an annualised basis). However, over the same period the distribution has been negatively skewed and fat-tailed. We also compute conditional densities to evaluate the sensitivity of the exchange rate index to movements in one currency pair.

Our analysis extends the toolkit of approximation methods available to price multivariate contingent claims beyond the assumption of multivariate Gaussianity. In particular, given risk-neutral densities for the implied cross-rate we can price claims with several underlying assets such as quantos and spread or basket options. In a second empirical application of this paper we

look at the problem of pricing index options, ie options whose payout depends on the geometric average of individual currency pairs.[1]

The paper is structured as follows. Section 2 gives a brief review of how copula functions provide an intuitive and practical way to simplify the problem of modelling multivariate distributions. Section 3 describes our approach and compares it to the existing literature. In Section 4 we compare the performance of several copula functions and relate it to the issue of different forms of asymmetric dependence. In Section 5 we report our results for the sterling ERI and present a short application to the problem of pricing index options. Section 6 concludes and discusses several possible avenues of future research.[2]

2 Modelling joint dependence

The sterling effective exchange rate index (ERI) is based on a geometric mean of 21 currencies.[3] In order to make statements about the probability of particular outcomes for the index, we therefore would need to model the joint distribution of all the individual currency pairs. There are two reasons why this approach is difficult in practice. First, in our case we do not have risk-neutral marginal distributions for all 21 inputs, as (liquid) options are only traded for the major currencies. This constraint may not be binding when pricing more general index options. Second, any attempt to model the joint distribution of 21 variables could, at best, be described as tricky without the use of some *ad hoc* assumption such as multivariate Gaussianity or Student-t. Not wishing to make such a compromise and in order to demonstrate our methodology, we model a simplified sterling ERI (SERI), which consists of a weighted average of the euro-sterling and the dollar-sterling bilaterals. As Chart 1 shows, the simplified index has tracked the ERI very closely over the past five years.

Reducing the problem to only two variables greatly simplifies matters. However, it is still not obvious, *a priori*, how we should model the joint distribution. As mentioned earlier, a simple benchmark would be to assume a jointly normal distribution which would be fully described by

(1) Examples of index options include options on futures contracts on a geometric weighted average of US dollar exchange rates which are traded on the CME (CME$INDEX^{TM}) and NYBOT (USDX$^{©}$).
(2) Four appendices contain more detailed background information on the relationship between risk-neutral distributions under different numeraire measures, the Bernstein copula and other copula functions as well as the estimation of option-implied marginal densities.
(3) See Lynch and Whitaker (2004) for more details.

Chart 1: Sterling ERI and simplified ERI (set equal for 1 Jan 2002)

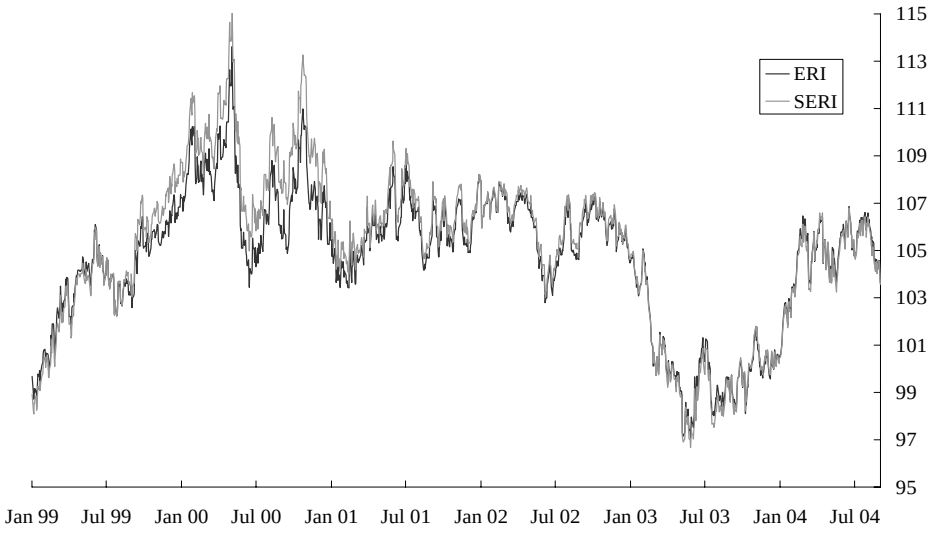

the forward rates, the implied volatilities and the implied correlation coefficient.[4] However, this approach only uses information from ATM (at-the-money) option contracts and ignores the higher moment structure in the data. Alternatively, we could use the product of option-implied marginal distributions, which would be able to account for departures from normality given by excess skewness and kurtosis but this would implicitly assume that the two marginals are independent from each other. The next section describes how copulas enable us to overcome this problem by allowing us to construct a joint distribution that allows for a general form of dependence *given* our marginal distributions.

2.1 Copulas

The central result in copula theory is a theorem by Sklar (1959), which states that any continuous N-dimensional cumulative distribution function F, evaluated at point $x = (x_1, ..., x_N)$[5] can be represented as

$$F(x) = C(F_1(x_1), ..., F_N(x_N)) \qquad \textbf{(1)}$$

(4) See Campa and Chang (1998) for more details on FX implied correlation coefficients.
(5) With some abuse of notation we fail to distinguish between random variables and their realisations. However, the meaning will be clear from context.

where C is called a copula function and F_i, $i = 1, ..., N$ are the margins.[6] The use of copulas therefore splits a complicated problem (finding a multivariate distribution) into two simpler tasks. The first task is to model the univariate marginal distributions and the second task is finding a copula that summarises the dependence structure between them.

It is also useful to think of copulas as joint distribution functions of standard uniform random variables $U_1 = F(X_1)$ and $U_2 = F(X_2)$:

$$\mathcal{C}(u, v) = \Pr(U_1 \leq u, U_2 \leq v).$$

The outcome of uniform random variables falls into the interval $[0, 1]$; therefore the domain of a copula must be the N-dimensional unit cube. Similarly, because the mapping represents a probability, the range of the copula must also be the unit interval. Also, it is easy to determine the value of a copula on the border of its domain. When one argument equals zero, the probability of any joint event must also be zero. Similarly, when all but one of the inputs are equal to one (that is, they are certain), the joint probability must be equal to the (marginal) probability of the argument that does not equal one. Finally, the function must be increasing in all its arguments.

Formally, a two-dimensional copula is a function $\mathcal{C} : [0, 1] \times [0, 1] \rightarrow [0, 1]$, such that

(i) $\mathcal{C}(u, 0) = \mathcal{C}(0, v) = 0$, ($\mathcal{C}$ is *grounded*),

(ii) $\mathcal{C}(u, 1) = u$ and $\mathcal{C}(1, v) = v$, (*consistent with margins*)

(iii) for any $u_1, u_2, v_1, v_2 \in [0, 1]$ with $u_1 \leq u_2$ and $v_1 \leq v_2$,

$$\mathcal{C}(u_2, v_2) + \mathcal{C}(u_1, v_1) - \mathcal{C}(u_1, v_2) - \mathcal{C}(u_2, v_1) \geq 0 \quad (\textit{2-increasing})$$

Intuitively, this last property ensures that the density of a copula (where it exists) is non-negative.

Furthermore, any copula function has a lower and an upper bound, C^- and C^+, which are called the minimum and the maximum copula, respectively. For any point $(u, v) \in [0, 1] \times [0, 1]$ the copula must lie in the interval

$$C^-(u, v) \equiv \max(u + v - 1, 0) \leq \mathcal{C}(u, v) \leq \min(u, v) \equiv C^+(u, v).$$

[6] We use the terms *marginal distribution* and *margin* interchangeably.

Finally, if two random variables are statistically independent, their distribution is given by the product copula, denoted by the symbol for orthogonality, or independence $\perp$;

$$F(u,v) = C^{\perp}(F_1(u), F_2(v)) = F_1(u)F_2(v).$$

A property that makes copulas extremely useful is that they are invariant under increasing non-linear transforms. To see this, note that for any monotonically increasing function $g(x)$; $\Pr(x \leq X) = \Pr(g(x) \leq g(X))$. In the context of this study, this is important, because we can use the same copula to describe the distribution of exchange rates in levels as in returns (log-differences from the current spot or forward rate).

As with standard distribution functions, copulas have associated densities [7]

$$c(u,v) = \frac{\partial^2 C(u,v)}{\partial u \partial v}$$

which permit the *canonical representation* of a bivariate density $f(u,v)$ as the product of the copula density and the density functions of the margins

$$f(u,v) = c\left(F_1(u), F_2(v)\right) f_1(u) f_2(v) \qquad \textbf{(2)}$$

This expression indicates how the simple product of two marginal distributions will fail to properly measure the joint distribution of two asset prices unless they are in fact independent and the dependence information captured by the copula density, $c\left(F_1(u), F_2(v)\right)$, is normalised to unity.

2.2 Measuring association

While a copula function fully describes the dependence structure between two (or more) random variables, we are often interested in simple summary statistics that indicate the extent and the direction of comovement. The most basic of these measures are concordance measures, which assume the value of -1 for perfect negative association (counter-monotonic, corresponding to the lower Fréchet bound or the minimum copula C^{-}), 0 for independence (corresponding to $C^{\perp}$) and 1 for perfect positive association (comonotonic, corresponding to the upper Fréchet bound or the maximum copula C^{+}). Concordance measures are completely determined by the underlying copula and independent of the marginal distributions unlike Pearson's correlation. As such concordance measures are also invariant under monotonically increasing transforms which in the

[7] The density exists a.e. in the interior of the domain. See Cherubini *et al* (2004, page 81) for further details.

context of this study implies that the levels of exchange rates have the same concordance measure as their returns (measured as log-difference from the current price or the forward rate).

The two most common concordance measures are Kendall's tau and Spearman's rho (Kendall (1938)). For two pairs of $i.i.d.$ random vectors (X_1, Y_1), and (X_2, Y_2), Kendall's tau has an interpretation as the difference between the probability of a joint outcome with the same sign $(\Pr[(X_1 - X_2)(Y_1 - Y_2) > 0])$ and the probability of a joint outcome with different signs $(\Pr[(X_1 - X_2)(Y_1 - Y_2) < 0])$, in other words the difference between the probability of the variables either rising or falling together (concordance) minus the probability of them moving in different directions (discordance). Similarly, Spearman's rho measures the correlation between rank ordered data.[8] The standard Pearson correlation coefficient is not a concordance measure and in particular, in contrast to Kendall's tau and Spearman's rho, it will be changed by non-monotonic transformations of the data since it is not independent of the marginal distributions. Despite these failings of correlation as a measure of non-linear dependence, foreign exchange analysts frequently use the correlation coefficient derived from implied volatilities to infer the dependence relationship between different exchange rates. We will compare the use of correlation with the two concordance measures described in this section in our application below.

3 The methodology

Our methodology builds on earlier unpublished work by Bikos (2000), who uses one-parameter copulas[9] such as the Gaussian and the Frank copula to model the joint distribution of the dollar-sterling and euro-sterling bilaterals. The marginal distributions are given by univariate risk-neutral densities estimated using the Malz (1997) method and the parameter of the copula function is chosen in such a way that the empirical correlation coefficient (computed from the variances of the three margins) equals the implied correlation coefficient (computed from ATM volatilities). A very similar approach has been taken in a recent contribution by Taylor and Wang

[8] In terms of a copula function C, Kendall's tau can be expressed as

$$\tau = 4 \int \int_{I^2} C(u, v) dC(u, v) - 1.$$

Similarly, Spearman's rho can be written as

$$\rho = 12 \int \int_{I^2} uv dC(u, v) - 3.$$

[9] These are functional forms that are described by a single parameter which is typically directly related to the concordance measures described in Section 2.2.

(2004), who also fit the implied correlation coefficient, but use a more refined set-up which ensures that the implied joint density belongs to a common risk-neutral numeraire measure. Both studies (Bikos, and Taylor and Wang) find that one-parameter copulas provide a good fit to the data but essentially use one observation to fit a single parameter.[10]

A more general approach is taken by Bennett and Kennedy (2004), who use copulas in conjunction with a triangular no-arbitrage condition to price quanto FX options, ie FX options whose payout is in a third currency. Similarly to Bikos, and Taylor and Wang, they use option-implied[11] densities as margins for the bivariate distribution. However, they estimate their copula function by fitting an entire set of option contracts in the third bilateral (over different strike prices) instead of fitting just the implied correlation coefficient. This additional information leads them to use a Gaussian copula which is perturbed by a cubic spline, and which therefore allows for a more flexible dependence structure between the three currency pairs. In the context of the quanto pricing problem this approach is appealing because the perturbation function indicates the extent of departure from the standard Black model which corresponds to a jointly lognormal distribution.

We extend these previous methods by estimating a joint distribution that is consistent with the option-implied marginal distribution of the third bilateral *over its entire support*. In order to do this we proceed in the following steps:

1. We use a similar line of reasoning as Taylor and Wang to derive a relationship between the density of the euro-dollar exchange rate (denoted by z), under the risk-neutral measure corresponding to a euro-denominated discount bond, and the bivariate density of the euro-sterling and the dollar-sterling exchange rates (denoted by x and y, respectively), under the risk-neutral measure corresponding to a sterling-denominated discount bond. However, we formulate our problem in return-space rather than in level-space by defining each of the three exchange rates (x, y, and z) in terms of their log-deviations from the forward rate. We argue that this set-up is more natural, because it preserves the natural symmetry of exchange rate

[10] Rosenberg (2003) follows a different route by using a non-parametric method and a copula which is estimated from *historical* exchange rate movements.

[11] While Bennett and Kennedy, as well as Taylor and Wang, work with a mixture of two log-normals distributions, Bikos (2000) and this paper use the Malz method to compute risk-neutral densities. The advantages and disadvantages of parametric and smile-interpolation approaches are discussed in Bliss and Panigirtzoglou (2002). We emphasise, however, that the general methodology discussed in this paper is independent of the way the marginal distributions are computed.

movements. As we show in Appendix A, this relationship is given by

$$f_z^{\mathbb{Q}_E}(s) = \int_{-\infty}^{\infty} f_{xy}^{\mathbb{Q}_S}(u, u-s)e^{-u}du \tag{3}$$

where $f_z^{\mathbb{Q}_E}$ denotes the risk-neutral density of euro-dollar (under a euro numeraire measure) and $f_{xy}^{\mathbb{Q}_S}$ is the bivariate distribution of euro-sterling and dollar-sterling (under a sterling numeraire measure). Note that by triangular arbitrage it must be the case that

$$x = y + z \tag{4}$$

2. By Sklar's theorem we may assume the existence of a copula $C(\cdot)$ with density $c(\cdot)$ in order to obtain the bivariate distribution of x and y in its canonical representation as shown above

$$f_{xy}^{\mathbb{Q}_S}(u, v) = c\left(F_x^{\mathbb{Q}_S}(u), F_y^{\mathbb{Q}_S}(v)\right) f_x^{\mathbb{Q}_S}(u) f_y^{\mathbb{Q}_S}(v) \tag{5}$$

3. We estimate a parametric representation, $\hat{c}(\cdot, \hat{\theta})$, of the copula density by minimising the L^2-distance between the option-implied third bilateral $f_z^{\mathbb{Q}_E}$ and $\hat{f}_z^{\mathbb{Q}_E}(\cdot, \hat{\theta})$

$$\hat{\theta} = \text{arginf}_\theta \left[\int_{-\infty}^{\infty} \left(f_z^{\mathbb{Q}_E}(s) - \hat{f}_z^{\mathbb{Q}_E}(s;\theta)\right)^2 ds\right]^{\frac{1}{2}} \tag{6}$$

where

$$\hat{f}_z^{\mathbb{Q}_E}(s, \theta) = \int_{-\infty}^{\infty} \hat{c}\left(F_x^{\mathbb{Q}_S}(u), F_y^{\mathbb{Q}_S}(u-s); \theta\right) f_x^{\mathbb{Q}_S}(u) f_y^{\mathbb{Q}_S}(u-s)e^{-u}du \tag{7}$$

is the distribution of the third bilateral implied by parameters θ.

4. We can then use change of variable techniques to obtain bivariate distributions of the exchange rates in levels, as well as the distribution of the simplified sterling ERI. In terms of the deviation from its forward rate (implied by the UIP condition) the SERI can be expressed as

$$\xi = \omega_E x + \omega_D y \tag{8}$$

where ω_E and ω_D are the weights of the euro and the dollar, currently set to 0.8 and 0.2. By substituting for y, the implied distribution of ξ under the sterling risk-neutral measure can then be computed as

$$f_\xi(s) = \int_{-\infty}^{\infty} f_{xy}^{\mathbb{Q}_S}\left(u, \frac{s - \omega_E u}{\omega_D}\right) \frac{1}{\omega_D} du \tag{9}$$

Once we have estimated the joint distribution of x and y under the sterling risk-neutral measure, it is possible to also derive the corresponding distributions $f_z^{\mathbb{Q}_S}$, $f_{xz}^{\mathbb{Q}_S}$ and $f_{yz}^{\mathbb{Q}_S}$, as well as the corresponding PDFs and CDFs in level-space.

4 Choosing the right copula

In the following we apply a variety of different copula functions to the empirical framework outlined in the previous section. Ultimately this leads us to choose a very flexible representation in the form of a Bernstein expansion, as developed recently by Sancetta (2003) and Sancetta and Satchell (2004).

4.1 A comparison of different copula functions

Most existing applications of copulas in finance have focused on simple alternatives such as the Gaussian, Frank, Gumbel or Plackett copula,[(12)] since typically there is relatively little information available about the margins and their dependence relationship. When historical data are used, the distribution is only observed one realisation per time period and therefore has to be estimated over a moving window, which makes it especially difficult to identify changes in the dependence structure. For example, it is often assumed that the margins have a Student t-distribution and the dependence is Gaussian (described by the normal copula). In comparison, by using option-implied densities, we are able to observe the entire estimated marginal distributions on a daily basis. Furthermore, the triangular nature of exchange rates imposes an additional restriction on the dependence structure.

Our implied densities are estimated using one-month, three-month and twelve-month over-the-counter option contracts which were obtained from the British Bankers' Association (BBA) and Citibank. For three and twelve-month contracts we use the Malz (1997) method, which fits a quadratic smile to the at-the-money volatility and the $\pm$25-delta volatilities implied by risk reversals and strangles. For one-month contracts our data set includes additional $\pm$10-delta contracts and we use a spline-interpolation which is briefly described in Appendix C.

We find that one-parameter copulas provide a fairly good fit when the distributions of the three exchange rates are roughly symmetrical. An example is given in Chart 2, which shows the distribution of the return on the euro-dollar exchange rate for 23 July 2004, as well as the approximations discussed in the previous section. From Table A we see that the normal copula underestimates the extent of kurtosis. The Frank and Plackett copulas have approximation errors

(12) See Appendix C and Joe (1997) for a description of parametric copula families.

Chart 2: Actual and fitted risk-neutral distribution for one-month euro-dollar contracts on 23 July 2004

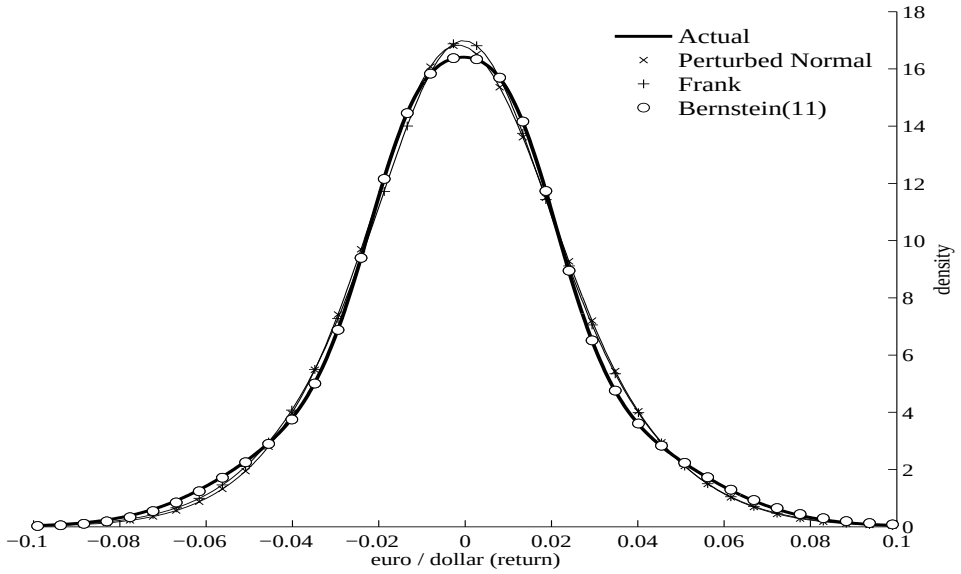

(under the L^2 criterion) which are somewhat smaller than that of the normal copula.

However, as the density of the dollar-euro exchange rate becomes more skewed we find that these one-parameter copulas frequently fail to provide an accurate description of the data. An example is given in Chart 3 which shows approximations to the distribution of one-month euro-dollar contracts on 18 December 2002. In this case the best one-parameter copula has an approximation error of only slightly less than 30% (Table B) which is largely due to the fact that the approximated densities fail to account for the extra probability mass in the left tail of the distribution (corresponding to dollar depreciation). We therefore considered more flexible functional forms, such the multi-parameter BB1 and BB7 (Joe-Clayton) copulas (see Joe (1997) for more details), as well as the perturbed normal specification of Bennett and Kennedy. However, as can be seen from the table, the results indicate that these copulas were only slightly more successful in matching the skew in the third bilateral than their one-parameter counterparts.

4.2 Modelling asymmetric dependence

The reason behind this failure, we believe, lies in the inability of standard copulas to account for a general degree of asymmetry in the dependence structure. We consider two different forms of

Table A: Comparison of fit (23 July 2004)

PDF	L^2-dist (%)	mean	std	skew	kurt
Actual	–	**-0.000**	**0.0268**	**0.023**	**3.661**
Normal	16.15	-0.000	0.0250	0.042	3.258
Frank	12.58	-0.000	0.0261	0.038	3.605
Plackett	12.76	-0.000	0.0264	0.036	3.653
BB1	13.98	-0.000	0.0256	0.072	3.463
BB7	14.64	-0.000	0.0254	0.059	3.407
Asy. Gumbel	12.18	-0.000	0.0262	0.029	3.637
Pert. Normal	13.82	0.000	0.0257	0.071	3.501
Bernstein(2)	15.06	-0.000	0.0254	0.044	3.319
Bernstein(5)	8.94	0.000	0.0273	0.074	4.073
Bernstein(7)	4.51	0.000	0.0273	0.108	3.971
Bernstein(9)	1.70	0.000	0.0269	0.107	3.744
Bernstein(11)	1.50	0.000	0.0269	0.124	3.752
Bernstein(13)	1.48	0.000	0.0269	0.133	3.763

asymmetry. In the first case joint negative events (the lower left quadrant) are more dependent than joint positive events (the upper right quadrant), or *vice versa*. In terms of a copula's density this can written as

$$c(u, v) \neq c(1 - v, 1 - u).$$

We call this (for lack of a better expression) *asymmetry along the 45 degree line*. This concept is useful for many asset classes such as equities, because the degree of comovement across different markets is often thought to depend on its direction. For instance, a frequently made observation is that dependency is greater during downturns than during expansions (see Cherubini *et al* (2004), for examples) and there are several copulas, such as Gumbel's copula or the Joe-Clayton copula that can model this form of asymmetry. In the second case of asymmetry, however, dependence for a negative outcome of the first variable and a positive outcome in the second variable (the upper left quadrant) is greater than dependence for a positive outcome in the first variable and a negative outcome in the second variable (the lower right quadrant), or *vice versa*. We can write this as

$$c(u, v) \neq c(v, u).$$

In this case, the copula is not *interchangeable*, that is $C(u, v) \neq C(v, u)$. We argue that for exchange rates, this second form of asymmetry (which we call *asymmetry across the 45 degree line*), is as equally important as that *along the 45 degree line*.[(13)] This is because unlike equities,

[(13)] Patton (2006) finds asymmetry in the exchange rates of the dollar with the yen and the mark during the 1990s, which he explains by competitiveness between Japanese and German exporters. This argument relies on the assumption that the US assumes the position of a passive (numeraire) player.

Chart 3: Actual and fitted risk-neutral distribution for one-month euro-dollar contracts on 18 December 2002

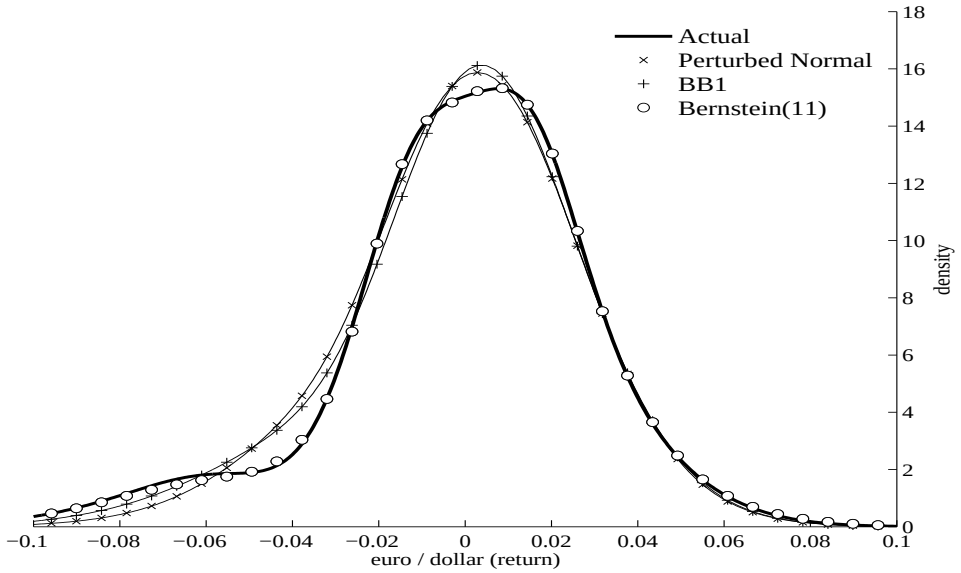

changes in exchange rates do not have a natural interpretation in terms of upward or downward movements; the appreciation of one currency pair corresponds to the depreciation of its reciprocal.

To our knowledge there has been very little work on parametric representations for non-interchangeable copulas. Exceptions are the doctoral thesis of Khoudraji (1995) and a recent article by Mari and Monbet (2004) in which they develop asymmetric extensions to the Clayton and the Gumbel families.

When we fit the asymmetric Gumbel copula to the skewed distribution of the one-month maturity euro-dollar contracts of 18 December 2002, we find that it lowers the approximation error from 30% to about 27%. However, this degree of error is still far too large to be acceptable; it is apparent that we need to find a different copula to capture the dependence structure between foreign exchange rates.

A very general starting point for constructing a copula is to define its values on a discrete set of points such that the three conditions in the definition of a copula (grounded, consistence with margins and two-increasing) are satisfied. We can then use some form of interpolation to obtain a functional form on the unit square. One possible way to obtain such an interpolation is to use a

Table B: Comparison of fit (18 December 2002)

PDF	L^2-dist (%)	mean	std	skew	kurt
Actual	–	**-0.000**	**0.0297**	**-0.682**	**4.248**
Normal	29.95	-0.000	0.0273	-0.294	3.456
Frank	29.79	-0.000	0.0278	-0.310	3.614
Plackett	29.87	-0.000	0.0279	-0.311	3.621
BB1	23.96	-0.001	0.0284	-0.498	3.721
BB7	28.24	-0.000	0.0273	-0.339	3.570
Asy. Gumbel	27.31	-0.000	0.0277	-0.440	3.645
Pert. Normal	27.38	-0.000	0.0271	-0.319	3.462
Bernstein(2)	29.52	-0.000	0.0277	-0.307	3.587
Bernstein(5)	19.19	-0.000	0.0283	-0.529	4.038
Bernstein(7)	9.91	0.000	0.0295	-0.649	4.418
Bernstein(9)	5.55	0.000	0.0299	-0.685	4.511
Bernstein(11)	3.59	0.000	0.0297	-0.629	4.340
Bernstein(13)	4.89	0.000	0.0297	-0.659	4.420

Bernstein expansion.[(14)] This has been proposed by Li *et al* (1998) and the ensuing Bernstein copula and its theoretical characteristics are discussed by Sancetta (2003) and Sancetta and Satchell (2004). One very desirable property of the Bernstein copula is its ability to approximate any possible copula function. We refer the reader to Appendix B for a more detailed description of the Bernstein copula as well as an efficient way to fit it to a given set of triangular marginal distributions.

The results of our application of the Bernstein copula to the euro-dollar distributions for 23 July 2004 and 18 December 2002 are shown in the bottom sections of Tables A and B. We use approximations with Bernstein polynomials of order 2 to 11 and find that for both dates the L^2 goodness-of-fit criterion improves quickly and that an 11^{th} order expansion provides a fit which is dramatically better than that of the parametric and perturbed normal copulas.

The top panels in Chart 4 show the level curves of the Bernstein copula and its density for 24 June 2002. For this particular day the dependence structure is more or less symmetric across the 45 degree axis. However, it is asymmetric along the 45 degree axis, in the sense that dependence between joint negative outcomes is stronger than dependence between joint positive outcomes. The middle row shows the cumulative bivariate distribution and its density, which is obtained by

(14) The Bernstein expansion is a linear combination of polynomials (Bernstein polynomials) which are described in more detail in Appendix B. A family of Bernstein polynomials is shown in Chart B1.

Chart 4: Distributions of euro-sterling and dollar-sterling for one-month contracts on 24 June 2002

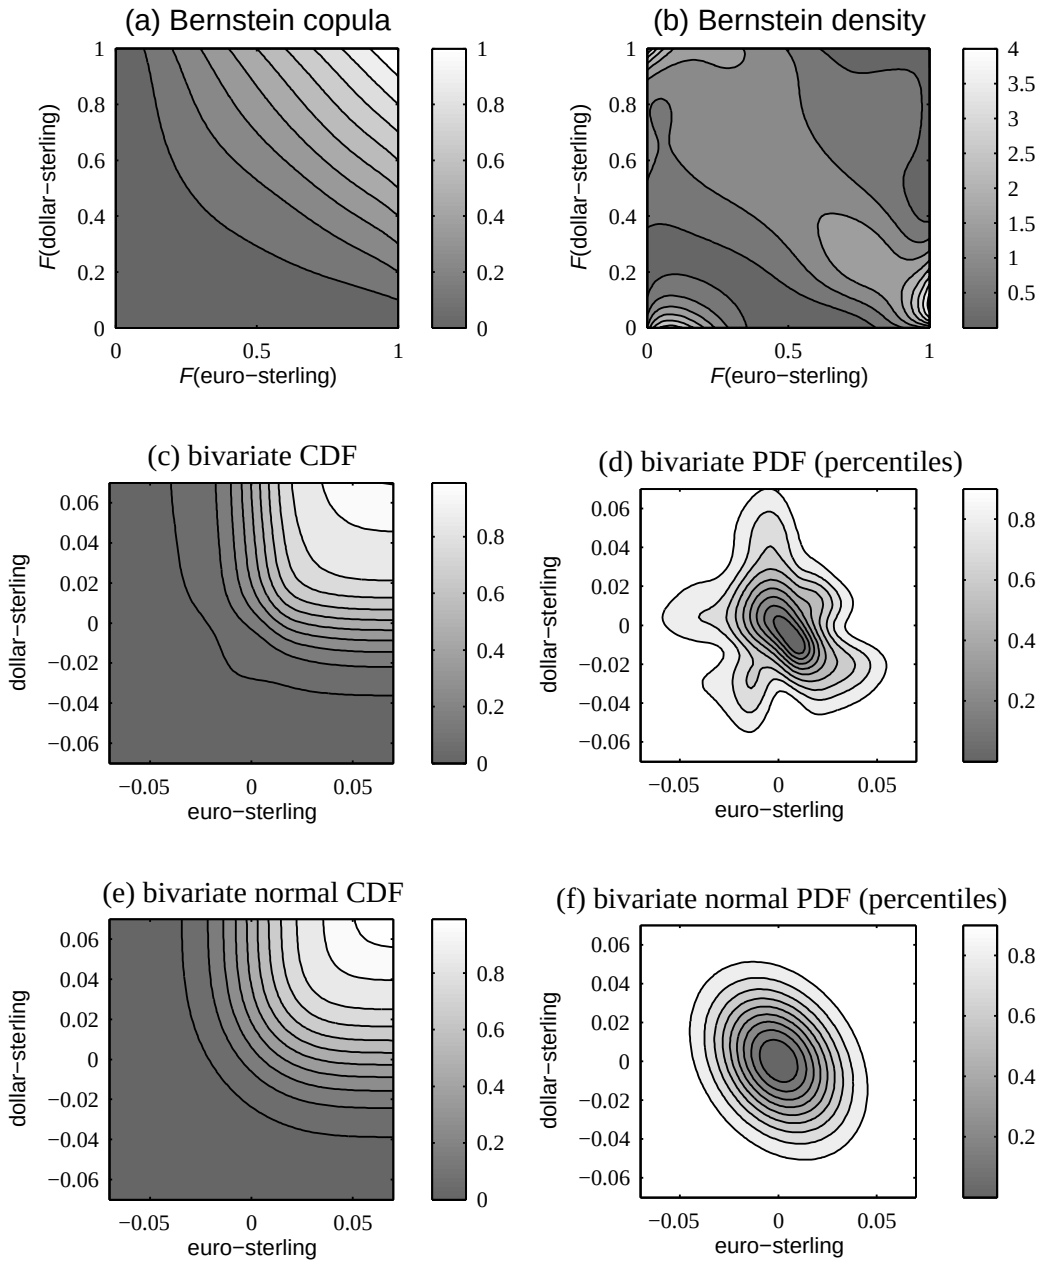

Panels (a) and (b) contain level curves of the Bernstein(11) copula and its density. Panels (c) and (d) show distributions implied by the Bernstein copula, while panels (e) and (f) show the bivariate normal benchmark distributions (based on the implied correlation coefficient). The variables are expressed in terms of their return relative to the current forward rate.

connecting the copula with the option-implied marginal distributions for euro-sterling and dollar-sterling. Overall there is a negative slope in the dependence between the two bilaterals. This is reflected in negative values of the dependence measures ($\tau_{\text{Kendall}} = -0.20$,

$\rho_{\text{Spearman}} = -0.27$ and $\rho_{\text{correlation}} = -0.29$). Finally, the two lower panels show what the joint CDF and PDF would look like if we assumed joint normality (using only information from implied volatilities and forward rates). The density has the characteristic elliptical shape of a bivariate normal distribution, which is downward-tilted, corresponding to a negative implied correlation coefficient. Compared to the estimated copula distribution, the bivariate normal distribution can be seen to provide a highly simplified description of the data.

5 Empirical results

In the following we present empirical estimates of SERI densities, as well as a short application to the problem of pricing FX index options.

5.1 Assessing foreign exchange risks

The estimated bivariate distribution provides a compact summary about the market's consensus about the magnitude and balance of future exchange rate risks. Here we describe several statistics and graphical tools which we believe may be helpful for policymakers and market participants in order to assess these risks.

5.1.1 Summary statistics

In order to get a first idea how the distribution of the SERI has changed over time, Chart 5 shows the annualised standard deviation, skewness and kurtosis of one-month return PDFs. It is evident that volatility has seen large changes, with the (annualised) standard deviation fluctuating between 4% and 13%. In general, volatility is considerably lower in the second half of our sample (after the end of 2001). Turning to skewness, we find that the balance of risk is typically close to symmetric, except for the period from 2002 to the middle of 2003, in which the probability of large depreciations was seen to be larger than that of large appreciations. This result is interesting insofar as the negative skew precedes the actual strong depreciation of sterling towards the beginning of 2003. Finally, return PDFs have been moderately, but consistently fat-tailed, with an average fourth moment of about 4. The third and fourth moments appear to be closely related as negative skews tend to coincide with large levels of kurtosis.

Chart 6 shows a time series of the implied correlation coefficient for one-month maturities, as well

Chart 5: First four moments of the implied simplified sterling ERI (log-returns)

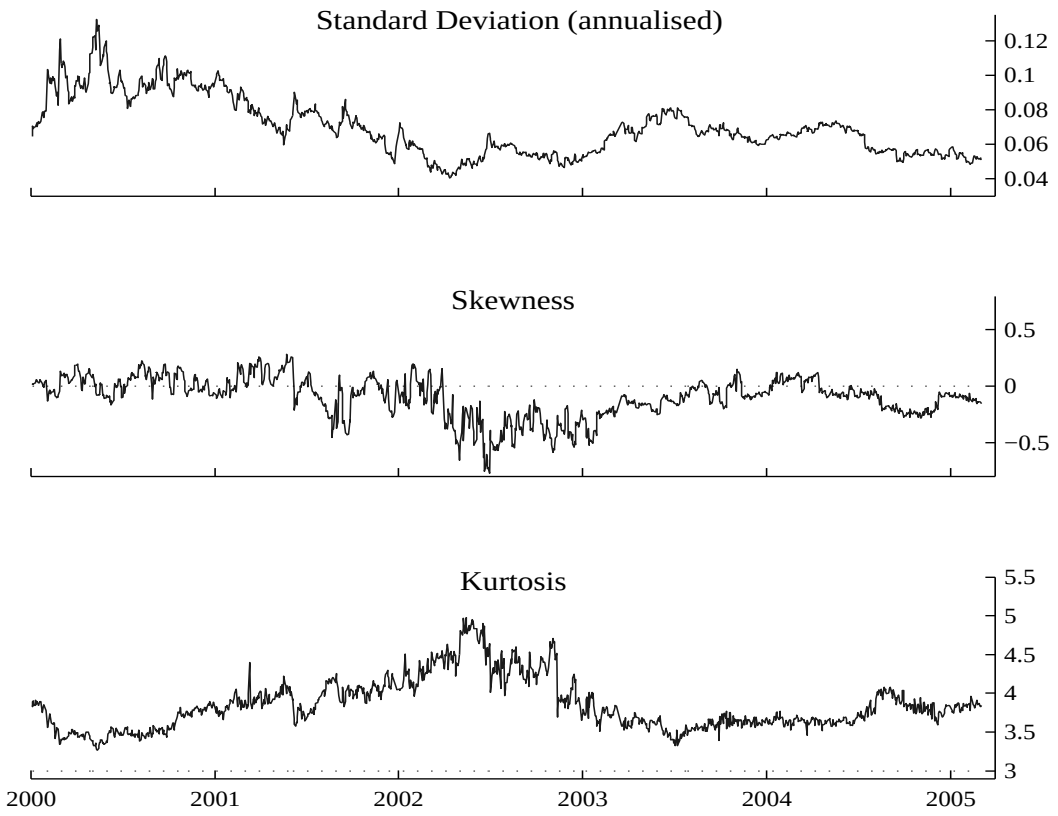

Chart 6: Different measures for the dependence between euro-sterling and dollar-sterling

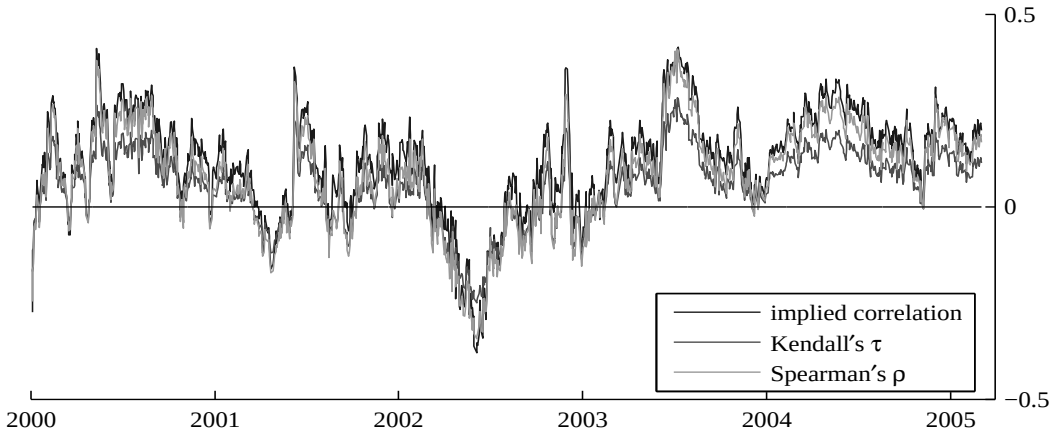

as Kendall's tau and Spearman's rho, the two dependence measures described in Section 2.2. Most of the time these measures indicate a (slightly) positive relationship between euro and the dollar

versus sterling. However, there is a marked period during the first half of 2002 when this relationship turned negative. In general, we find that the three measures are closely related. However, the linear correlation coefficient is virtually always larger in absolute size than Spearman's rho and, especially, Kendall's tau. There is also evidence that the (return) standard deviation of the SERI is positively related to the magnitude of dependence between the two variables.[15] A low (or negative) dependence between sterling's two main bilaterals is often viewed as a desirable property, as it insulates the effective exchange rate against external shocks.

5.1.2 Unconditional and conditional risk-neutral PDFs of the ERI

Having estimated the joint distribution of the euro-sterling and the dollar-sterling bilaterals we can compute distributions of the (simplified) exchange rate index (as we described in Section 3). These distributions can be unconditional (ie taking into account all possible constellations of the three exchange rates) or conditional on certain outcomes. Unconditional distributions of the level of the ERI for the first business day in the years 2001, 2003 and 2005 are shown in the left panel of Chart 7. We can see that the level as well as the shape of the distributions have seen considerable variation over the past five years. For example, the implied variance of the 2001 curve is much larger than that of the 2003 curve, leading to a flatter and wider PDF. The right panel of this chart shows return PDFs for one, three and twelve-month maturities.

We can further condition the distribution of the ERI on certain outcomes of the exchange rates involved. For example we can compute a distribution under the assumption that the dollar appreciates or depreciates by a certain percentage against the euro. Say, that at the maturity of the option contracts, the level of the euro-dollar bilateral equals $z = \bar{z}$. In this case, by Bayes' law, the conditional joint distribution of euro-sterling (x) and dollar-sterling (y) would be

$$f_{xy}\left(u, v \,|z = \bar{z}\right) = \frac{f_{xy}\left(u, u - \bar{z}\right)}{\int_{-\infty}^{\infty} f_{xy}\left(s, s - \bar{z}\right) ds} = f_x\left(u \,|z = \bar{z}\right) \qquad \textbf{(10)}$$

From equation **(8)** we can then derive the conditional distribution of the ERI as

$$f_\xi\left(u \,|z = \bar{z}\right) = f_x\left(u + \omega_D \bar{z} \,|z = \bar{z}\right) \qquad \textbf{(11)}$$

Chart 8 plots the one-month distribution of the ERI under different scenarios for the euro-dollar exchange rate. Chart 9 illustrates how each of these scenarios corresponds to an intersection of the bivariate distribution between the euro-sterling and dollar-sterling bilaterals. It can be seen that an

[15] The regression coefficient for Spearman's rho in a regression explaining the return standard deviation has a coefficient of 4.6 with a t-value of 13.6.

Chart 7: One-month SERI level PDFs (left panel) and return PDFs for 24 June 2004 and different maturities (right panel)

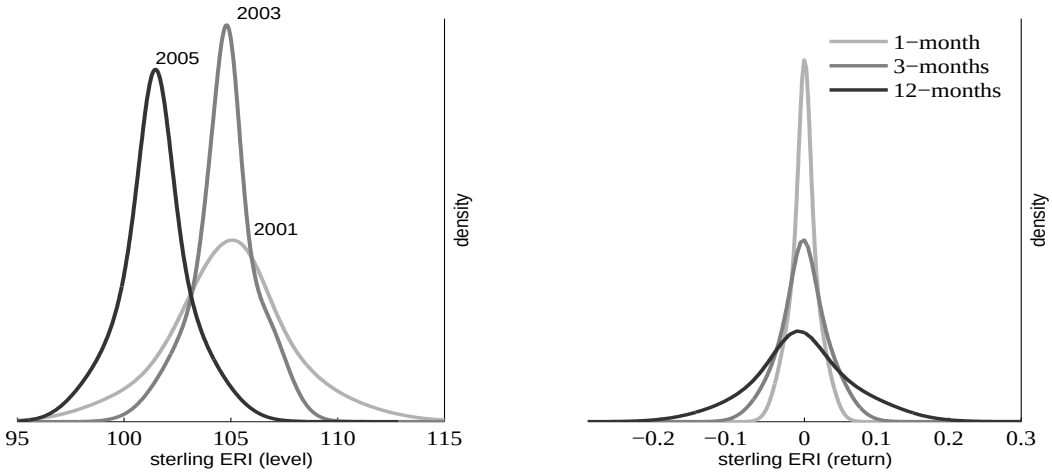

appreciation of the dollar versus the euro (decrease in z) shifts the location of the conditional ERI PDF downwards, while a depreciation of the dollar increases the level of the conditional density and also leads to a bimodal shape. The reason for this behaviour is the asymmetry of the bivariate distribution across the 45 degree axis.

5.2 Pricing index options and deriving synthetic contracts for the ERI

Knowledge of the implied bivariate risk-neutral distribution of two assets immediately allows us to price multi-factor contingent claims such as basket options, spread options, and quantos. In particular, consider a European-style pay-off in currency c which depends on some functional form $G(S_T^{c,a}, S_T^{c,b})$ of currency pairs $S_t^{c,a}, S_t^{c,b}$ at time T. Then it is well known that a contingent claim V at time t can be priced via the integral

$$V_t\left(G(S_t^{c,a}, S_t^{c,b})\right) = e^{-r_c(T-t)} \int_0^\infty \int_0^\infty G(u,v) \tilde{f}^{\mathbb{Q}_c}_{S_T^{c,a}, S_T^{c,b}, t}(u,v) du dv \qquad \textbf{(12)}$$

where $\tilde{f}^{\mathbb{Q}_c}_{S_T^{c,a}, S_T^{c,b}, t}(\cdot,\cdot)$ is the bivariate risk-neutral distribution of $S_T^{c,a}, S_T^{c,b}$ under numeraire-measure c at time t for time T.

Here we consider the case of index options, which is directly related to the problem of estimating risk-neutral distributions for effective exchange rate indices. Because of this close analogy, we can interpret the prices of standard contracts such as straddles, risk-reversals and butterflies of an index option as those belonging to hypothetical contracts on the (simplified) exchange rate index

Chart 8: One-month level PDFs of the SERI on 24 September 2004, conditional on different outcomes for euro-dollar

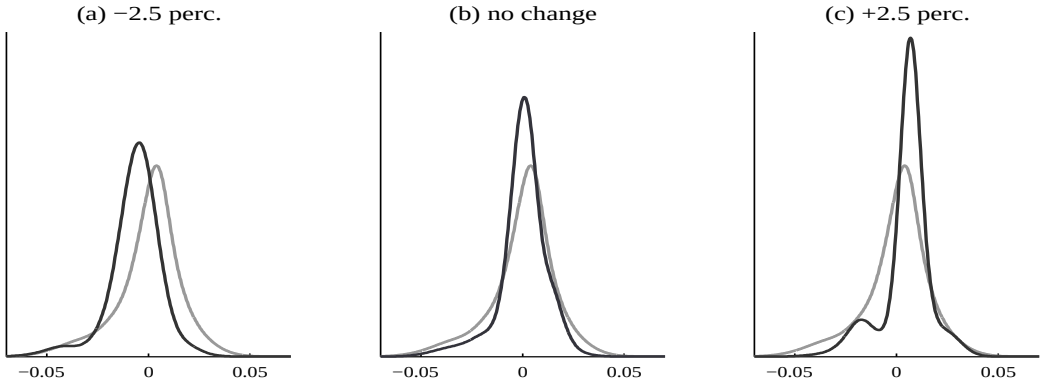

The unconditional distributions are shown in light grey.

itself. This is useful because policymakers such as central banks often use time series of these contracts to obtain insights on market expectations about exchange rate risks.

A possible way to formulate the pay-off of a two-asset index call option for currencies $S_t^{c,a}$, $S_t^{c,b}$ with strike K and expiry T is

$$\text{payoff} = \max\left(\left(\frac{S_T^{c,a}}{M_{t,T}^{c,a}}\right)^{\omega_a}\left(\frac{S_T^{c,b}}{M_{t,T}^{c,b}}\right)^{(1-\omega_a)} - K, 0\right) \qquad \textbf{(13)}$$

where $M_{t,T}^{c,a}$ and $M_{t,T}^{c,b}$ are the forward rates at time t. In this case, the index is the geometric average of the deviation of the two exchange rates from their forward rates. In order to price such an option, we take the prices of vanilla options on the individual currency pairs as given and compute an option-implied PDF for an index with weights ω_a and $(1-\omega_a)$ using the same method as for the simplified sterling ERI.[16] We can then obtain the call-price function by integrating the implied PDF twice with respect to the strike price (essentially reversing the result of Breeden and Litzenberger (1978)).

Since standard OTC FX contracts are typically expressed in terms of delta (first derivative with respect to the underlying) rather than strikes, and priced in terms of vols (implied volatility) rather than currency units, we transform the call prices into implied volatilities (using Black's formula) and the strikes into deltas (using the implied volatility of a strike equal to the forward rate).

[16] The relationship to the simplified sterling ERI is easily seen by comparing equations **(8)** and **(13)** and setting $x = \log(S_T^{c,a}/M_{t,T}^{c,a})$, $y = \log(S_T^{c,b}/M_{t,T}^{c,b})$, $\omega_a = \omega_E$, and $(1-\omega_a) = \omega_D$.

Chart 9: Bivariate one-month distribution of the SERI on 1 September 2004

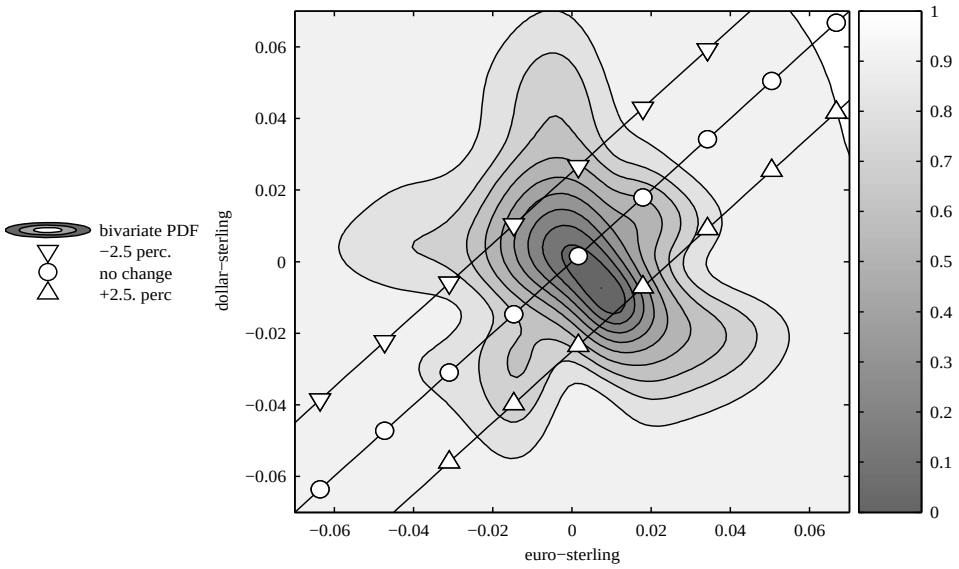

Chart 10 shows the implied volatility smile for an index option with 50:50 weights of euro and yen versus the dollar and a maturity of one month for 24 June 2004. For comparison we also show the constant smile corresponding to the multivariate Black model (assuming joint log-normality) [17] and an *ad hoc* adjustment. For the latter, we take the average of the smiles implied by vanilla options for the euro-dollar and yen-dollar bilaterals and adjust its mean in a way that it equals the implied volatility of the Black model at a delta of 0.5. Compared to the Black model which corresponds to a straight smile, the *ad hoc* adjustment therefore recognises the fact that implied volatilities are strike-dependent. We then compute the straddle (50 delta call + 50 delta put), risk reversal (25 delta call - 25 delta put) and butterfly (25 delta call + 75 delta call - straddle) for each model. We also compute risk reversals and butterflies corresponding to 10-delta contracts. The prices are given in Table C and indicate that the *ad hoc* adjustment performs well for the straddle and the 10-delta contracts, but poorly for the 25-delta contracts (especially the risk reversal).

Chart 10 compares synthetic implied volatilities (straddles) and risk reversals for the SERI with actual contracts for the euro-sterling and dollar-sterling bilaterals. SERI implied volatilities are similar in magnitude to the annualised standard deviation of the return PDF shown in Chart 5. Since the middle of 2000 they have been consistently lower than the implied volatilities of either

[17] The problem of pricing options on geometric indices has been discussed *inter alia* by Eytan and Harpaz (1986) in the context of the Value Line Index (VLCI) and by Bhargava and Clark (2003) in the context of the US Dollar Index (USDX). Both studies use multivariate log-normal extensions of the Black model.

Chart 10: Smile of a 50:50 index option of euro and yen versus the dollar for one-month contracts on 24 June 2004

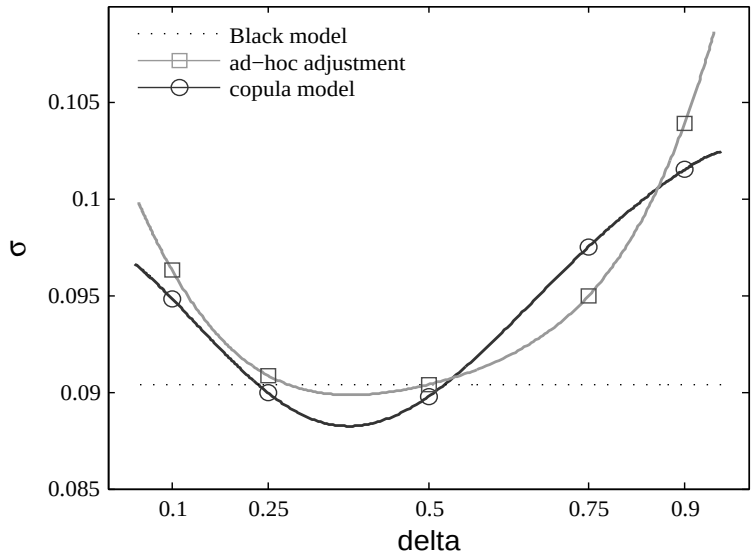

Table C: Prices of index option contracts for 24 June 2004 (implied volatility)

	straddle	risk rev.$_{25\Delta}$	risk rev.$_{10\Delta}$	butterfly$_{25\Delta}$	butterfly$_{10\Delta}$
Black model	9.04	0.00	0.00	0.00	0.00
ad-hoc adj.	9.04	-0.41	-0.76	0.25	0.97
copula model	8.99	-0.75	-0.67	0.38	0.83
error$_{\text{Black}}$	*-1%*	*100%*	*100%*	*100%*	*100%*
error$_{\text{ad-hoc}}$	*-1%*	*45%*	*-14%*	*34%*	*-17%*

of the two individual exchange rates, which is due to the fact that the two series are not perfectly correlated. This analogous to the fact that a basket or index option is cheaper than a set of options on the individual underlyings. The time series of SERI risk reversals has similar characteristics as the skew of the return PDF. It is also interesting to see that SERI risk reversals are very similar to sterling-euro risk reversals and often move in opposite direction to the sterling-dollar risk reversal. This close alignment is likely to be a consequence of the large weight of the euro in the currency index.

6 Conclusions and further research

In this paper we have proposed a new approach to derive bivariate risk-neutral distributions from option prices in foreign exchange markets. A distinguishing feature of our method is that it

Chart 11: One-month implied volatilities and risk reversals

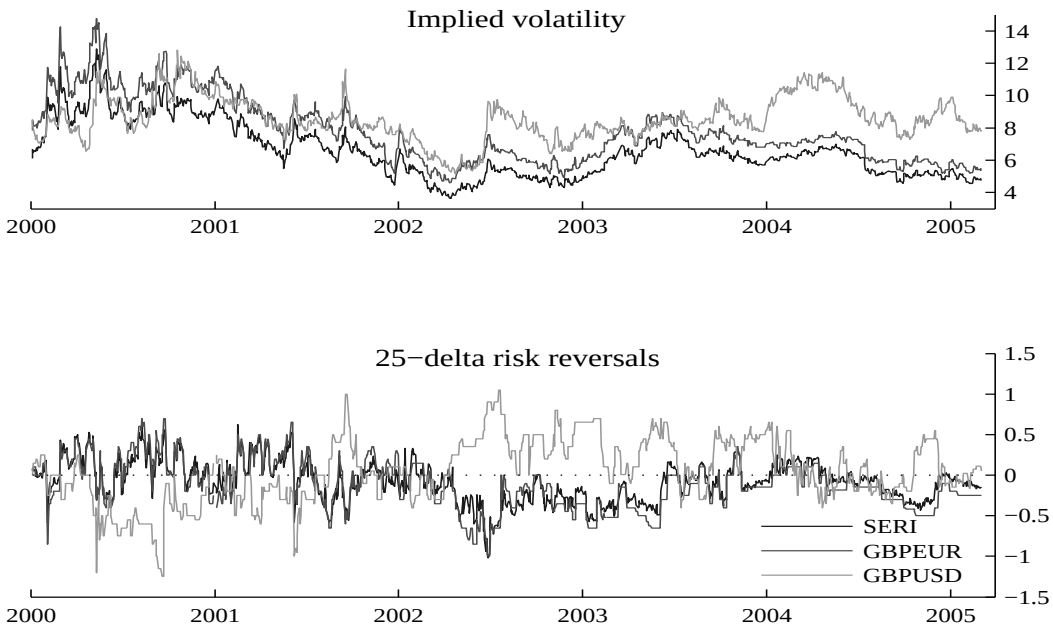

provides a joint density that is consistent with the marginals of three currency pairs over their entire support. This is achieved by minimising the L^2-distance between the actual PDF of the cross-rate (third bilateral) and the corresponding PDF implied by a copula function. Previous studies have fitted the joint distribution either to the implied correlation coefficient (Bikos (2000) and Taylor and Wang (2004)) or to a set of individual option contracts (Bennett and Kennedy (2004)).

One of our main findings is that standard one-parameter copulas lack the flexibility to consistently model the relationship between risk-neutral FX densities. This is especially true when one or more of the margins is heavily skewed. The same failure occurs with more sophisticated two-parameter copulas and the perturbed normal copula of Bennett and Kennedy that can account for asymmetry along the 45 degree line. A more general dependence structure is clearly needed to capture the relationships between foreign exchange rates. We solve this problem by using a Bernstein copula, which permits a close approximation to any possible dependence function.

We apply our method to estimate the risk-neutral distribution of a simplified exchange rate index for sterling. Our results indicate that the shape, as well as the location of this distribution has seen

considerable change over time. An interesting observation is that the large depreciation of sterling in 2003 was pre-dated by a negative skew and an increase in kurtosis of the SERI distribution. We also show how the estimated bivariate distribution can be used to construct densities of the exchange rate index that are conditional on hypothetical movements of one currency pair.

In a second application we use our empirical framework to price currency index options, ie options whose underlying is a geometric mean of individual exchange rates. In an example we demonstrate how standard contracts such as straddles, risk reversals and butterflies can be priced for options of an index of the euro and the yen versus the dollar. We compare our method to the standard multivariate Black model, which assumes joint log-normality, as well as an *ad hoc* adjustment based on the actual smiles and implied correlation. In our example we find that volatility quotes differ significantly across different methods.

We believe that there are several reasons why the methodology of this paper is useful for practitioners. Our framework is general in the sense that it does not make any assumptions on the marginal distributions. It is also straightforward to extend our option-pricing example to other multi-asset currency options such as quantos, basket options, spread options and rainbow options. Finally, the Bernstein copula exhausts the space of all possible parametric copulas. A failure to provide a match therefore implies that the individual margins are not arbitrage-free.

We envisage several avenues of future research. First, it would be interesting to extend our methodology to more then three exchange rates. For example, a three-dimensional copula could be used to link the six bilaterals between the dollar, the euro, the yen and sterling. Second, we believe that our method could be adapted to price options whose payout is path-dependent, such as barriers.

From the perspective of a policymaker, an important caveat is the fact that the option-implied distributions discussed in this paper are computed under a risk-neutral probability measure. In the presence of risk premia, they are therefore different from real-world density forecasts. Finding a relationship between the risk-neutral and the real-world distributions is an interesting, as well as challenging, area of ongoing research,[18] but beyond the scope of this paper.

[18] See, for example, Bliss and Panigirtzoglou (2004).

Finally, we would like to emphasise the fact that the statistical model of this paper has the unique feature that option-implied densities in conjunction with the triangular nature of exchange rates provide us with information on the entire marginal distributions as well as their dependence structure on a daily frequency. For many other applications much less information is available and the margins, as well as the copula, need to be estimated over a slow-moving time horizon. We therefore believe that the framework of this paper may serve as a laboratory environment to evaluate the reliability of other copula methods which rely on less information.

Appendix A: The relationship between risk-neutral densities under different numeraire measures

Let $z_T^{a,b}$, $z_T^{a,c}$, $z_T^{b,c}$ denote the logarithmic deviation of three triangular exchange rates $S_T^{a,b}$, $S_T^{a,c}$, $S_T^{b,c}$ from their respective forward rates $M_{t,T}^{a,b}$, $M_{t,T}^{a,c}$, $M_{t,T}^{b,c}$. We show that at any time $t \leq T$ the relationship between the univariate PDF of $z_T^{a,b}$ under the risk-neutral measure $\mathbb{Q}_a$ [19] and the bivariate PDF of $z_T^{a,c}$ and $z_T^{b,c}$ under the risk-neutral measure $\mathbb{Q}_c$ is given by

$$f^{\mathbb{Q}_a}_{z_T^{a,b},t}(x) = \int_{-\infty}^{\infty} f^{\mathbb{Q}_c}_{z_T^{a,c},z_T^{b,c},t}(u, u-x)e^{-u}du \qquad \textbf{A-1}$$

To see this,[20] consider the following no-arbitrage condition between ratios of spots and forwards ($Z_T^{i,j} \equiv S_T^{i,j}/M_{t,T}^{i,j}$),

$$\left(Z_T^{a,b} - K\right)^+ = Z_T^{a,c}\left(Z_T^{c,b} - Z_T^{c,a}K\right)^+ \qquad \textbf{A-2}$$

That is, the pay-off of a European call option on currency a, b with strike K is the same as $Z_T^{a,c}$ times the pay-off of a European call option on currency c, b with strike $Z_T^{c,a}K$.

Discounting both sides under the respective equivalent martingale measures we can write this equality in terms of expectations at time t as

$$e^{-r_a(T-t)}\mathbb{E}_t^{\mathbb{Q}_a}\left[\left(Z_T^{a,b} - K\right)^+\right] = \frac{S_t^{a,c}}{M_{t,T}^{a,c}}e^{-r_c(T-t)}\mathbb{E}_t^{\mathbb{Q}_c}\left[\left(Z_T^{c,b} - Z_T^{c,a}K\right)^+\right] \qquad \textbf{A-3}$$

---

[19] More precisely the risk-neutral measure $\mathbb{Q}_j$ is the equivalent martingale measure associated with a discount bond in currency j.

[20] We follow a similar line of reasoning as Taylor and Wang (2004), who derive a relationship between distributions of exchange rates in levels.

which simplifies to

$$\mathbb{E}_t^{\mathbb{Q}_a}\left[\left(Z_T^{a,b}-K\right)^+\right] = \mathbb{E}_t^{\mathbb{Q}_c}\left[\left(Z_T^{c,b}-Z_T^{c,a}K\right)^+\right], \quad \textbf{A-4}$$

$$\text{or} \quad \int_K^\infty (s-K) f^{\mathbb{Q}_a}_{Z_T^{a,b},t}(s)ds = \int_0^\infty \int_{uK}^\infty (v-uK) f^{\mathbb{Q}_c}_{Z_T^{c,a},Z_T^{c,b},t}(u,v)dvdu$$

Differentiating both sides twice with respect to K gives

$$f^{\mathbb{Q}_a}_{Z_T^{a,b},t}(K) = \int_0^\infty f^{\mathbb{Q}_c}_{Z_T^{c,a},Z_T^{c,b},t}(u,uK)u^2du \quad \textbf{A-5}$$

By a change of variable, $z^{i,j} = \ln Z^{i,j}$, we then get

$$f^{\mathbb{Q}_a}_{z_T^{a,b},t}(x) = \int_{-\infty}^\infty f^{\mathbb{Q}_c}_{z_T^{c,a},z_T^{c,b},t}(u,u+x)e^u du \quad \textbf{A-6}$$

or

$$f^{\mathbb{Q}_a}_{z_T^{a,b},t}(x) = \int_{-\infty}^\infty f^{\mathbb{Q}_c}_{z_T^{a,c},z_T^{b,c},t}(u,u-x)e^{-u} du \quad \textbf{A-7}$$

Appendix B: The Bernstein copula and a simple representation of the optimisation problem

The underlying idea of the Bernstein copula is to define a function $\alpha(\omega)$ on a set of grid points and then use a polynomial expansion to extend the function to all points in the unit square. In our application we use an evenly spaced grid of $(m+1)^2$ points, $\omega = \frac{k}{m} \times \frac{l}{m}$, $k,l = 0,...,m$. The bivariate Bernstein copula is then defined as

$$C^B(u,v) = \sum_{k=0}^m \sum_{l=0}^m \alpha\left(\frac{k}{m},\frac{l}{m}\right) P_{k,m}(u)P_{l,m}(v) \quad \textbf{B-1}$$

where

$$P_{j,m}(x) = \begin{pmatrix} m \\ j \end{pmatrix} x^j (1-x)^{m-j}$$

is the j-th Bernstein polynomial of order m (for $j = 0,...,m$). Chart B1 plots the Bernstein polynomial basis $\{P_{j,m}\}_{j=0}^m$ for $m = 5$. Sancetta and Satchell (2004) show that this function will be a copula as long as $\alpha(\omega)$ satisfies the basic three conditions of a copula (grounded, consistent with margins and two increasing) for all points on the grid.

Similarly, the density of the bivariate Bernstein copula is given by

$$c^B(u,v) = m^2 \sum_{k=0}^{m-1} \sum_{l=0}^{m-1} \beta\left(\frac{k}{m},\frac{l}{m}\right) P_{k,m-1}(u)P_{l,m-1}(v) \quad \textbf{B-2}$$

Chart B1: Bernstein polynomials for *m*=5

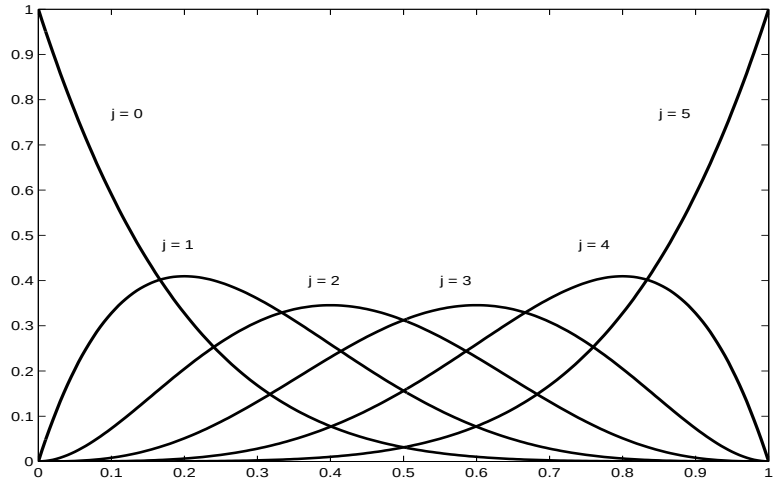

where

$$\beta\left(\frac{k}{m},\frac{l}{m}\right)=\alpha\left(\frac{k+1}{m},\frac{l+1}{m}\right)-\alpha\left(\frac{k+1}{m},\frac{l}{m}\right)-\alpha\left(\frac{k}{m},\frac{l+1}{m}\right)+\alpha\left(\frac{k}{m},\frac{l}{m}\right)$$

Note that the two-increasing property of α ensures that the density is non-negative.

We now show how we can write our optimisation problem **(6)** as a simple constrained quadratic programming problem. Recall from equation **(7)** that we approximate the euro-dollar density as

$$\hat{f}_z^{\mathbb{Q}_{\mathrm{E}}}(s,\hat{\theta})=\int_{-\infty}^{\infty}\hat{c}\left(F_x^{\mathbb{Q}_{\mathrm{S}}}(u),F_y^{\mathbb{Q}_{\mathrm{S}}}(u-s),\hat{\theta}\right)f_x^{\mathbb{Q}_{\mathrm{S}}}(u)f_y^{\mathbb{Q}_{\mathrm{S}}}(u-s)e^{-u}du \qquad \textbf{B-3}$$

Using the Bernstein copula, this can be written as

$$\hat{f}_z^{\mathbb{Q}_{\mathrm{E}}}(s,\hat{\theta})=\sum_{k=0}^{m-1}\sum_{l=0}^{m-1}\theta_{k,l}\psi_{k,l}(s) \qquad \textbf{B-4}$$

where $\theta_{k,l}=\beta\left(\frac{k}{m},\frac{l}{m}\right)$ and

$$\psi_{k,l}(s)=m^2\int_{-\infty}^{\infty}P_{k,m-1}\left(F_x^{\mathbb{Q}_{\mathrm{S}}}(u)\right)P_{l,m-1}\left(F_y^{\mathbb{Q}_{\mathrm{S}}}(u-s)\right)f_x^{\mathbb{Q}_{\mathrm{S}}}(u)f_y^{\mathbb{Q}_{\mathrm{S}}}(u-s)e^{-u}du \qquad \textbf{B-5}$$

These functions have the properties that $\psi_{k,l}(\cdot)\geq 0$ and $\int_{-\infty}^{\infty}\psi_{k,l}(s)ds=1$, for all $k,l=0,...,m-1$.

Proposition 1 (Restrictions on θ) Expression (**B-5**) is a copula density if and only if

$$\theta_{k,l} \geq 0, \quad k,l = 0,...,m-1 \qquad \textbf{B-6}$$

$$\sum_{k=0}^{m-1} \theta_{k,l} = \frac{1}{m}, \quad l = 0,...,m-1 \quad \text{and} \qquad \textbf{B-7}$$

$$\sum_{l=0}^{m-1} \theta_{k,l} = \frac{1}{m}, \quad k = 0,...,m-1 \qquad \textbf{B-8}$$

These restrictions also imply that the sum of all coefficients equals unity.

The optimisation problem can be restated as

$$\inf_{\{\theta_{k,l}\}_{k,l=1}^{m-1}} \int_{-\infty}^{\infty} \left(\sum_{k=0}^{m-1} \sum_{l=0}^{m-1} \theta_{k,l} \psi_{k,l}(s) - \tilde{f}_z^{\mathbb{Q}_E}(s) \right)^2 ds \qquad \textbf{B-9}$$
$$s.t. \text{ restrictions on } \{\theta_{k,l}\}_{k,l=0}^{m-1}$$

which can be simplified to

$$\inf_{\theta} \theta' \mathbf{H} \theta - 2\mathbf{g}\theta, \qquad s.t.\ \mathbf{R}_1\theta \leq \mathbf{q}_1,\ \mathbf{R}_2\theta = \mathbf{q}_2 \qquad \textbf{B-10}$$

where

$$\mathbf{H} = \int_{-\infty}^{\infty} \psi(s)\psi'(s)ds, \qquad \mathbf{g} = \int_{-\infty}^{\infty} f_z^{\mathbb{Q}_E}(s)\psi'(s)ds$$

and $\mathbf{R}_j$ and $\mathbf{q}_j$ impose the equality ($j = 1$) and inequality ($j = 2$) constraints of Proposition 1. This allows the problem to be solved by a standard Lagrangian approach.

Appendix C: Description of other copula functions

The Normal copula:

$$C(u, v; \rho) = \Phi_\rho(\zeta_1, \zeta_2) \qquad \textbf{C-1}$$

and

$$c(u, v; \rho) = \frac{1}{\sqrt{1-\rho^2}} \exp\left(\frac{\zeta_1^2 + \zeta_2^2}{2} + \frac{2\rho\zeta_1\zeta_2 - \zeta_1^2 - \zeta_2^2}{2(1-\rho^2)} \right) \qquad \textbf{C-2}$$

where $\zeta_1 \equiv \Phi^{-1}(u)$, $\zeta_2 \equiv \Phi^{-1}(v)$, $\Phi(\cdot)$ is the standard normal CDF and $\Phi_\rho(\cdot,\cdot)$ is the bivariate normal CDF with correlation coefficient $\rho \in [-1,1]$. A joint distribution is normal if, and only if, both the copula and the margins are normal.

The Frank copula:

$$C(u,v;\delta) = -\frac{1}{\delta}\log\left\{\left[\eta - g_u g_v\right]/\eta\right\} \qquad \textbf{C-3}$$

and

$$c(u,v;\delta) = \frac{-\delta g_1(1+g_{u+v})}{(g_u g_v + g_1)^2} \qquad \textbf{C-4}$$

where $g_z \equiv e^{-\delta z} - 1$ and $\delta \geq 0$.

The Plackett copula:

$$C(u,v;\delta) = \frac{1}{2}\eta^{-1}\left\{1+\eta(u+v) - \left[(1+\eta(u+v))^2 - 4\delta\eta uv\right]^{1/2}\right\} \qquad \textbf{C-5}$$

and

$$c(u,v;\delta) = \left[(1+\eta(u+v))^2 - 4\delta\eta uv\right]^{-3/2}\delta\left[1+\eta(u+v-2uv)\right] \qquad \textbf{C-6}$$

where $\eta = (\delta - 1)$ and $\delta \geq 0$.

The BB1 copula:

$$C(u,v;\theta,\delta) = \phi^{-1}\left(\phi(u) + \phi(v)\right) \qquad \textbf{C-7}$$

where $\phi(s;\theta,\delta) = \left(s^{-\theta} - 1\right)^{\delta}$ is an Archimedian generator with $\theta > 0$, $\delta \geq 1$, and

$$c(u,v;\theta,\delta) = -\frac{\phi''(C(u,v))\phi'(u)\phi'(v)}{\phi'(C(u,v))^3} \qquad \textbf{C-8}$$

The BB7 copula: This copula is also known as the Joe-Clayton copula and has a similar structure as the BB1 copula, except that the Archimedian generator is given by
$\phi(s;\theta,\delta) = \left[1-(1-s)^{\theta}\right]^{-\delta} - 1$, for $\theta \geq 1$, $\delta > 0$.

The (Generalised) Asymmetric Gumbel copula:

$$C(u,v;\alpha_0,\alpha_1,\alpha_2,\delta) = u^A v^B e^{-\alpha_0 H^{\frac{1}{\delta}}} \qquad \textbf{C-9}$$

where

$$A = \frac{\alpha_1}{\alpha_0+\alpha_1}, \quad B = \frac{\alpha_2}{\alpha_0+\alpha_2}, \quad H = \left(\frac{\tilde{u}}{F}\right)^{\delta} + \left(\frac{\tilde{v}}{G}\right)^{\delta}$$

and

$$F = \alpha_0 + \alpha_1, \quad G = \alpha_0 + \alpha_2, \quad , \tilde{u} = -\log u, \quad \tilde{v} = -\log v$$

The density is given by

$$c(u,v;\alpha_0,\alpha_1,\alpha_2,\delta) = \frac{C}{uv}\left(T_1T_2+T_3\right) \qquad \textbf{C-10}$$

where

$$\begin{aligned} T_1 &= A + \frac{\alpha_0}{F}\left(\frac{\tilde{u}}{F}\right)^{\delta-1} H^{\frac{1}{\delta-1}} \\ T_2 &= B + \frac{\alpha_0}{G}\left(\frac{\tilde{b}}{G}\right)^{\delta-1} G^{\frac{1}{\delta-1}} \\ T_3 &= \frac{\alpha_0(\delta-1)}{FG}\left(\frac{\tilde{u}\tilde{v}}{FG}\right)^{\delta-1} H^{\frac{1}{\delta-2}} \end{aligned}$$

The Perturbed Normal copula: Bennett and Kennedy (2004) construct a perturbed version of the normal copula C_N by using the transformation

$$C_\varphi(u,v) = \varphi^{-1}\left(C(\varphi(u),\varphi(v))\right)$$

where the transformation function $\varphi : [0,1] \to [0,1]$ is modelled as a cubic spline that satisfies (i) $\varphi(0) = 0$, (ii) $\varphi(1) = 1$ and (iii) $\varphi(\cdot)$ is concave. We use the same specification as them with knotpoints at $[0, 0.1, 0.5, 0.9, 1]$.

Appendix D: Estimating the marginal distributions

We use two data sets of over-the-counter (OTC) option contracts. The first data set, provided by Citibank, consists of 10 and 25 delta risk reversals and strangles, as well as ATM and forward contracts with one-month maturity. The second data set, provided by the British Bankers' Association (BBA), contains 25 delta risk reversals and strangles, as well as ATM contracts with maturities of one, three and twelve months. We further use spot and Libor rates, which are also provided by the BBA.

For the BBA data set we estimate PDFs using the Malz (1999) method, which fits a quadratic function to the volatility smile as a function of (spot) delta. The relationship between call prices and deltas are computed by the Black (1976) model.

For the Citibank data set, which includes additional 10 delta contracts, we extend the Malz method

in the sense that we fit the five available contracts by a spline consisting of two cubics (in the intervals between 0.1 and 0.25 and 0.75 and 0.9, respectively) and a quartic (in the interval between 0.25 and 0.75). We impose the restriction that the first three derivatives are continuous.

In both cases the PDF is computed as the (numerical) second derivative of the call price function corresponding to the interpolated volatility smile.

In order to adapt the (level) PDFs to the set-up connecting the three margins (equation **(3)**), we use the transformation

$$f^{\mathbb{Q}_d}_{S^{b,a}}(x) = f^{\mathbb{Q}_d}_{S^{a,b}}\left(\frac{1}{x}\right)\frac{1}{x^2} \qquad \textbf{D-1}$$

to exchange numerator and denominator, and the transformation

$$f^{\mathbb{Q}_a}_{S^{a,b}}(x) = f^{\mathbb{Q}_b}_{S^{a,b}}(x)\frac{M^{a,b}_{t,T}}{x} \qquad \textbf{D-2}$$

where $M^{a,b}_{t,T}$ is the forward rate, to change the numeraire measure. For example, we initially compute a PDF for dollar per euro under a dollar numeraire measure and then transform it into a PDF for euro per dollar under a euro numeraire measure. We use the relationship

$$f_{z^{a,b}}(x) = f_{S^{a,b}}\left(\ln\left(\frac{x}{M^{a,b}_{t,T}}\right)\right)\frac{1}{x} \qquad \textbf{D-3}$$

to transform the level PDF $f_{S^{a,b}}(\cdot)$ into the return PDF $f_{z^{a,b}}(\cdot)$.

Appendix E: Black's model for pricing an index option

We consider a European call option whose payout depends on the value of the index

$$I_t = \left(\frac{M^{c,a}_{t,T}}{M^{c,a}_{0,T}}\right)^{\omega_a}\left(\frac{M^{c,b}_{t,T}}{M^{c,b}_{0,T}}\right)^{(1-\omega_a)} = A\left(M^{c,a}_{t,T}\right)^{\omega_a}\left(M^{c,b}_{t,T}\right)^{(1-\omega_a)} \qquad \textbf{E-1}$$

on expiry ($t = T$). Under standard assumptions, the forward rates follow driftless diffusion processes under the numeraire measure $\mathbb{Q}_c$ (time-subscripts are suppressed for ease of notation)

$$dM^{c,i} = \sigma_{c,i}M^{c,i}dW^{\mathbb{Q}_c}_i, \quad i \in \{a, b\} \qquad \textbf{E-2}$$

where $W^{\mathbb{Q}_c}_a$ and $W^{\mathbb{Q}_c}_b$ are standard Brownian motions with correlation coefficient ρ (such that $dW^{\mathbb{Q}_c}_a dW^{\mathbb{Q}_c}_b = \rho dt$). Applying Itô's lemma gives

$$\begin{aligned} dI &= \sum_i \frac{\partial I}{\partial M^{c,i}}dM^{c,i} + \frac{1}{2}\sum_i\sum_j \frac{\partial^2 I}{\partial M^{c,i}\partial M^{c,j}}dM^{c,i}dM^{c,j} \\ &= \delta_I I dt + \sigma_I I dW^{\mathbb{Q}_c}_I \end{aligned} \qquad \textbf{E-3}$$

where $\delta_I = (\omega_a^2 - \omega_a)\left(\frac{1}{2}\sigma_{c,a}^2 + \frac{1}{2}\sigma_{c,b}^2 - \sigma_{c,a}\sigma_{c,b}\rho\right)$,
$\sigma_I^2 = \omega_a^2\sigma_{c,a}^2 + (1-\omega_a)^2\sigma_{c,b}^2 + 2\omega_a(1-\omega_a)\sigma_{c,a}\sigma_{c,b}\rho$ and $W_I^{\mathbb{Q}_c}$ is a standard Brownian motion under the equivalent martingale measure. Setting $\mathbb{E}_t^{\mathbb{Q}_c}[I_T] - M_{t,T}^I$ equal to zero gives the arbitrage-free forward rate for the index

$$M_{t,T}^I = I_t e^{\delta_I(T-t)} \qquad \textbf{E-4}$$

The call price for an option with strike K is then given by

$$C(I_t, K, \sigma_I, T) = e^{-r_c(T-t)}\left[M_{t,T}^I\Phi(d_1) - K\Phi(d_2)\right] \qquad \textbf{E-5}$$

where $\Phi(\cdot)$ denotes the standard normal distribution and

$$d_1 = \frac{\log\left(M_{t,T}^I/K\right) - \frac{\sigma_I^2}{2}(T-t)}{\sigma_I\sqrt{T-t}} \quad \text{and} \quad d_2 = d_1 - \sigma_I\sqrt{T-t}.$$

References

Bennett, M and Kennedy, J (2004), 'Quanto pricing with copulas', *Journal of Derivatives*, Vol. 12, No. 1, pages 26–45.

Bernstein, S (1912), 'Démonstration du théorème de Weierstrass fondée sur le calcul des probabilités', *Comm Soc Math Kharkov*, Vol. 13, pages 1–2.

Bhargava, V and Clark, J (2003), 'Pricing U.S. dollar index futures options: an empirical investigation', *The Financial Review*, Vol. 38, pages 571–90.

Bikos, A (2000), 'Bivariate FX PDFs: a sterling ERI application', *mimeo*.

Black, F (1976), 'The pricing of commodity contracts', *Journal of Financial Economics*, Vol. 3, pages 167–79.

Bliss, R and Panigirtzoglou, N (2002), 'Testing the stability of implied probability density functions', *Journal of Banking and Finance*, Vol. 23, No. 2-3, pages 381–422.

Bliss, R and Panigirtzoglou, N (2004), 'Option-implied risk aversion estimates', *Journal of Finance*, Vol. 59, pages 407–46.

Breeden, D and Litzenberger, R (1978), 'Prices of state-contingent claims implicit in option prices', *Journal of Business*, Vol. 51, pages 621–51.

Campa, J and Chang, K (1998), 'The forecasting ability of correlations implied in foreign exchange options', *Journal of International Money and Finance*, Vol. 17, No. 6, pages 855–80.

Cherubini, U, Luciano, E and Vecchiato, W (2004), *Copula methods in finance*, London: John Wiley & Sons.

Eytan, T and Harpaz, G (1986), 'The pricing of futures and options contracts on the Value Line Index', *Journal of Finance*, Vol. 41, No. 4, pages 843–55.

Joe, H (1997), *Multivariate models and dependence concepts*, London: Chapman & Hall.

Kendall, M (1938), 'A new measure of rank correlation', *Biometrica*, Vol. 30, pages 81–93.

Khoudraji, A (1995), 'Contribution à l'étude des copules et à la modélisation de valeurs extrêmes multivariées', Doctoral thesis, Université Laval.

Li, X, Mikusinski, P and Taylor, M (1998), 'Strong approximations of copulas', *Journal of Mathematical Analysis and Applications*, Vol. 225, No. 2, pages 608–23.

Lynch, B and Whitaker, S (2004), 'The new sterling ERI', *Bank of England Quarterly Bulletin*, Winter, pages 429–41.

Malz, A (1997), 'Estimating the probability distribution of the future exchange rate from option prices', *Journal of Derivatives*, Vol. 5, pages 18–36.

Mari, D and Monbet, V (2004), 'Construction de copules asymétriques avec dépendance extrême. Application à des données océanoqraphiques', *mimeo*, Laboratoire Sabres.

Patton, A (2006), 'Modelling asymmetric exchange rate dependence', *International Economic Review*, Vol. 47, No. 2, pages 527–56.

Rosenberg, J (2003), 'Non-parametric pricing of multivariate contingent claims', *Journal of Derivatives*, Vol. 10, pages 9–26.

Sancetta, A (2003), 'Nonparametric estimation of multivariate distributions with given marginals: L^2 theory', Cambridge University, Department of Applied Economics, *Working Paper No. 320*.

Sancetta, A and Satchell, S (2004), 'The Bernstein copula and its applications to modelling and approximations of multivariate distributions', *Econometric Theory*, Vol. 20, No. 3, pages 535–62.

Sklar, A (1959), 'Fonctions de répartition à n dimensions et leurs marges', *Publications de l'Institut Statistique de l'Université Paris*, Vol. 8, pages 229–31.

Taylor, S and Wang, Y-H (2004), 'Option prices and risk-neutral densities for currency cross-rates', University of Lancaster, *mimeo*.