

The most general methodology to create a valid correlation matrix for risk management and option pricing purposes

Riccardo Rebonato

Peter Jäckel

Quantitative Research Centre, NatWest Group, 135 Bishopsgate, EC2M 3UR

2nd November 1999

Abstract

Ad-hoc correlation matrices, those calculated from incomplete data, and those taken from news services sometimes don't comply with the requirement of symmetry and positive semi-definiteness. Whilst it is easy to amend the symmetry requirement by manual intervention, it is not always straightforward to see how to adjust the given correlation matrix to become usable for factor analysis or simulation purposes. In this document, two methods are described that can be used to best-match an invalid correlation matrix given the constraint of positive-semidefiniteness.

1 Introduction and motivation

The problem of how to specify a correlation matrix occurs in several important areas of finance and of risk management. A few of the important applications are, for instance, the specification of a (possibly time-dependent) instantaneous correlation matrix in the context of the BGM interest-rate option models, stress-testing and scenario analysis for market risk management purposes, or the specification of a correlation matrix amongst a large number of obligors for credit-derivative pricing or credit risk management.

For those applications where the most important *desideratum* is the recovery of the real-world correlation matrix, the problem is in principle well defined and readily solvable by means of well-established statistical techniques. In practice, however, the estimation problems can be severe: a small number of outliers, for instance, can seriously “pollute” a sample; non-synchronous data can easily destroy or hide correlation patterns; and the discontinuities in the correlation surface amongst forward rates when moving from deposit rates to the future strip, and from the latter to the swap market are well known to practitioners. In all these cases, the user often resorts to best-fitting the “noisy” elements of the sample correlation matrix by means of a plausible parametric function. This is, for instance, the route taken by Rebonato [9] for his calibration of the BGM model.

Similarly, when faced with the problem of pricing credit derivatives such as, for instance, n^{th} -to-default credit swaps, practitioners often posit or adjust (rather than simply estimate) a correlation matrix between intuitively understandable factors, such as country, industrial sector, *etc.*

Even when the correlation matrix econometrically estimated is thought to be reliable, a risk manager often wants to alter it in an *ad hoc* fashion, either as a stress test or in the context of scenario analysis. It is well known, for instance, that, in the event of equity market crashes, the correlation between different equity markets can dramatically increase (see Boyer *et al.* [1]). A risk manager in charge of a portfolio made up of, say, several national equity indices would greatly overestimate the degree of diversification in his portfolio in the event of a crash if he used the matrix estimated during “normal” periods.

Another important situation where straightforward econometric estimation of a correlation matrix is not possible is the case of legacy currencies (such as French Franc, the Lira, *etc.*, which have all been subsumed in the Euro). To tackle this problem both the option pricer and the risk manager have to invent, on the basis of financial intuition, a correlation matrix for a currency for which no history exists.

Finally, and perhaps most importantly, statistical estimation techniques based on the analysis

of historical data are intrinsically ill suited to provide predictions of future quantities when a paradigm shift has taken place. Such a paradigm shift could be, for instance, the imposition of currency controls for an emerging market, the granting of independence of a central Bank, a financial crisis such as Russia’s default in 1998, *etc.* In all these cases, it is to some extent possible to glimpse, from the traded prices of options, forward-looking estimates of volatilities that can supplement or replace the intrinsically backward-looking econometric estimates. The information available in derivative products about the future correlation matrix is, however, far too convoluted¹ to be of practical use. Whenever one therefore believed that the past might not be a reliable guide to the future, one could not realistically hope to distil a market consensus about the future correlation from the options market, and one would have to resort to altering an estimated matrix on the basis of financial intuition.

What all these different situations (and many more) have in common is therefore the desirability of altering a given correlation matrix deemed, for any reason, inadequate or inappropriate or of creating *tout court* a new one. The task, unfortunately, is not as simple as it might seem, given the requirement of a correlation matrix to be not only real symmetric, but also positive-semidefinite. This requirement, it must be stressed, does not belong to the “nice-to-have” list, but is an absolute requirement: if it were not fulfilled, for instance, a risk manager would have no guarantee that his VaR calculation, carried out using a variance-covariance normal approximation according to

$$VaR = k \cdot \mathbf{a}^T \cdot \mathbf{C} \cdot \mathbf{a} \quad (1)$$

will yield a positive number. In equation (1) above, \mathbf{a} is the vector of portfolio sensitivities to the market factors whose covariance matrix is denoted by \mathbf{C} and the constant k depends on the

¹The most promising area might appear to be the caplet/swaption markets. Rebonato [7, 8, 9] however shows that, in the absence of serial options, it is impossible to pin down uniquely from the prices of traded options both the time-dependent volatility of the forward rates and the correlation matrix amongst them. Furthermore, even if one assumed perfect knowledge of the instantaneous volatility functions for all the forward rates, swaption prices have very poor discriminatory power with respect to different correlation matrices, because they depend more strongly on the functional form of the instantaneous volatilities.

required percentile. Whilst the guarantee of the VaR as given by equation (1) being non-negative is usually not sufficient for it to be a meaningful risk measure², it is strictly necessary to make any kind of simulation possible when correlation between underlying assets is to be taken into account.

In practice, the problem of an invalid correlation matrix, *i.e.* one that has negative eigenvalues, can also very easily arise in the context of risk analysis for equity portfolios. This is because there are frequently asynchronous gaps in the historical stock exchange time series. The chance that slight inconsistencies in the data from which historical correlation coefficients are calculated can lead to negative eigenvalues grows rapidly as the size of the correlation matrix increases. This has recently been pointed out by Ju and Pearson [11]. Intuitively, it can be understood to be an effect of the characteristic polynomial that determines the eigenvalues becoming of higher order as the dimension of the correlation matrix grows, and thus displaying a stronger nonlinear response to slight changes in the polynomial's coefficients. Since equity index or portfolio analysis typically involves many underlying assets, the risk of negative eigenvalues of the correlation matrix calculated from historical data is particularly large.

The problem of how to specify a valid correlation matrix has been recognized in the literature and several papers have appeared on the subject (see, *e.g.* Kupiec [5], Finger [4], Brooks *et al.* [2]). The solutions proposed so far, however, have been only partial answers. The technique proposed by Finger, for instance, is designed to increase portions of the correlation matrix. Finger's approach (also reported in Brooks *et al.* [2]) has the drawback that some portions of the correlation matrix are altered in the desired fashion, but, in order to retain positive-semidefiniteness, the slack is taken up by other portions of the matrix in an uncontrolled fashion. The problems with

²The VaR can in principle be zero due to a non-vanishing kernel of the correlation matrix C . In this case, a trader or fund manager can in theory construct, by optimisation over the constituents, an apparently risk-free portfolio, *i.e.* a portfolio with zero Value at Risk. This fundamental shortcoming is not handled by the method proposed in this article. In order to reduce the risk of traders "arbitraging" a VaR measure by means of a perverse optimisation, it is rather recommendable to use one of the class of *coherent risk measures* [3] such as *expected Value at Risk* which is given by the conditional average over the tail of the profit & loss distribution beyond the required percentile.

this approach is clearly presented by Brooks *et al.*:

“...It [Finger’s] is not a simple procedure in practice, and represents the main drawback in J.P. Morgan’s approach [*i.e.* RiskMetrics] ... Note also that the other correlations in the matrix (that we had no *a priori* view on) have also changed, but this is the price that must be paid to ensure a consistent revised correlations matrix.”

The shrinkage technique (Kupiec [5]) is different in spirit and definitely more general. It proceeds by iterating towards a feasible solution that modifies as much as possible (in a non-easily quantifiable sense) a pre-existing well defined positive-semidefinite matrix towards a desired target correlation matrix. Despite its appeal and greater generality, its main drawbacks are that it can be rather time consuming (since it requires a full matrix orthogonalisation at each iteration); that it requires as a starting point and “anchor” a positive-semidefinite matrix to start with; and that, above all, there is no way of determining to what extent the resulting matrix is optimal in any easily quantifiable sense.

In order to obviate these shortcomings we present a method which:-

- i) is guaranteed to produce a positive-semidefinite matrix;
- ii) does not require a pre-existing acceptable matrix to start with;
- iii) is fast to implement even for large matrices;
- iv) allows the determination of a feasible matrix that most closely approximates a target real symmetric (but not positive-semidefinite) matrix in a well defined and quantifiable sense.

As for the last property of the method here presented, the user is at liberty to specify a suitable metric under which the technique is guaranteed to yield the best possible answer.

Furthermore, we present a second, and even faster, method, which shares properties i), ii) and iii) above, but which is not guaranteed to enjoy property iv)³. In all the empirical studies we have

³More precisely, there probably exists a metric with respect to which this method does provide the optimal solution but we have not been able to identify it.

carried out, however, the results obtained using this second approach have been extremely close, albeit not identical, to the ones obtained using the first technique. After presenting the main method, we therefore propose that the latter can be used either as an excellent fast approximation to the solution obtained using what is probably the most intuitive metric or as an initial guess for the general problem.

2 Hypersphere decomposition

The starting point are the well known results from linear algebra that every $n \times n$ matrix M given by

$$M = WW^T \quad (2)$$

for any $W \in \mathbb{R}^{n \times n}$ is positive-semidefinite and that, conversely, every positive-semidefinite matrix $M \in \mathbb{R}^{n \times n}$ can be decomposed as in equation (2).

The method we propose for the construction of a valid correlation matrix

$$\hat{C} = BB^T \quad (3)$$

that best-matches a given, *not* positive-semidefinite, target matrix C is to view the elements of the row vectors of matrix B in equation (3) as coordinates lying on a unit hypersphere [8, 9, 10].

If we denote by b_{ij} the elements of the matrix B , the key is to obtain the $n \times n$ coordinates b_{ij} from $n \times (n - 1)$ angular coordinates θ_{ij} according to

$$\begin{aligned} b_{ij} &= \cos \theta_{ij} \cdot \prod_{k=1}^{j-1} \sin \theta_{ik} \quad \text{for } j = 1 \dots n - 1 \\ &\text{and} \\ b_{ij} &= \prod_{k=1}^{j-1} \sin \theta_{ik} \quad \text{for } j = n \quad . \end{aligned} \quad (4)$$

For an arbitrary set of angles $\{\theta_{ij}\}$, a matrix \hat{C} formed from B as in equation (3) satisfies all the given constraints required of a correlation matrix by construction. In particular, thanks to the trigonometric relationship (4) and to the requirement that the radius of the unit hypersphere

should be equal to one, the main diagonal elements are guaranteed to be unity, see appendix A for a proof.

In general, matrix \hat{C} will bear no resemblance to the target matrix C . However, after using the above transformation and after defining a suitable error measure ε in the resulting approximate correlation matrix \hat{C}

$$\varepsilon = \left\| C - \hat{C} \right\| , \quad (5)$$

one can use an optimisation procedure over the angles θ_{ij} to find the best possible fit given the chosen error measure. Sensible choices for the error measure are:-

- The sum of squares of the elements of the difference matrix $(C - \hat{C})$,

$$\chi_{\text{Elements}}^2 := \sum_{ij} (c_{ij} - \hat{c}_{ij})^2 . \quad (6)$$

Since both C and \hat{C} have unit diagonal elements, this error norm is equal to twice the sum of squares of errors in the free correlation coefficients.

- The elementwise sum of squares of errors in the sorted sets of eigenvalues of C and \hat{C} ,

$$\chi_{\text{Eigenvalues}}^2 := \sum_i (\lambda_i - \hat{\lambda}_i)^2 . \quad (7)$$

Naturally, the above suggestions are only examples and various other choices are conceivable. If, in particular, a risk manager felt that certain portions of the target correlation matrix C should be recovered with particularly high accuracy, then correspondingly large weights could be assigned to the relative elements $(c_{ij} - \hat{c}_{ij})^2$. In this context, the error norm given by equation (6), where every element has exactly the same weight, has been shown by Rebonato [8, 9, 10] to have desirable global features in so far as the calibration of the BGM model is concerned.

The fundamental benefits of this method are twofold: first, when the underlying space over which the optimisation is carried out is expressed in terms of angle vectors describing coordinates on a unit hypersphere, no constraints have to be satisfied. This can be of substantial benefit for the numerical fitting procedure. More importantly, despite the fact that the procedure is still

iterative, unlike the shrinking method [5], it simply requires a matrix multiplication per iteration, rather than a full matrix diagonalisation⁴. This can make a big difference, especially for large matrices. Furthermore, the approach presented in the next section requires no iterations and provides a solution very similar to the one obtained using error metric (6). It can therefore be used to provide the starting point for the search procedure, thereby further reducing the overall computational cost.

3 Spectral decomposition⁵

Given the right-hand-side eigensystem S of the real and symmetric matrix C and its associated set of eigenvalues $\{\lambda_i\}$ such that

$$C \cdot S = \Lambda \cdot S \quad \text{with} \quad \Lambda = \text{diag}(\lambda_i) \quad , \quad (8)$$

define the non-zero elements of the diagonal matrix Λ' as

$$\Lambda' \quad : \quad \lambda'_i = \begin{cases} \lambda_i & : \lambda_i \geq 0 \\ 0 & : \lambda_i < 0 \end{cases} \quad . \quad (9)$$

If the target matrix C is not positive-semidefinite, it has at least one negative eigenvalue and at least one of the λ'_i will be zero.

Also, define the non-zero elements of the diagonal scaling matrix T with respect to the eigensystem S by

$$T \quad : \quad t_i = \left[\sum_m s_{im}^2 \lambda'_m \right]^{-1} \quad . \quad (10)$$

Now, let⁶

$$B' := S\sqrt{\Lambda'} \quad (11)$$

⁴It should be noted that for most optimisation procedures, the function defining the error norm has to be evaluated many times per iteration since typically some kind of Hessian or Jacobian matrix is also needed. Thus, the reduction in computational effort at this level can be crucial.

⁵also known as principal component analysis

⁶Please note that the notation \sqrt{D} for a diagonal matrix D with non-negative elements is a symbolic description of the diagonal matrix whose non-zero elements are the positive roots of the diagonal elements of D .

and

$$\mathbf{B} := \sqrt{\mathbf{T}}\mathbf{B}' = \sqrt{\mathbf{T}}\mathbf{S}\sqrt{\mathbf{\Lambda}'} \quad . \quad (12)$$

For normalised row vectors of \mathbf{S} , the truncation of the negative eigenvalues results in row vectors of \mathbf{B}' that are not of unit length. This is rectified in equation (12) by the aid of matrix \mathbf{T} which contains the required normalisation factors. By construction,

$$\hat{\mathbf{C}} := \mathbf{B}\mathbf{B}^T \quad (13)$$

is now both positive-semidefinite and has unit diagonal elements since its elements are

$$\begin{aligned} \hat{c}_{ij} &= \sum_{klmnp} (\sqrt{\mathbf{T}})_{ik} \cdot s_{kl} \cdot (\sqrt{\mathbf{\Lambda}'})_{lm} \cdot (\sqrt{\mathbf{\Lambda}'})_{mn} \cdot s_{pn} \cdot (\sqrt{\mathbf{T}})_{pj} \\ &= \sum_l \sqrt{t_i} \cdot s_{il} \cdot \lambda'_l \cdot s_{jl} \cdot \sqrt{t_j} \\ &= \frac{\sum_l s_{il}s_{jl}\lambda'_l}{\sqrt{\sum_m s_{im}^2 \lambda'_m \cdot \sum_k s_{jk}^2 \lambda'_k}} \quad . \end{aligned} \quad (14)$$

A procedural description of the above method may clarify what actually has to be done:

- Calculate the eigenvalues λ_i and the right-hand-side eigenvectors \mathbf{s}_i of \mathbf{C} .
- Set all negative λ_i to zero.
- Multiply the vectors \mathbf{s}_i with their associated “corrected” eigenvalues λ'_i and arrange as the columns of \mathbf{B}' .
- Finally, \mathbf{B} results from \mathbf{B}' by normalising the *row vectors* of \mathbf{B}' to unit length.

By following this procedure we obtain an acceptable correlation matrix which is intuitively “similar” to the target one (the more so, the fewer the eigenvalues which have to be set to zero). The crucial point, however, is not so much the plausibility of the metric but the fact that empirically we have always observed the results obtained using equations (8) to (13) to be very similar to those from the angular method discussed in section 2. How close the results are in practice is shown in the next section. This is significant because one can use the result of the method described here either as an accurate approximation to the best (in a χ_{Elements}^2 sense) solution, or as

the starting point for the optimisation discussed in section 2, thereby substantially reducing the computational burden of the hypersphere decomposition approach.

4 Examples

A risk manager retrieves from the middle office's reporting system the following correlation matrix of three world equity indices:

$$\tilde{\mathbf{C}} = \begin{pmatrix} 1 & 0.9 & 0.7 \\ 0.9 & 1 & 0.4 \\ 0.7 & 0.4 & 1 \end{pmatrix} .$$

The eigenvalues of $\tilde{\mathbf{C}}$ are $\{ 2.35364, 0.616017, 0.0303474 \}$ and the correlation matrix can be split up as in

$$\tilde{\mathbf{C}} = \tilde{\mathbf{B}}\tilde{\mathbf{B}}^T$$

with

$$\tilde{\mathbf{B}} = \begin{pmatrix} 0.98742 & 0.08718 & -0.13192 \\ 0.88465 & 0.45536 & 0.10021 \\ 0.77203 & -0.63329 & 0.05389 \end{pmatrix} .$$

The risk manager is aware of the Value at Risk calculated under the assumption of this correlation between the three indices. In order to assess the change in Value at Risk resulting from a decrease in correlation between two of the three underlying variables, the risk manager wishes to adjust the matrix to

$$\mathbf{C}' = \begin{pmatrix} 1 & 0.9 & 0.7 \\ 0.9 & 1 & 0.3 \\ 0.7 & 0.3 & 1 \end{pmatrix} .$$

Unfortunately, the eigenvalues of \mathbf{C}' are now $\{ 2.29673, 0.710625, -0.00735244 \}$, and despite its plausible appearance, matrix \mathbf{C}' is no longer an acceptable correlation matrix. This highlights how a minor change can lead to the violation of the requirement of positive-semidefiniteness of a correlation matrix. The system will now fail when trying to construct a split-up

matrix B for the purpose of Monte Carlo simulations⁷ in order to calculate the Value at Risk under the new assumptions.

Using the method outlined in section 2 with the error measure chosen to be χ_{Elements}^2 as given by equation (6), we can calculate

$$\hat{B} = \begin{pmatrix} 0.99804 & 0.06265 & 0 \\ 0.86482 & 0.50209 & 0 \\ 0.74020 & -0.67239 & 0 \end{pmatrix}$$

with

$$\hat{C} = \hat{B}\hat{B}^T = \begin{pmatrix} 1 & 0.89458 & 0.69662 \\ 0.89458 & 1 & 0.30254 \\ 0.69662 & 0.30254 & 1 \end{pmatrix}$$

and a total error of $\chi_{\text{Elements}}^2 = 0.946 \cdot 10^{-4}$.

In comparison, the method outlined in section 3 above, yields

$$\hat{B} = \begin{pmatrix} 0.99805 & 0.06238 & 0 \\ 0.86434 & 0.50292 & 0 \\ 0.73974 & -0.67290 & 0 \end{pmatrix}$$

to give us

$$\hat{C} = \hat{B}\hat{B}^T = \begin{pmatrix} 1 & 0.89402 & 0.69632 \\ 0.89402 & 1 & 0.30100 \\ 0.69632 & 0.30100 & 1 \end{pmatrix} .$$

Not only is the total error only $\chi_{\text{Elements}}^2 = 1.0 \cdot 10^{-4}$ but also the individual elements are remarkably close to the values obtained by optimisation. Despite the fact that there is in general no guarantee that the results of the two methods are as close together as in this example, we have always found very good agreement between the two approaches.

⁷The construction of correlated normal variates from a vector of uncorrelated normal variates z is done by the transformation $x = B \cdot z$ with $C = BB^T$.

5 Conclusion

We have presented two simple methods to produce a feasible (*i.e.* real, symmetric, and positive-semidefinite) correlation matrix when the econometric one is either noisy, unavailable, or inappropriate. The first method is to the knowledge of the authors more general than any of the approaches which have been proposed in the literature, and computationally faster. It can actually produce the optimal feasible solution in a sense specified by the user. The second method is, in principle, not as general, but we show that i) it is extremely fast and ii) it produces results very close to those obtained using the general procedure. It can therefore be used in its own right, or as a starting point for the general optimisation procedure, thereby making the latter even faster.

A Angular coordinates on a hypersphere of unit radius

The i -th row vector of B as specified by equation (4) is given by

$$\begin{aligned} & (b_{i1}, b_{i2}, \dots, b_{in-1}, b_{in}) \\ & = \\ & (\cos \theta_{i1}, \sin \theta_{i1} \cos \theta_{i2}, \dots, \prod_{k=1}^{n-2} \sin \theta_{ik} \cos \theta_{in-1}, \prod_{k=1}^{n-2} \sin \theta_{ik} \sin \theta_{in-1}) \quad . \end{aligned} \tag{15}$$

The sum of the squares $\sum_{j=1}^{j=n} b_{ij}^2$ will definitely be unity if the recursive relation

$$\sum_{k=j+1}^{k=n} b_{ik}^2 = b_{ij}^2 \cdot \frac{\sin^2 \theta_{ij}}{\cos^2 \theta_{ij}} \tag{16}$$

holds since then

$$\sum_{k=2}^{k=n} b_{ik}^2 = b_{i1}^2 \cdot \frac{\sin^2 \theta_{i1}}{\cos^2 \theta_{i1}} = \cos^2 \theta_{i1} \cdot \frac{\sin^2 \theta_{i1}}{\cos^2 \theta_{i1}} = \sin^2 \theta_{i1} \tag{17}$$

and thus

$$\sum_{k=1}^{k=n} b_{ik}^2 = b_{i1}^2 + \sum_{k=2}^{k=n} b_{ik}^2 = \cos^2 \theta_{i1} + \sin^2 \theta_{i1} = 1 \quad . \tag{18}$$

To start the induction, we see that equation (16) is satisfied for $j = n - 1$ since then we have

$$\prod_{k=1}^{n-2} \sin^2 \theta_{ik} \sin^2 \theta_{in-1} = \prod_{k=1}^{n-2} \sin^2 \theta_{ik} \cos^2 \theta_{in-1} \cdot \frac{\sin^2 \theta_{in-1}}{\cos^2 \theta_{in-1}} \quad . \tag{19}$$

It remains to be shown that

$$b_{ij}^2 = \frac{\cos^2 \theta_{ij}}{\sin^2 \theta_{ij}} \cdot \sum_{k=j+1}^{k=n} b_{ik}^2 \quad (20)$$

for $j < n - 1$. Using the recursion relation (16) itself, we obtain

$$\begin{aligned} b_{ij}^2 &= \frac{\cos^2 \theta_{ij}}{\sin^2 \theta_{ij}} \cdot \left(b_{i,j+1}^2 + \sum_{k=j+2}^{k=n} b_{ik}^2 \right) \\ &= \frac{\cos^2 \theta_{ij}}{\sin^2 \theta_{ij}} \cdot \left(b_{i,j+1}^2 + \frac{\sin^2 \theta_{i,j+1}}{\cos^2 \theta_{i,j+1}} \cdot b_{i,j+1}^2 \right) \\ &= \frac{\cos^2 \theta_{ij}}{\sin^2 \theta_{ij} \cos^2 \theta_{i,j+1}} \cdot b_{i,j+1}^2 \quad , \end{aligned}$$

i.e.

$$b_{i,j+1} = b_{ij} \cdot \frac{\sin \theta_{ij}}{\cos \theta_{ij}} \cos \theta_{i,j+1} \quad (21)$$

which is identical to the construction description (4) for $j < n - 1$. Hence, all row vectors of \mathbf{B} are of unit length. The elements of $\hat{\mathbf{C}} = \mathbf{B}\mathbf{B}^T$ are the pairwise scalar products of the row vectors of \mathbf{B} . Since the scalar product of two vectors of unit length is by definition $\in [-1, 1]$, $\hat{\mathbf{C}}$ satisfies the requirements of unit diagonal elements, symmetry, and all elements being $\in [-1, 1]$.

References

- [1] B. H. Boyer, M. S. Gibson, and M. Loretan. Pitfalls in Tests for Changes in Correlations. International Finance Discussion paper #597, December 1997.
- [2] C. Brooks, B. Scott-Quinn, and J. Whalmsey. Adjusting VaR Models for the impact of the Euro. Working paper of the ISMA Centre, 1998.
- [3] F. Delbaen, J. M. Eber, and D. Heath. Thinking Coherently. *Risk*, 10(11), November 1997.
- [4] C. Finger. A Methodology for Stress Correlation. *RiskMetrics Monitor*, pages 3–11, Fourth Quarter 1997.
- [5] P. H. Kupiec. Stress Testing in a Value at Risk Framework. *Journal of Derivatives*, pages 7–24, Fall 1998.

- [6] W. H. Press, S. A. Teukolsky, W. T. Vetterling, and B. P. Flannery. *Numerical Recipes in C*. Cambridge University Press, 1992. www.nrbook.com/a/bookcpdf.php.
- [7] R. Rebonato. *Interest Rate Option Models*. John Wiley and Sons, 1998.
- [8] R. Rebonato. Calibration of the BGM model. *Risk*, March:74–79, 1999.
- [9] R. Rebonato. On the simultaneous calibration of multifactor lognormal interest rate models to Black volatilities and to the correlation matrix. *Journal of Computational Finance*, 2(4):5–27, Summer 1999.
- [10] R. Rebonato. *Volatility and Correlation*. John Wiley and Sons, 1999.
- [11] J. Xiongwei and N. Pearson. Using Value-at-Risk to Control Risk Taking: How wrong can you be? *Journal of Risk*, 1(2):5–36, Winter 1999.