Calibrating probability distributions used for risk measurement purposes to market-implied data

Abstract

The following pages provide an introduction to how it is possible to calibrate probability distributions used for risk measurement purposes to market-implied data.

Contents

1. Introduction

1.1 Kemp (2005) and Kemp (2009) argue in favour of greater use of market-consistent risk management. This involves greater focus on market implied probability distributions and other risk parameters (consistent with the market prices of derivatives that might protect against the relevant risks) and lesser focus on estimating such parameters using time series analysis of historic market behaviour.

1.2 In practice, as is explained in Kemp (2009), there is insufficient market implied data available to be able to build up a *full* specification of the relevant market implied probability distribution. However, there is often *some* market implied data that we can use for calibration purposes. We explain in these pages how we can combine this data with an assumed prior distribution in a Bayesian or credibility-weighted type manner, thereby deriving risk measures that are as market consistent as possible, but still coloured by our prior views in cases where market implied data does not exist. A natural prior to adopt in this context, if we think the historic data is relevant, is one based on the relevant historic dataset.

1.3 A prototype methodology for market consistent risk measurement is described in Kemp (2009) involving an *analytical weighted Monte Carlo*, aka *analytical relative entropy* approach. This involves identifying a distribution that is, in some suitable sense, as 'close as possible' to the original prior distribution but that has characteristics that match calibration characteristics derived from market implied data (or otherwise).

1.4 As its name suggests, the analytical weighted Monte Carlo methodology has its genesis in the weighted Monte Carlo simulation approach, see Elices & Gimenez (2006) and Avellandeda et al (2001). In this approach, we calibrate a simulation exercise not by changing the draws from some previously specified prior probability distribution but by changing the assumed likelihoods ascribed to each draw in the computation of statistics in which we are interested (e.g. mean, standard deviation, other moments, quantiles etc.). Typically we choose these revised likelihoods of occurrence to minimise the *relative entropy* between the original prior distribution and the calibrated output distribution. The relative entropy between two discrete distributions $\mathbf{p}$ and $\mathbf{q}$ is given by the following, where i here indexes the possible outcomes (the formula for continuous distributions is derived in an analogous manner):

$$D(\mathbf{p},\mathbf{q}) = \sum_{i=1}^{N} p_i \log \frac{p_i}{q_i}$$

1.5 The analytical weighted Monte Carlo methodology notes that in the limit where the number of simulations $N \to \infty$ we will recover an exact output probability distribution and that in a number of cases of interest this probability distribution can be expressed in 'analytical' form, i.e. we may be able to recover relatively simply what the methodology would have produced had we been able to use an infinitely large Monte Carlo sample size.

1.6 In the special case of single instrument calibration involving a *univariate* normal prior distribution, say, $N(\mu,\sigma^2)$ we find, as we might expect, that application of an analytical weighted Monte Carlo returns a normal distribution with the following characteristics, see Kemp (2009):

(a) If we have no market data whatsoever to calibrate to then the calibrated distribution is the same as the prior distribution, i.e. $N(\mu,\sigma^2)$;

(b) If we have just a mean to calibrate to, say μ_0, then the calibrated distribution is $N(\mu_0,\sigma^2)$; and

(c) If we have both a mean and a variance to calibrate to, say, μ_0 and σ_0^2, then the calibrated distribution is $N(\mu_0,\sigma_0^2)$.

For risk management purposes, it is usual to adopt the assumption that the return on all assets is the same (i.e. to discount the possibility of 'manager skill', because it is prudent to do so). It can be argued that (b) provides the theoretical justification for this, i.e. in effect we are 'calibrating' some assumed prior distribution (that might include differential returns) to fit a constraint that requires any prior assumed manager skill in asset selection not actually to be present in practice.Without loss of generality the mean return can for the purposes of this paper be set to zero (since our focus will be on relative returns). Hence our focus will be on second and higher moments (or just second moments in the case of normally distributed variables, which is the main focus of these pages).

2. Multi-instrument calibration

2.1 Risk models in practice need to cater for multiple instruments. The most common framework involves assuming that the underlying (log) return distribution is multivariate normal. Traditionally, the corresponding covariance matrix is derived from historical observations although usually a parsimonious factor structure is imposed to limit the number of terms in the covariance matrix that need to be estimated from past history, see e.g. Kemp (2005) and Kemp (2009).

2.2 Calibrating a multivariate normal prior distribution to market-implied (or other) data is typically different to the univariate case because we will usually have fewer calibration points than we have degrees of freedom in relation to the number of terms in the now multi-dimensional covariance matrix. However, some of the principles noted in the univariate case still carry through to the multivariate case. In particular, if we are calibrating a multivariate normal prior distribution merely to market implied volatilities and covariances for a given fixed period (i.e. merely to second moments) then the resulting calibrated distribution will still be multivariate normal distribution, just with a different covariance matrix.

2.3 A multivariate normal distribution $N(\boldsymbol{\mu}, \boldsymbol{\Omega})$ with mean $\boldsymbol{\mu}$ (a vector of random variables) and covariance matrix $\boldsymbol{\Omega}$ has a probability density function, $q(\mathbf{x})$ as follows, where n is the number of entries in the vector $\boldsymbol{\mu}$ and n^2 the number of entries in $\boldsymbol{\Omega}$):

$$q(\mathbf{x}) = \frac{1}{(2\pi)^{n/2}|\Omega|^{1/2}} \exp\left(-\frac{1}{2}(\mathbf{x}-\boldsymbol{\mu})^T \boldsymbol{\Omega}^{-1}(\mathbf{x}-\boldsymbol{\mu})\right)$$

2.4 We note that:

(a) Any n-dimensional multivariate normal distribution has a probability density function expressible as $C \times \exp(-F)$ where C is some suitable constant and F is a positive definite symmetric quadratic form (with possibly non-zero drift) in n different variables, and vice versa.

(b) Applying analytical weighted Monte Carlo (using relative entropy) to the sort of calibration problem referred to above will therefore return (unless the calibration problem is ill-posed) a calibrated probability distribution which also has multivariate normal form. This is because the problem can be restated using Lagrange multipliers to one that involves minimising L defined as follows, where the λ_j refer to whatever calibrations there are on the means and λ_{jk} to those on covariance terms (in general there will be fewer than n^2 of the λ_{jk}):

$$L = \int p.\log\left(\frac{p}{q}\right) d\mathbf{x} + \sum_{i=1}^{n} \kappa_i \left(\int p \, d\mathbf{x} - 1\right) + \sum \lambda_j \left(\int x_j p \, d\mathbf{x} - m_j\right) + \sum \lambda_{jk} \left(\int (x_j - m_j)(x_k - m_k) p \, d\mathbf{x} - s_{jk}^2\right)$$

The solution to this minimisation problem is given by the following:

$$\frac{\partial L}{\partial p} = 0, \frac{\partial L}{\partial \kappa_i} = 0, \frac{\partial L}{\partial \lambda_j} = 0, \frac{\partial L}{\partial \lambda_{jk}} = 0$$

$$\Rightarrow \log(p) - \log(q) + 1 + \sum_{i=1}^{n} \kappa_i + \sum \lambda_j x_j + \sum \lambda_{jk}(x_j - m_j)(x_k - m_k) = 0$$

$$\Rightarrow p(\mathbf{x}) = q(\mathbf{x}) \exp\left(-\left(1 + \sum_{i=1}^{n} \kappa_i + \sum \lambda_j x_j + \sum \lambda_{jk}(x_j - m_j)(x_k - m_k)\right)\right)$$

subject to $\int p \, d\mathbf{x} = 1$ (i.e. that p is a probability distribution) and other constraints derived directly from calibration requirements, e.g. that $\int x_j p \, dx = m_j$ etc.

Thus if $q(\mathbf{x})$ is expressible as $C \times \exp(-F)$ as above, then $p(\mathbf{x})$ will be too, just for a different F.

(c) Applying the principle of no arbitrage we may therefore expect $p(\mathbf{x})$ to have zero mean (more precisely for each element of $\boldsymbol{\mu}$ to be the same, which without loss of generality we may take as zero if we are focusing on relative returns) and therefore to have the form:

$$p(\mathbf{x}) = q(\mathbf{x})\exp\left(-\left(1 + \sum_{i=1}^{n} \kappa_i + \sum \lambda_{jk} F_{jk}\right)\right)$$

where the F_{jk} are symmetric zero-drift quadratic forms ($= \mathbf{x}^T V_{jk} \mathbf{x}$, say) corresponding to each of the implied volatilities/implied correlations to which we wish to calibrate.

(d) The calibrated distribution will therefore be multivariate normal with zero mean and probability distribution as follows, for suitably chosen λ_{jk} that reproduce for the calibrations the relevant market implied variances or covariances (where c_p is some constant the value of which ensures that $\int p \, d\mathbf{x} = 1$):

$$p(\mathbf{x}) = \frac{1}{(2\pi)^{n/2}} \exp\left(-\frac{1}{2}\mathbf{x}^T \boldsymbol{\Omega}^{-1} \mathbf{x} - \sum \lambda_{jk} \mathbf{x}^T V_{jk} \mathbf{x}\right)$$

(e) Thus the calibrated probability distribution will be characterised by a covariance matrix $\overline{\boldsymbol{\Omega}}$ as follows:

$$-\frac{1}{2}\mathbf{x}^T \overline{\boldsymbol{\Omega}}^{-1} \mathbf{x} = -\frac{1}{2}\mathbf{x}^T \boldsymbol{\Omega}^{-1} \mathbf{x} - \sum \lambda_{jk} F_{jk}$$
$$\Rightarrow \overline{\boldsymbol{\Omega}} = \left(\boldsymbol{\Omega}^{-1} + 2\sum \lambda_{jk} V_{jk}\right)^{\mathbf{-1}}$$

2.5 What in practice does this mean in the n-instrument case? Suppose we wish to calibrate to m different variances s_j^2 $(j = 1, \dots, m)$ exhibited by instrument baskets described by vectors $\mathbf{w}_a$, where each $\mathbf{w}_a$ is a vector of n elements, the first element of which is the weight in the basket of the first instrument etc. For example, suppose we have implied volatilities for each instrument in isolation and for an equally weighted portfolio of the instruments. We would then have $m = n + 1$ calibrations, the first n of which involve weight vectors of the form $\mathbf{w}_a = (0, \dots, 0,1,0, \dots 0)$ (with the j'th element of the weight vector being 1, other terms being zero) and the last calibration having $\mathbf{w}_{n+1} = (1/n, 1/n, \dots, 1/n)$. If instead of calibrating to the implied volatility of an equally weighted basket we wished to calibrate to the implied volatility of a market cap weighted index implied volatility then $\mathbf{w}_{n+1}$ would be a vector of index weights.

2.6 The calibrated probability distribution will then have a covariance matrix as follows, where each $\mathbf{W}_a = \mathbf{w}_a \mathbf{w}_a^T$ is an $n \times n$ dimensional matrix:

$$\overline{\boldsymbol{\Omega}} = \left(\boldsymbol{\Omega}^{-1} + 2\sum_{a=1}^{m} \lambda_a \mathbf{W}_a\right)^{\mathbf{-1}}$$

subject to the m calibration equations $s_a^2 = \mathbf{w}_a^T \overline{\boldsymbol{\Omega}} \mathbf{w}_a$.

2.7 As long as this problem is not ill posed (e.g. because there are too many calibrations relative to the number of terms in the covariance matrix, or because there are no feasible solutions to the equations) calibration involves solving a set of m simultaneous equations $s_a^2 = \mathbf{w}_a^T \overline{\boldsymbol{\Omega}} \mathbf{w}_a$ in a unknowns, i.e. the λ_a.

2.8 An example of such a calibration is set out in the Appendix.

3. Other comments

3.1 If $m > 1$ then the simultaneous equations generated by the above procedure do not appear to have analytic solutions, see e.g. the Appendix. Instead they must in general be solved numerically by some iterative root search algorithm, e.g. one that finds the value of $\boldsymbol{\lambda} = (\lambda_1, \dots, \lambda_m)$ that satisfies the following equation:

$$\sum_{a=1}^{m} (s_a^2(\boldsymbol{\lambda}) - \hat{s}_a^2) = 0$$

where:

$$s_a^2(\boldsymbol{\lambda}) = \mathbf{w}_a^T \left(\boldsymbol{\Omega}^{-1} + 2 \sum_{a=1}^{m} \lambda_a \mathbf{W}_a \right)^{-1} \mathbf{w}_a$$

If the problem involves a relatively small number of instruments/calibration implieds then the problem can be solved using relatively straightforward tools such as the Excel Solver add-in. For larger instrument universes and/or calibration sets, more sophisticated root searching algorithms may be needed.

3.2 However, it may not always be appropriate to adopt a rote formulaic approach to calibration. For example, suppose we were trying to calibrate a UK equity risk model. Easily observable implied volatilities are available from listed equity derivatives on the FTSE100 index, the FTSE250 index but only for an incomplete range of individual equities. What should we do for securities for which there are no readily available implied volatilities?

3.3 The problem of incomplete or missing data is, of course, a generic problem with calibration, and not specific to the above problem. We could of course calibrate solely to those instrument volatilities that are easily observable. However, this would typically disproportionately affect the volatility assigned to the instruments included in the calibration set, see Appendix. Suppose that in fact general levels of implied volatilities are materially higher than those in the uncalibrated prior distribution, e.g. because there is an overall perception within the market that the "world is uncertain at the moment". Would we want calibration disproportionately to mark up volatilities of securities on which there were readily available option prices versus those on which there were not?

3.4 An alternative would perhaps be to introduce just two λ_a 's, i.e. two calibration equations. One might calibrate the volatility of the main market index to its current implied volatility (or actually in the spirit of above the variance to its current implied variance). The other might calibrate the average volatility (variance) of individual instruments to the average of their individual implied volatilities (variances) to the extent that these exist, with the same overall volatility (variance) adjustment then also applied to instruments where there is no observable implied volatility. Or perhaps we could adopt an intermediate approach of applying averaging within individual sectors rather than across the market as a whole. Of course, such averaging approaches would be less effective at calibrating individual instrument volatilities so that they exactly matched their own implied volatilities where these exist. That is the nature of averaging! We might also want to impose further constraints on the calibration to force retention of any parsimonious factor structure imposed on the prior distribution.

3.5 Analytical weighted Monte Carlo can also be used to calibrate risk models to cater for different future time periods, as long as there is a suitable term structure of implied volatilities available. We merely need to repeat the exercise separately for each term. Again, if necessary, 'missing' calibration data can be handled using averaging approaches as per 3.4.

3.6 It is worth noting that calibration of risk models to market implieds may make resulting ex ante tracking errors, VaR's and other similar risk measures more volatile (because implied volatilities are themselves volatile over even quite short periods of time). This may be important if explicit ex ante tracking error or VaR style risk limits are present in investment management agreements. Use of risk statistics calibrated to market implieds might then create greater likelihood of breach of such mandate restrictions merely because of market movements. We could dampen the impact of this 'volatility of volatility' by applying some sort of credibility weighting to current implied volatilities versus volatilities extrapolated from past history but of course only at the expense of calibration quality. Alternatively, it might be appropriate to quote more than one set of risk statistics, e.g. some 'longer term' ones (perhaps based solely on extrapolating past history using a relatively long time window) and some 'shorter-term' ones more fully calibrated to market implieds. The former might then be used more for mandate limit purposes and the latter more for day-to-day management of the portfolio.

References

Avellandeda M., Friedman C., Buff R, Granchamp N., Kruk L. and Newman J. (2001). Weighted Monte Carlo: A New Technique for Calibrating Asset-Pricing Models*. International Journal of Theoretical and Applied Finance* 4(1), pages 91 – 119.

Elices & Gimenez (2006). Weighted Monte Carlo. *Risk Magazine*. May 2006.

Kemp, M.H.D. (2005). Risk Management in a Fair Valuation World. British Actuarial Journal, 11, No 4, pp. 595-712

Kemp, M.H.D. (2009). Market consistency: Model calibration in imperfect markets. John Wiley & Sons [for futher information on this book please see MarketConsistency]

APPENDIX: Example calibration to market implied volatilities using the analytical weighted Monte Carlo approach

As at some date in the past a leading commercial vendor's risk system included the following standard deviations and correlations within its covariance matrix risk model for the following four UK equities:

Security	Standard Deviation (%pa)	Correlation with			
		1	0.29	0.33	0.30
AL	16.92	0.29	1	0.35	0.29
BLT	28.56	0.33	0.35	1	0.45
AVZ	36.64	0.30	0.29	0.45	1
BAY	32.85	1	0.29	0.33	0.30

The predicted tracking error of a model portfolio consisting of 35% Alliance and Leicester, 35% BHP Billiton, 15% Amvescap and 15% British Airways versus an equally weighted benchmark of these four stocks using this risk model was 5.35%pa.

As at approximately the same date the (at-the-money) implied volatilities for (call) options on these equities were approximately as follows:

Security	Implied Volatility (%pa)
AL	22
BLT	31
AVZ	30
BAY	27

Calibrating simultaneously to these four pieces of market information using the analytical weighted Monte Carlo approach gives a calibrated covariance matrix as shown below. The calibrated volatility of each individual security now matches its implied volatility. There are also some changes to individual correlations. The tracking error of the model portfolio using this calibrated covariance matrix is 4.77%pa.

Security	Standard Deviation (%pa)	Correlation with			
		AL	BLT	AVZ	BAY
AL	22	1	0.35	0.33	0.30
BLT	31	0.35	1	0.33	0.26
AVZ	30	0.33	0.33	1	0.36
BAY	27	0.30	0.26	0.36	1

If we were also to calibrate to an 'index' implied volatility, then the individual security implied volatilities would remain the same, since they are already calibrated to the market. What instead would happen is that correlations between securities would change. There is no listed option relating to a basket of these four stocks. If the above model fully matched market implieds then the 'index' implied volatility would be 19.3%pa. But suppose that there was an observable 'index' implied volatility and that it was 15%pa. The covariance matrix would then be as follows. The tracking error of the model portfolio would increase to 5.73%pa.

Security	**Standard Deviation (%pa)**	**Correlation with**			
		AL	BLT	AVZ	BAY
AL	22	1	0.13	0.09	0.06
BLT	31	0.13	1	0.03	-0.05
AVZ	30	0.09	0.03	1	0.12
BAY	27	0.06	-0.05	0.12	1

Conversely, if our 'index' basket actually had an implied volatility of 21%pa then the covariance matrix would become:

Security	**Standard Deviation (%pa)**	**Correlation with**			
		AL	BLT	AVZ	BAY
AL	22	1	0.46	0.45	0.41
BLT	31	0.46	1	0.46	0.40
AVZ	30	0.45	0.46	1	0.47
BAY	27	0.41	0.40	0.47	1

The average correlation between individual securities is quite sensitive to divergent movements between index implied volatility and average single security implied volatility, as is the tracking error of the model portfolio, which would now be 4.30%pa.

As we might also expect, each individual security calibration point disproportionately affects the volatility of that particular security. For example, suppose we only calibrated the original risk model to one implied volatility, namely the one for Amvescap. The calibrated volatilities would then be as follows:

	Volatility prior to calibration (%pa)	**Volatility post calibration (%pa)**	**Change**
AL	16.92	16.62	-2%
BLT	28.56	27.97	-2%
AVZ	36.64	30.00	-18%
BAY	32.85	31.75	-3%