
ISSN 1544-6123
Volume 6, Issue 3, September 2009
ELSEVIER
Finance Research Letters
Editor-in-Chief:
Ramazan Gençay, Simon Fraser University
Co-Editors:
Sugato Bhattacharyya, University of Michigan
Dirk Hackbarth, University of Illinois
Amir Yaron, University of Pennsylvania
Associate Editors:
Franklin Allen, University of Pennsylvania
Joshua D. Coval, Harvard University
Jean-Pierre Danthine, University of Lausanne
Espen Eckbo, Dartmouth College
Murray Frank, University of Minnesota
Kenneth A. Froot, Harvard University
Christian Gollier, University of Toulouse
John Graham, Duke University
Campbell Harvey, Duke University
Joel Hasbrouck, New York University
Christopher Hennessy, University of California, Berkeley
Robert A. Jarrow, Cornell University
Christopher S. Jones, University of Southern California
Francis Longstaff, University of California, Los Angeles
Richard K. Lyons, University of California, Berkeley
Enrico C. Perotti, University of Amsterdam
Gordon M. Phillips, University of Maryland
Michael R. Roberts, University of Pennsylvania
Stephen A. Ross, Massachusetts Institute of Technology
Matthew Spiegel, Yale University
Sheridan Titman, The University of Texas at Austin
Peter Tufano, Harvard University
Toni Whited, University of Rochester, New York
Available online at
ScienceDirect
www.sciencedirect.com

Finance Research Letters 6 (2009) 122–129

Contents lists available at ScienceDirect

Finance Research Letters

journal homepage: www.elsevier.com/locate/frl

Options on portfolios with higher-order moments

Rishabh Bhandari [a], Sanjiv R. Das [b,*,1]

[a] Indian School of Business, Gachibowli, Hyderabad 500032, India
[b] Santa Clara University, Finance (Leavey School of Business), 500 El Camino Real, Santa Clara, CA 95053, USA

ARTICLE INFO

Article history:
Received 17 January 2009
Accepted 8 April 2009
Available online 14 April 2009

JEL classification:
C14
C16
C63
G13

Keywords:
Co-moments
Tensors
Options
Portfolios

ABSTRACT

We develop a simple calibration approach to generate return distributions for multivariate asset distributions and use this technique to price options on portfolios given the first four co-moments of the joint distribution of returns. The technique is fast and captures the impact of covariance, and the co-skewness and co-kurtosis tensors on the value of these options. Given the technique works for a portfolio, the same is also applicable to options on individual securities as a special simpler case.

1. Introduction

Higher-order moments are important considerations in pricing securities and in valuing portfolios comprised of these securities. To quote Gene Fama[2]: "Many of the market tragedies that you see are the result of extreme events that people take to be unusual but that really aren't that unusual."

The Black and Scholes (1973) model assumes that the underlying security return is normally distributed. That this assumption is empirically rejected is now well established, and extensions to the model in the form of the class of jump models of Merton (1976) as well as stochastic volatility models

* Corresponding author.
E-mail addresses: rishabh_bhandari@isb.edu (R. Bhandari), srdas@scu.edu (S.R. Das).

[1] The author acknowledges financial support from the Dean Witter Foundation and a Breetwor Fellowship.

[2] In "Ideas that changed the theory and practice of investing, a conversation with Eugene F. Fama," Journal of Investment Consulting, Fall 2008, 6–14.

1544-6123/$ - see front matter
doi:10.1016/j.frl.2009.04.002

such as the one by Heston (1993) are widely used. Das and Sundaram (1999) showed that the first four conditional return moments have specific mappings into jump-diffusion and stochastic volatility models. That is, these models generate non-zero skewness and excess kurtosis (for single securities) with precise modulation of their parameters.

If options are written on *portfolios* then we need to model covariance, co-skewness and co-kurtosis, and obtain the conditional joint distribution of returns of the portfolio. These portfolio moments depend on the asset weights in the portfolio. The prices of options on this portfolio depend on the risk-neutral distribution of portfolio returns. We develop a simple method that takes the covariance and higher-dimensional matrices (tensors) of co-skewness and co-kurtosis as inputs, and produces prices of options on entire portfolios, correctly embedding the skewness and kurtosis of portfolio returns.

The increasing need for this technology is predicated on the growth of derivatives on baskets of securities. As securitization has increased, the underlying asset for many derivatives is often a portfolio and not a single security. Tranches of asset pools are derivatives that require slicing up the payoff profile of the underlying pool. And in many cases the underlying basket's return is far from being normally distributed, such as ABS pools, mortgage pools, credit baskets, options on indexes, etc. Our approach here is therefore especially apt for we price options on portfolios accounting correctly for higher return co-moments.

The model is computationally fast. It takes a few seconds to compute all the inputs and generate the prices of portfolio options. Given this, it is easy to calibrate the model to historical data or to prices of traded securities. And even though the model is aimed at options on portfolios, it may as well be used to price options on single assets accounting for higher-order moments.

There is a rich literature that has also attempted to model higher-order moments in the pricing of derivative securities. Starting from the classic models of Merton (1976) (jump-diffusions) and Heston (1993) (stochastic volatility), we also have implied tree models that calibrate exactly to the implied volatility surface (see Derman and Kani, 1994; Dupire, 1994; Rubinstein and Mark, 1994). It is possible to start from a Gaussian return distribution and then include additional terms in it to match higher moments. Jarrow and Rudd (1982) showed how an Edgeworth series technique might be used to extend a known probability density function to higher-order moments by including these as additional terms in the series. In this manner an unknown distribution whose moments are known may be approximated to a high degree of accuracy by adding as many terms to the series as desired. Rubinstein (1998) showed how to take the idea of Edgeworth series and incorporate the model on a tree so as to enable the pricing of American options with higher-order moments (see also Tian (1993) for an implementation with the same motivations). This technique, known as Edgeworth trees is simple and easy to implement. The Edgeworth series approach uses cumulants, and an alternate approach is to use Gram–Charlier series, where moments are used instead of cumulants (see Johnson et al. (1994) for a statistical overview, and Backus et al. (1997) for an application).

In this paper, we extend these ideas in two ways. First, we develop fat-tailed distributions for portfolio returns, not just that of individual securities. Second, rather than use a Gram–Charlier or Edgeworth expansion, we use an exponential-affine approach to transforming any "base" distribution into one that matches exactly the first l moments of the desired return distribution. In the applications below we work with the first four ($l = 4$) co-moments of the returns of all assets within a portfolio. Our work in this paper is also related to the work on modeling joint distributions using copula functions, as described in Sklar (1959), Sklar (1973), Clayton (1978), Nelsen (1999), Frees and Valdez (1998), and Frey and McNeil (2003). For applications of copulas with regards to credit portfolios see Embrechts et al. (2003), and Das and Geng (2004). A very general exposition of Hermite polynomial expansions for multivariate distributions related to high-dimensional stochastic processes is developed in Ait-Sahalia (2008).

In Section 2 we present the notation used in the paper, and the initial computations for co-skewness and co-kurtosis. Section 3 provides the main innovation of the paper – the technique for generating the conditional portfolio return distribution accounting correctly for all the moments. We provide illustrative examples of how the return distribution is modulated for skewness and kurtosis. Section 4 offers option pricing examples. The first example shows how skewness and kurtosis impact the value of put options. The second example uses real data and illustrates call option pricing. A third example shows how options are priced on a portfolio where prices are shown to respond to changing portfolio weights. Section 5 provides suggestions for further research and extensions of the ideas in this paper.

2. Notation

The portfolio comprises N assets. The portfolio weights are denoted $w = [w_1, w_2, \ldots, w_N]'$. We price an option $V(w)$ of maturity τ that is a function of the random return on the portfolio, i.e.

$$V(w) = e^{-r\tau} E[g(R(w))] \tag{1}$$

where $R(w)$ is the vector of returns in each state of the world, conditional on choosing a portfolio w. The payoff function of the derivative security is denoted $g(R)$. The risk free interest rate is r. The expectation $E(\cdot)$ above is taken across all states of the world with respect to the risk-neutral measure. The state space is generated by the joint returns of the N assets.

2.1. Portfolio inputs

The time series of returns on these assets are represented as r_{it} where $i = 1 \ldots N$ and $t = 1 \ldots T$, where T is the number of periods of returns in the data.

The inputs to the model comprise the vector of mean returns on these N assets, denoted $\mu = \{\mu_i\}_{i=1\ldots N} = [\mu_1, \ldots, \mu_N]' \in \mathscr{R}^N$. The covariance matrix of these assets is denoted $\Sigma = \{\sigma_{ij}\}_{i,j=1\ldots N} \in \mathscr{R}^{N\times N}$. Both these central moments are calculated in the usual way.

Likewise, we define the non-central co-skewness (S) and co-kurtosis (K) of returns as follows:

$$S = \{S_{ijk}\}_{i,j,k=1\ldots N} \in \mathscr{R}^{N\times N\times N} \tag{2}$$

$$K = \{K_{ijkl}\}_{i,j,k,l=1\ldots N} \in \mathscr{R}^{N\times N\times N\times N} \tag{3}$$

These tensors are easy to compute from the data. We note that

$$S_{ijk} = E[r_i \times r_j \times r_k] \tag{4}$$

$$K_{ijkl} = E[r_i \times r_j \times r_k \times r_l] \tag{5}$$

These comprise the raw moments from the data.

2.2. Portfolio moments

Given portfolio weights w, the mean (μ_p) and variance (σ_p^2) of the portfolio are obtained via the usual calculation:

$$m_1 = \mu_p(w) = w'\mu, \quad \sigma_p^2(w) = w'\Sigma w \tag{6}$$

The non-central second moment is $m_2 = \sigma_p^2 + m_1^2$. The non-central third and fourth moments of the portfolio are:

$$m_3 = \sum_{i=1}^{N}\sum_{j=1}^{N}\sum_{k=1}^{N} w_i w_j w_k S_{ijk} \tag{7}$$

$$m_4 = \sum_{i=1}^{N}\sum_{j=1}^{N}\sum_{k=1}^{N}\sum_{l=1}^{N} w_i w_j w_k w_l K_{ijkl} \tag{8}$$

The portfolio skewness (S_p) and excess kurtosis (K_p) are then given by the usual expressions:

$$S_p(w) = \frac{1}{\sigma_p(w)^3}[m_3 - 3m_2 m_1 + 2m_1^3] \tag{9}$$

$$K_p(w) = \frac{1}{\sigma_p(w)^4}[m_4 - 4m_3 m_1 + 6m_2 m_1^2 - 3m_1^4] - 3 \tag{10}$$

Note that all four moments are necessarily functions of portfolio weights $w \in \mathscr{R}^N$.

3. Portfolio distribution with higher moments

Given the moments $\{\mu_p(w), \sigma_p^2(w), S_p(w), K_p(w)\}$, conditional on weights w, we develop an algorithm to create the portfolio return distribution. The algorithm uses a "base" distribution $F(\cdot)$ for the generation of returns, and this distribution should be chosen such that its support covers an empirically acceptable range of returns. Distributions such as the normal and Student-T are easily specified to cover all return ranges given that their support lies on $(-\infty, +\infty)$. The Beta distribution may also be used with any chosen support. It allows for varied shapes. The steps in the algorithm are as follows:

(1) Generate a vector of values $u = [0, \Delta u, 2\Delta u, \ldots, 1] \in \mathscr{R}^m$, where m is sufficiently large so that Δu is sufficiently small.

(2) Transform the vector $u \in \mathscr{R}^m$ into a vector $x \in \mathscr{R}^m$ as follows: $x = F^{-1}(u)$. This assumes that the inverse function of the distribution $F(\cdot)$ exists. For the normal distribution this is easily available.

Important: we highlight the fact that the vector x comprises m equiprobable outcomes. This helps speed up calculations in the model.

(3) Transform the vector x into a portfolio return vector $R \in \mathscr{R}^m$ as follows:

$$R = a + bxe^{cx+dx^2} \in \mathscr{R}^m \tag{11}$$

where each outcome of R is equiprobable. The parameters $\{a, b, c, d\}$ are chosen such that the mean (μ_R), variance (σ_R^2), skewness (S_R) and kurtosis (K_R) of the vector R are equal to the portfolio moments, i.e., $\{\mu_R, \sigma_R^2, S_R, K_R\} = \{\mu_p(w), \sigma_p^2(w), S_p(w), K_p(w)\}$. This fit is undertaken numerically. For a detailed treatise on generating non-uniform random variates, see the excellent book by Devroye (1986).

Since all derivative pricing is undertaken using the risk-neutral probability measure, and the expected return of all assets under this measure is the risk free rate, we set the mean return to be equal to $r \times \tau$.

Whereas the first four moments of the portfolio return are functions of all parameters $\{a, b, c, d\}$, we note that the these parameters primarily modulate the mean, variance, skewness and kurtosis, respectively.

Our experiments reveal that this fitting exercise takes only a few seconds (even when undertaken on a spreadsheet). The portfolio return distribution depends on the choice of inversion distribution F. This allows the financial engineer flexibility in choosing a range of shapes of the portfolio return distribution while ensuring that the first four moments are preserved. An analogous situation arises in the use of copula functions for modeling credit portfolios.

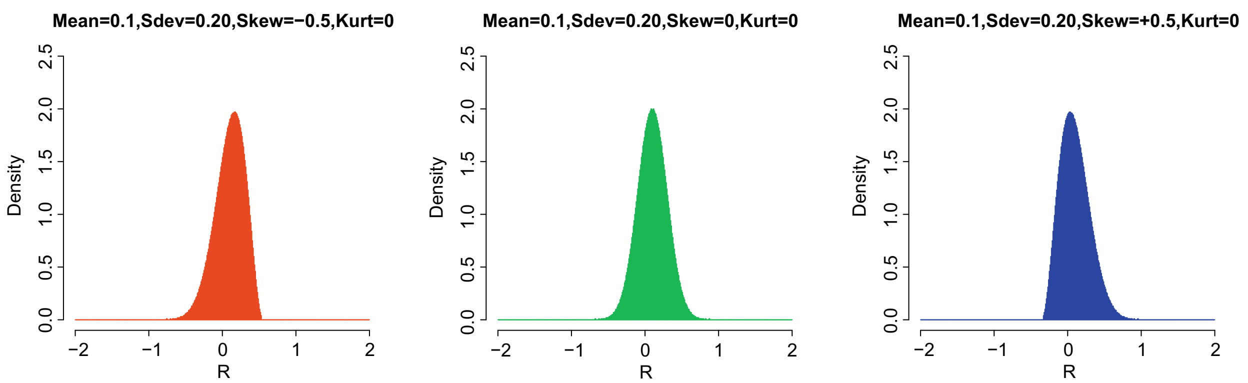

Fig. 1. Distribution shapes when skewness is varied. The mean and standard deviation of returns are chosen to be 10% and 20%, respectively. Skewness is chosen to take three values: $\{-0.5, 0, +0.5\}$ as shown in the three plots from left to right. Excess kurtosis is zero. The inversion distribution F is chosen to be Gaussian.

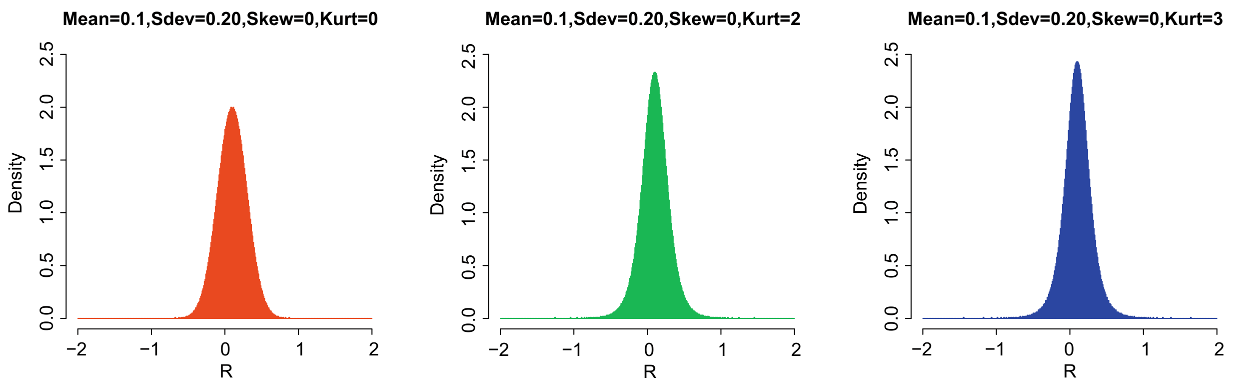

Fig. 2. Distribution shapes when kurtosis is varied. The mean and standard deviation of returns are chosen to be 10% and 20%, respectively. Excess kurtosis is chosen to take three values: $\{0, 2, 3\}$ as shown in the three plots from left to right. Skewness is zero. The inversion distribution F is chosen to be Gaussian.

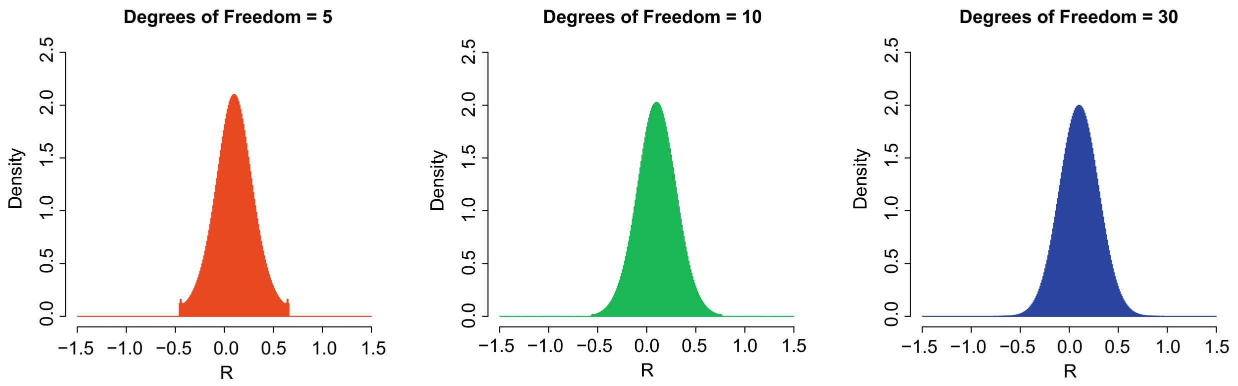

Fig. 3. Distribution shapes when degrees of freedom of the inversion distribution F is varied. The mean and standard deviation of returns are chosen to be 10% and 20%, respectively. Skewness and kurtosis are zero. The inversion distribution F is chosen to be Student-T with $\{10, 30\}$ degrees of freedom.

Figs. 1 and 2 show how the distribution shape changes when the method is applied to incorporate skewness and kurtosis. Fig. 1 shows different distributions when skewness is varied, keeping all other moments the same. Fig. 2 shows different distributions when kurtosis is varied, keeping all other moments the same.

Depending on the "base" inversion distribution $F(\cdot)$, we may obtain different shapes of the portfolio return distribution, keeping the first four moments the same. To illustrate this, see Fig. 3. The transformation causes truncation of the T-distribution at lower degrees of freedom to reduce the already high kurtosis.

4. Option pricing

4.1. A first example

To fix ideas, assume that we wish to price a put option on a portfolio with a current normalized price of \$1. The value of the portfolio at maturity τ will be given by $1 \times e^R$ where R is the continuous portfolio return over this time period and is generated using Eq. (11). Since our procedure generates an entire m-dimensional supporting vector of portfolio returns $R \in \mathscr{R}^m$, where each return is equiprobable, the value of the put option becomes

$$P = e^{-r\tau} \frac{1}{m} \sum_{j=1}^{m} \max[0, K - e^{R_j}] \tag{12}$$

where K is the strike price of the option. Illustrative pricing results are shown in Table 1.

Table 1
Put option prices with higher moments. The maturity of the options is 1 year. The risk free rate is 5%. Strikes are chosen ranging from 0.7 (OTM) to 1.3 (ITM) times the at-the-money-forward (ATMF) strike price. The ATMF strike is equal to the forward price of the portfolio. We varied skewness (Sk) to be $\{-0.5, 0, +0.5\}$ and excess kurtosis (Kt) to be $\{0, 3, 6\}$. The inversion distribution is Gaussian.

Skewness/kurtosis	Multiple of the at-the-money-forward strike						
	OTM			ATM		ITM	
	0.70	0.80	0.90	1.00	1.10	1.20	1.30
$Sk = -0.5; Kt = 0$	0.0079	0.0215	0.0465	0.0854	0.1392	0.2071	0.2871
$Sk = -0.5; Kt = 3$	0.0073	0.0188	0.0418	0.0807	0.1372	0.2090	0.2922
$Sk = -0.5; Kt = 6$	0.0070	0.0178	0.0401	0.0791	0.1364	0.2095	0.2935
$Sk = 0; Kt = 0$	0.0055	0.0182	0.0442	0.0861	0.1437	0.2147	0.2959
$Sk = 0; Kt = 3$	0.0059	0.0170	0.0406	0.0814	0.1400	0.2133	0.2968
$Sk = 0; Kt = 6$	0.0060	0.0157	0.0375	0.0775	0.1369	0.2120	0.2972
$Sk = +0.5; Kt = 0$	0.0021	0.0145	0.0430	0.0887	0.1493	0.2217	0.3029
$Sk = +0.5; Kt = 3$	0.0044	0.0152	0.0399	0.0828	0.1434	0.2175	0.3010
$Sk = +0.5; Kt = 6$	0.0048	0.0151	0.0387	0.0808	0.1413	0.2160	0.3003

Even with this simple example, we see interesting results. As skewness and kurtosis change, the prices of out-of-the-money (OTM) options experience the biggest percentage change. As expected OTM puts are more valuable as skewness turns negative. Interestingly, as kurtosis increases, OTM puts become more valuable when skewness is zero or positive, but OTM puts become less valuable if skewness is negative. For ITM puts, as kurtosis increases, put values increase for zero and negative skewness, but decrease if skewness is positive. Hence, kurtosis whittles away one tail to enhance the other and offsets skewness to some extent. When skewness is zero, increasing kurtosis diminishes the price of puts in a broad range around the ATM strike, but increases the price of deep OTM and ITM options.

4.2. A second example

We also considered an example using real-world data. We constructed an equally-weighted portfolio of five securities: the 10-year Treasury bond, the 90-day Treasury bill, the composite S&P index, and two stocks, Citigroup and Apple. We downloaded historical monthly returns on these five securities for the period January 2000–December 2007, and used the data to construct the covariance matrix, and the co-skewness and co-kurtosis tensors. From these, we computed the moments of the portfolio's monthly return. The mean return is 0.0090, the standard deviation of return is 0.0414, and the skewness and excess kurtosis are 1.8713 and 0.9830, respectively. We used these moments to price call options on the unit portfolio at varied strike prices. Results are shown in Table 2.

The entire pricing procedure takes only a few seconds. We see that the deep OTM calls are worthless, as the probability of crossing the strike of 1.20 in 1 month is negligible, given that this is a five-sigma event.

Table 2
Option prices using a real-world portfolio. Calls are priced on an equally-weighted portfolio of five securities: the 10-year Treasury bond, the 90-day Treasury bill, the composite S&P index, and two stocks, Citigroup and Apple. Data is the historical monthly returns on these five securities for the period January 2000–December 2007. Portfolio statistics are: the mean return is 0.0090, the standard deviation of monthly return is 0.0414, and the skewness and excess kurtosis are 1.8713 and 0.9830, respectively. The risk free rate per month is 0.004167 or 5% per year. The option maturity is 1 month. The inversion distribution is Gaussian.

	ITM			ATM		OTM	
Strike price	0.70	0.80	0.90	1.00	1.10	1.20	1.30
Call price	0.3038	0.2042	0.1046	0.0180	0.0011	0.0000	0.0000

4.3. A third example

We now proceed to analyze the impact of co-skewness and co-kurtosis on option prices. To do so, we use a model with only 2 assets in the portfolio, thereby keeping ideas simple. The assets are the 10-year Treasury bond and the 90-day Treasury bill. The annualized mean return vector (μ) and covariance matrix (Σ) of returns are computed and are as follows:

$$\mu = \begin{bmatrix} 0.06760987 \\ 0.03419100 \end{bmatrix}; \quad \Sigma = \begin{bmatrix} 0.058884101 & 0.0004640430 \\ 0.000464043 & 0.0003643663 \end{bmatrix} \tag{13}$$

We need to represent the non-central co-skewness matrix in two slices since it is of dimension $2 \times 2 \times 2$. These are as follows:

$$S_{ij1} = \begin{bmatrix} 0.004202153 & 0.0017095510 \\ 0.001709551 & 0.0001378430 \end{bmatrix}$$
$$S_{ij2} = \begin{bmatrix} 0.0017095510 & 1.378430e-04 \\ 0.0001378430 & 7.875703e-05 \end{bmatrix} \tag{14}$$

We represent the non-central co-kurtosis matrix in four slices since it is of dimension $2 \times 2 \times 2 \times 2$. These are as follows:

$$K_{ij11} = \begin{bmatrix} 0.0132277574 & 2.926168e-04 \\ 0.0002926168 & 6.810904e-05 \end{bmatrix}$$
$$K_{ij21} = \begin{bmatrix} 2.926168e-04 & 6.810904e-05 \\ 6.810904e-05 & 7.606382e-06 \end{bmatrix}$$
$$K_{ij12} = \begin{bmatrix} 2.926168e-04 & 6.810904e-05 \\ 6.810904e-05 & 7.606382e-06 \end{bmatrix}$$
$$K_{ij22} = \begin{bmatrix} 6.810904e-05 & 7.606382e-06 \\ 7.606382e-06 & 4.391194e-06 \end{bmatrix} \tag{15}$$

Taking these base levels of the moments, we price calls and puts on variously weighted portfolios of these two fixed-income instruments. The current portfolio value is normalized to \$1, and the options are priced for different ranges of strike prices, from 0.8 to 1.2, as is done in the preceding examples. The computed portfolio mean, standard deviation, skewness and excess kurtosis are reported in Table 3 which also shows the option pricing results. For calls, we see that at strikes greater than or equal to

Table 3

Option prices as co-moments change with changing portfolio weights. Portfolio weights are w_1 and w_2 in the 10-year Treasury bond and 90-day Treasury bill. Calls and puts are priced using data on the historical monthly returns on these two securities for the period January 2000–December 2007. The risk free rate per month is 0.004167 or 5% per year. The option maturity is 1 month. The inversion distribution is Gaussian.

w_1	w_2	Return Moments				Strikes				
		Mean	Sdev	Skew	Kurt	0.8	0.9	1.0	1.1	1.2
Calls										
0.1	0.9	0.0031	0.0026	1119.2569	0.9607	0.2242	0.1246	0.0251	0.0000	0.0000
0.3	0.7	0.0037	0.0063	108.6745	1.2490	0.2037	0.1042	0.0046	0.0000	0.0000
0.5	0.5	0.0042	0.0102	32.9040	1.1162	0.2006	0.1010	0.0014	0.0000	0.0000
0.7	0.3	0.0048	0.0142	15.3673	1.0375	0.2040	0.1044	0.0048	0.0000	0.0000
0.9	0.1	0.0054	0.0182	8.8656	0.9888	0.2049	0.1053	0.0071	0.0000	0.0000
Puts										
0.1	0.9	0.0031	0.0026	1119.2569	0.9607	0.0000	0.0000	0.0000	0.0745	0.1741
0.3	0.7	0.0037	0.0063	108.6745	1.2490	0.0000	0.0000	0.0000	0.0950	0.1946
0.5	0.5	0.0042	0.0102	32.9040	1.1162	0.0000	0.0000	0.0000	0.0982	0.1978
0.7	0.3	0.0048	0.0142	15.3673	1.0375	0.0000	0.0000	0.0000	0.0948	0.1943
0.9	0.1	0.0054	0.0182	8.8656	0.9888	0.0000	0.0000	0.0014	0.0939	0.1935

1.1, the options have no value since being in the money for those strikes is almost a nine-sigma event. The same is true for puts at strikes less than or equal to 0.9.

5. Concluding comments

The paper develops a simple approach to pricing options on portfolios accounting correctly for higher-order moments. There are several avenues for further research. The technique applies to European options and needs to be extended to pricing American type options. Since the approach develops the conditional distribution of returns, it can be implemented on a grid with smaller time intervals for the conditional distribution, making it possible to price American style options. Empirical examination of the model is a useful next step. How well does this simple approach fit the volatility surface for index options? Finally, the methodology is not restricted to option pricing. One may as well apply it to optimizing portfolios with higher moments, thereby extending the class of mean-variance models to one with mean-variance plus skewness and kurtosis.

References

Ait-Sahalia, Y., 2008. Closed-form likelihood expansions for multivariate diffusions. Annals of Statistics 36, 906–937.

Backus, D., Foresi, S., Li, K., Wu L., 1997. Accounting for Biases in Black–Scholes. Working Paper, Fordham University, CRIF Working Paper series. Paper 30.

Black, F., Scholes, M., 1973. The pricing of options and corporate liabilities. Journal of Political Economy 81, 637–654.

Clayton, D.G., 1978. A model for association in bivariate life tables and its application in epidemiological studies of familial tendency in chronic disease incidence. Biometrika 65, 141–151.

Das, S.R., Geng, G., 2004. Correlated default processes: a criterion-based copula approach. Journal of Investment Management 2, 44–70.

Das, S., Sundaram, R., 1999. Of smiles and smirks: a term structure perspective. Journal of Financial and Quantitative Analysis 34, 211–240.

Derman, E., Kani, I., 1994. Riding on a smile. Risk v7, 32–39.

Devroye, L., 1986. Non-Uniform Random Variate Generation. Springer-Verlag, New York.

Dupire, B., 1994. Pricing with a smile. Risk 7, 18–20.

Embrechts, P., Lindskog, F., McNeil, A., 2003. Modeling dependence with copulas and applications to risk management. In: Rachev, S.T. (Ed.), Handbook of Heavy Tailed Distributions in Finance. Elsevier, North-Holland, Amsterdam.

Frees, E.W., Valdez, E.A., 1998. Understanding relationships using copulas. North American Actuarial Journal 2, 1–25.

Frey, R., McNeil, A.J., 2003. Dependent defaults in models of portfolio credit risk. Journal of Risk 6, 59–92.

Heston, S., 1993. A closed-form solution for options with stochastic volatility with applications to bond and currency options. Review of Financial Studies 6, 327–343.

Jarrow, R., Rudd, A., 1982. Approximate option valuation for arbitrary stochastic processes. Journal of Financial Economics 10, 347–369.

Johnson, N., Kotz, S., Balakrishnan, N., 1994. Continuous Univariate Distributions. Wiley, New York.

Merton, R., 1976. Option pricing when underlying stock returns are discontinuous. Journal of Financial Economics 3, 125–144.

Nelsen, R.B., 1999. An Introduction to Copulas. Springer-Verlag, New York.

Rubinstein, Mark, 1994. Implied binomial trees. Journal of Finance 49, 393–440.

Rubinstein, M., 1998. Edgeworth binomial trees. Journal of Derivatives 5, 2027.

Sklar, A., 1959. Functions de repartition a n dimensions et leurs marges. Publications de l'Institut de Statistique de L'Université de Paris 8, 229–231.

Sklar, A., 1973. Random variables, joint distributions, and copulas. Kybernetica 9, 449–460.

Tian, Y., 1993. A modified lattice approach to option pricing. The Journal of Futures Markets 13, 563–578.