
Papers

Mean–variance versus full-scale optimisation: In and out of sample

Received: 9th May, 2006

Timothy Adler*

is vice president of research and development. He manages the development of Windham Portfolio Advisor™ investment technology and overseas research and development of IT-based tools and strategies to support our advisory and asset management service. Before his association with Windham, he was a member of the research group at State Street Associates, where he conducted portfolio construction and risk management analyses. He received Bachelor of Computer Systems Engineering and Bachelor of applied Computer Science degrees from the Royal Melbourne Institute of Technology and a Master of Business Administration degree from the Carroll Graduate School of Management at Boston College.

Mark Kritzman

is president and CEO of Windham Capital Management, LLC, where he is responsible for managing research activities and investment advisory services. He teaches a graduate course in financial engineering at MIT's Sloan School of Management. He serves on several corporate and non-profit boards, including the Institute for Quantitative Research in Finance, the International Securities Exchange, and State Street Associates. He is a member of several advisory and editorial boards, including *Emerging Markets Review*, *Financial Analysts Journal*, the International Association of Financial Engineers, the *Journal of Alternative Investments*, the *Journal of Derivatives*, and the *Journal of Investment Management*, where he is Book Review editor. He won Graham and Dodd awards in 1993 and 2002, the Research Prize from the Institute for Quantitative Investment Research in 1997, and the Bernstein-Fabozzi/Jacobs-Levy Award in 2003. In 2004, he was elected a Batten Fellow at the Darden Graduate School of Business Administration, University of Virginia. He received a Master of Business Administration degree from the Graduate School of Business at New York University.

*Windham Capital Management, LLC, 5 Revere Street, Cambridge, MA 02138, USA.
Tel: +1 617 234 9459; Fax: +1 617 576 7360; E-mail: tadler@windhamcapital.com

Abstract We present a recent innovation to portfolio construction called full-scale optimisation. In contrast to mean–variance analysis, which assumes that returns are normally distributed or that investors have quadratic utility, full-scale optimisation identifies the optimal portfolio given any set of return distributions and any description of investor preferences. It therefore yields the truly optimal portfolio in sample, whereas mean–variance analysis provides an approximation to the in-sample truth. Both approaches, however, suffer from estimation error. We employ a bootstrapping procedure to compare the estimation error of full-scale optimisation to the combined approximation and estimation error of mean–variance analysis. We find that, to a significant degree, the in-sample superiority of full-scale optimisation prevails out of sample.

Journal of Asset Management (2007) **7**, 302–311. doi:10.1057/palgrave.jam.2250042

Keywords: *full-scale, mean variance, optimization, estimation error, approximation error*

Part I: Introduction

For three decades, mean–variance analysis has served as the standard procedure for constructing portfolios. Recently, investors have experimented with a new optimisation procedure, called full-scale optimisation, to address certain limitations of mean–variance analysis. Specifically, mean–variance analysis

assumes that returns are normally distributed or that investor preferences are well approximated by mean and variance.¹ Full-scale optimisation relies on sophisticated search algorithms to identify the optimal portfolio given any set of return distributions and based on any description of investor preferences. Full-scale optimisation yields the

truly optimal portfolio in sample, whereas the mean–variance solution is an approximation to the in-sample truth.

Both approaches to portfolio formation, however, suffer from estimation error.² Mean–variance analysis requires investors to estimate the means and variances of all assets and the covariances of all asset pairs. To the extent the out-of-sample experience of these parameters departs from the in-sample parameter values, the mean–variance approximation will be even less accurate. Full-scale optimisation requires investors to estimate the entire multivariate return distribution. To the extent it varies from the in-sample distribution, full-scale optimisation will also yield sub-optimal results out of sample. We employ a bootstrapping procedure to compare the estimation error of full-scale optimisation to the combined approximation and estimation error of mean–variance analysis. We find that to a significant degree the in-sample superiority of full-scale optimisation prevails out of sample.

We organise the paper as follows. In Part II, we review mean–variance analysis and its limiting assumptions, and we describe full-scale optimisation. In Part III, we review our bootstrapping procedure for generating out-of-sample results. We present these results in Part IV, and we conclude the paper in Part V.

Part II: Alternative approaches to optimisation

Mean–variance analysis and its limitations

In his classic paper, ‘Portfolio Selection’ (1952), Markowitz submitted that investors should not choose portfolios that maximise expected return, because this criterion by itself ignores the principle of diversification. He proposed that investors should instead consider variances of returns, along with expected returns, and choose portfolios that offer the highest expected return for a given

level of variance. He called this rule the E–V maxim.

Markowitz demonstrated that, for given levels of risk, we can identify particular combinations of securities that maximise expected return. He deemed these portfolios ‘efficient’ and referred to a continuum of such portfolios in dimensions of expected return and standard deviation as the efficient frontier. According to Markowitz’s E–V maxim, investors should choose portfolios located along the efficient frontier.

This approach to portfolio formation is sufficient for maximising expected utility if portfolio returns are normally distributed or if investors have quadratic utility, which is defined as $E(U) = \mu - \lambda \sigma^2$, where μ equals portfolio expected return, λ equals risk aversion, and σ^2 equals portfolio variance. If returns are normally distributed, investors can infer the entire distribution of returns from its mean and variance; hence the irrelevance of specific periodic returns or higher moments. And even if returns are not normally distributed, quadratic utility assumes that investors are indifferent to other features of the distribution.

Many assets display return distributions that are approximately normal; however, no asset produces a perfectly normal distribution. Moreover, quadratic utility is not a realistic description of any investor’s attitude toward risk for several reasons. It assumes investors are as averse to upside

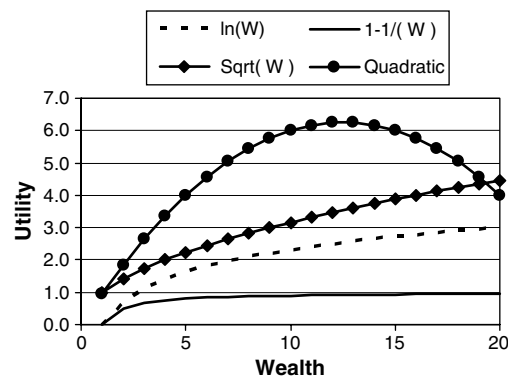


Figure 1 Utility functions

deviations as they are to downside deviations. Also at certain wealth levels it assumes investors prefer less wealth to more wealth. And finally quadratic utility assumes investors have increasing absolute risk aversion.

Financial economists usually assume that investors have power utility functions, which define utility as $1/\gamma \times \text{Wealth}^\gamma$. A log wealth utility function is a special case of power utility. As γ approaches 0, utility approaches the natural logarithm of wealth. A γ equal to $1/2$; implies less risk aversion than log wealth, while a γ equal to -1 implies greater risk aversion.³ These utility functions, along with a quadratic utility function, are shown in Figure 1.

Unlike quadratic utility, power utility functions assume a preference for upside deviations, they never slope downward, which would reflect a preference to reduce wealth, and they assume constant relative risk aversion. They are therefore a more plausible description of investor preferences than quadratic utility. Nonetheless, quadratic utility serves as an excellent approximation to most variations of power utility for a wide range of reasonable investment outcomes.⁴ Many investors, however, view power utility functions as overly simplistic. They believe that investor preferences are better described by more complex functions such as kinked utility functions or S-shaped value functions.

A kinked utility function changes abruptly at a particular wealth or return level and is relevant for investors who are concerned with breaching a threshold. Consider, for example, a situation in which an investor

requires a minimum level of wealth to maintain a certain standard of living. The investor's lifestyle might change drastically if she penetrates this threshold. Or she may be faced with a situation in which she will become insolvent if her wealth breaches some threshold. Or a particular decline in wealth may breach a covenant on a loan or a regulatory requirement. In these and similar situations, a kinked utility function as described below is more likely to describe one's attitude toward risk. Utility is defined by a log-wealth function above the threshold return and by a steeper function below the threshold return.

$$U(x) = \begin{cases} \ln(1+x), & \text{for } x \geq \theta \\ 10 \times (x - \theta) + \ln(1 + \theta), & \text{for } x < \theta \end{cases}$$

The symbol x represents portfolio return, and θ represents the return threshold.

Figure 2 shows a kinked utility function with the kink located at -5 per cent.

Proponents of behavioural finance also believe that utility is more complex than depicted by quadratic or power utility functions. In particular, Kahnemann and Tversky (1979) have found that people focus on returns more than wealth levels and that they are risk averse in the domain of gains but risk seeking in the domain of losses. For example, if a typical investor is confronted with a choice between a certain gain and an




Figure 2 Kinked utility with $\theta = -5$ per cent




Figure 3 S-shaped value function

uncertain outcome with a higher expected value, he will choose the certain gain. In contrast, when confronted with a choice between a certain loss and an uncertain outcome with a lower expected value, he will choose the uncertain outcome. This behaviour is captured by an S-shaped value function, which Kahnemann and Tversky modelled as follows.

$$U(x) = \begin{cases} -A(\theta - x)^{\gamma_1}, & \text{for } x \leq \theta \\ +B(x - \theta)^{\gamma_2}, & \text{for } x > \theta \end{cases}$$

Subject to:

$$A, B > 0$$

$$0 < \gamma^1, \gamma^2 \leq 1$$

The portfolio's return is represented by x , and A and B are parameters that together control the degree of loss aversion and the curvature of the function for outcomes above and below the return threshold, θ . Figure 3 shows an S-shaped value function with a threshold of 0 per cent.

Because returns are not perfectly normally distributed and because investor preferences do not conform precisely to quadratic utility, mean–variance analysis always yields a solution that is an approximation to the true in-sample utility-maximising portfolio. It therefore suffers from approximation error. In many cases this error is negligible; in other cases it is overwhelming. Beyond approximation error, mean–variance analysis also suffers from estimation error because the in-sample means, variances, and covariances will not prevail precisely out of sample.

Full-scale optimisation

Computational efficiency now allows us to perform full-scale optimisation as an alternative to mean–variance analysis. With this approach we calculate a portfolio's utility for every period in our sample, considering as many asset mixes as necessary to identify the weights that yield the highest expected utility, given any description of utility.

Table 1 Full-scale optimization with kinked utility

Year	Fund A returns (%)	Fund B returns (%)	Fund A weight (%)	Fund B weight (%)	Portfolio return (%)	Portfolio utility (%)
1	10.06	16.16	42.86	57.14	13.55	$\ln(1+0.1355)^{1/10}=1.2703$
2	1.32	-7.10	42.86	57.14	-3.49	$\ln(1-0.0349)^{1/10}=-2.5913$
3	37.53	29.95	42.86	57.14	33.20	$\ln(1+0.3320)^{1/10}=2.8668$
4	22.93	0.14	42.86	57.14	9.91	$\ln(1+0.0991)^{1/10}=0.9448$
5	33.34	14.52	42.86	57.14	22.59	$\ln(1+0.2259)^{1/10}=2.0365$
6	28.60	11.76	42.86	57.14	18.98	$\ln(1+0.1898)^{1/10}=1.7377$
7	5.00	-7.64	42.86	57.14	-2.22	$\ln(1-0.0222)^{1/10}=-1.3224$
8	-9.09	16.14	42.86	57.14	5.33	$\ln(1+0.0533)^{1/10}=0.5188$
9	-0.94	-7.26	42.86	57.14	-4.55	$\ln(1-0.0455)^{1/10}=-3.6514$
10	-22.10	14.83	42.86	57.14	-1.00	$\ln(1-0.0100)^{1/10}=-0.1005$
Average utility						1.7093

Suppose, for example, we wish to find the optimal blend between two funds whose returns are displayed in Table 1, assuming the investor has kinked utility. We compute utility for each period as $\ln(1+x)$ if the weighted portfolio return is greater than θ , and as $[10(x-\theta) + \ln(1+\theta)]$ if the weighted portfolio return is less than θ , where x equals the weighted portfolio return. We then shift the fund's weights using a numerical search procedure until we find the combination that maximises expected utility, which for this example equals a 42.86 per cent allocation to fund A and a 57.14 per cent allocation to fund B. The expected utility of the portfolio reaches a maximum at 1.7093 per cent. This approach implicitly takes into account all of the features of the empirical sample, including skewness, kurtosis, and any other peculiarities of the distribution.

We can apply full-scale optimisation to empirical distributions, theoretical distributions, or combinations based on empirical returns and theoretical assumptions. If we assume a theoretical distribution, we simply discretise it by randomly drawing returns from it and then applying the full-scale algorithm to these discrete returns. If we prefer to preserve the shape of an empirical sample but modify the assets' means to conform with our views about them prospectively, we simply adjust each return in the empirical sample by the difference between our view and the empirical mean.

Assuming our search algorithm is sufficiently effective, full-scale optimisation will yield the true in-sample utility-maximising portfolio. Unlike mean-variance analysis, it has no approximation error. But like mean-variance analysis it too suffers from estimation error. To the extent any of the features of the in-sample distribution do not prevail out of sample, the full-scale solution will be sub-optimal. In the next section, we describe the methodology we use to compare the combination of approximation and estimation error of

mean-variance analysis to the estimation error of full-scale optimisation.

Part III: Methodology

We base our analysis on a sample of monthly hedge fund returns covering a 10-year period from January 1994 to December 2003. We use this sample because hedge funds tend to display significantly non-normal higher moments. For evidence of non-normality in hedge funds returns, see Alexiev (2004), Davies *et al.* (2003), Fung and Hsieh (2000), Gregoriou and Gueyie (2003), Kat and Lu (2002), Lo (2001), Lo (2005), and McFall (2003).

We conduct two out-of-sample tests. In the first test, we use the monthly returns from January 1994 to December 1998 to solve for the in-sample utility-maximising portfolio based on full-scale optimisation, and we also solve for the mean-variance efficient portfolio with the same expected return.

Although we preserve the exact shape of the funds' in-sample distributions, we scale each return so that the average returns of the distributions conform to the implied returns of an equally weighted portfolio of hedge funds. By scaling the returns, we mitigate the effect of survivorship bias and selection bias on the means, and we mitigate the likelihood that the optimisations will yield corner solutions. This scaling does not affect the comparisons we wish to make, because both optimisation methods use the same scaled data and because our focus is on the shape of the return distributions and not the location of their means. We do not scale the out-of-sample returns.

We consider four utility functions: kinked with the kink set at -1 per cent; kinked with the kink set at -5 per cent; S-shaped with the inflection point set at 0 per cent; and S-shaped with the inflection point set at 0.5 per cent⁵. We do not consider power utility functions because the approximation errors of mean-variance analysis are arbitrarily small for these utility functions.⁶

We solve for the in-sample utility-maximising portfolio using our full-scale optimisation algorithm. Then we solve for the mean–variance portfolio on the efficient frontier that has the same expected return as the full-scale optimal portfolio. We record the asset weights of these portfolios, as well as their expected utility, expected return, standard deviation, kurtosis, and skewness.

We next bootstrap one-month vectors of cross-sectional returns with replacement 60 times from the period beginning January 1999 and ending December 2003 to generate a new five-year sample of returns. This sample does not overlap with the sample used to determine the portfolio weights, nor are the returns in this sample scaled. We repeat this procedure until we have 1,000 new five-year samples. We then apply the weights of the full-scale and mean–variance portfolios generated from the January

1994 to December 1998 sample to the 1,000 bootstrapped samples generated from the subsequent five-year period, and we compute portfolio metrics for these 1,000 out-of-sample histories.

We perform one additional test. We use the second five-year period (January 1999 to December 2003) to determine the full-scale and mean–variance weights, and we bootstrap from the first five-year period (January 1994 to December 1998) to generate the 1,000 out-of-sample histories. Again, we apply the weights from the in-sample full-scale and mean–variance optimisations to these new return samples and compute portfolio metrics for these out-of-sample histories.

These results allow us to compare how well the full-scale and mean–variance weights derived in sample perform out of sample. In the case of the full-scale optimal portfolios, the differences from

Table 2 In-sample results

In-sample periods	1994–1998		1999–2003	
	F-S	M-V	F-S	M-V
<i>Kinked at –1%</i>				
Expected return	5.85%	5.85%	8.34%	8.34%
Standard deviation	4.25%	2.79%	6.00%	5.29%
Skewness	0.43	–2.29	1.31	–0.05
Kurtosis	2.32	12.15	5.01	3.22
Expected utility	0.47%	0.19%	0.66%	0.16%
<i>Kinked at –5%</i>				
Expected return	6.78%	6.78%	9.63%	9.63%
Standard deviation	10.44%	9.06%	14.86%	14.60%
Skewness	0.01	–1.76	0.66	0.10
Kurtosis	2.49	9.62	3.45	3.18
Expected utility	0.52%	–0.09%	0.71%	0.12%
<i>S-shaped at 0%</i>				
Expected return	5.79%	5.79%	7.74%	7.74%
Standard deviation	2.72%	2.44%	1.84%	1.58%
Skewness	–2.74	–2.22	0.60	0.03
Kurtosis	16.09	11.30	3.04	2.82
Expected utility	113.85%	94.68%	126.48%	114.23%
<i>S-shaped at 0.5%</i>				
Expected return	7.01%	7.01%	8.85%	8.85%
Standard deviation	11.06%	10.64%	10.33%	8.94%
Skewness	–3.16	–1.75	–0.47	–0.03
Kurtosis	19.28	9.45	3.84	3.26
Expected utility	65.84%	23.44%	71.87%	14.74%

the in-sample results arise purely from estimation error. In the case of the mean–variance portfolios, the differences arise from a combination of approximation and estimation error.

Part IV: Results

Table 2 shows results for the in-sample optimal portfolios identified by full-scale optimisation and the mean–variance portfolios on the efficient frontier with the same expected returns.

The results in Table 2 demonstrate that full-scale optimisation yields significantly higher in-sample utility across all samples for investors who care about thresholds or who have different preferences with respect to gains and losses. These differences arise entirely from the approximation error of mean–variance analysis, because there is no estimation error in sample.

The first panel shows the results assuming a kinked utility function with the kink located at -1 per cent. It reveals that the full-scale portfolios generate higher expected utility than mean–variance portfolios.

Moreover, full-scale optimisation identifies portfolios with positive skewness, whereas their mean–variance counterparts display negative skewness, a characteristic that is undesirable to investors with kinked utility. The same general pattern prevails in the second panel, which assumes the kink is located at -5 per cent.

The third panel assumes an S-shaped function with the inflection point located at 0 per cent. Again, the full-scale portfolios perform better than the mean–variance portfolios as evidenced by their higher expected utility. Also, with one exception, full-scale optimisation produces more negative skewness than mean–variance analysis, an attribute that appeals to S-shaped investors. These investors are not attracted to extremely bad outcomes; they just do not especially mind them. They are, however, strongly attracted to the high density of

moderately good outcomes, which serves as the offset to the few extremely bad outcomes in a negatively skewed distribution. Full-scale optimisation does an excellent job of accommodating these preferences, while mean–variance analysis ignores them. This pattern prevails in the fourth panel, which assumes S-shaped utility with the inflection point located at 0.5 per cent.

It is important to note that in all cases the mean–variance portfolios have lower standard deviations than the full-scale portfolios. One might be tempted to conclude that this lower volatility is an advantage of mean–variance analysis, but such a conclusion would be wrong. The higher standard deviations of the full-scale portfolios might result from upside deviations, given the asymmetry of the return distributions. Moreover, even if lower volatility is an advantage, the utility functions may be relatively more sensitive to other features of the distributions, which mean–variance analysis ignores.

Next, we show we show how the full-scale and mean–variance weights derived in sample perform in the bootstrapped histories generated from the out-of-sample data.⁷

Table 3 reveals that the full-scale weights applied out of sample generate higher average utility than the mean–variance weights for every utility function across all samples, although as should be expected, the advantage of full-scale optimisation diminishes out of sample. It also reveals that the full-scale and mean–variance weights, on balance, generate differences in higher moments that are consistent with the differences of the in-sample results.

Table 4 provides additional support for the out-of-sample superiority of full-scale optimisation. It reveals that the full-scale weights generate higher realised utility than the mean–variance weights in a majority of the bootstrapped histories for all cases.

The frequency with which full-scale optimisation outperforms mean–variance analysis may seem relatively high given that

Table 3 Out-of-sample results

Out-of-sample periods	1999–2003		1994–1998	
	F-S	M-V	F-S	M-V
<i>Kinked at –1%</i>				
Expected return	14.45%	13.07%	12.13%	15.50%
Standard deviation	4.78%	3.45%	5.27%	6.51%
Skewness	0.33	–0.46	–0.64	–0.95
Kurtosis	4.02	4.36	4.34	6.09
Average utility	0.87%	0.82%	0.21%	0.17%
<i>Kinked at –5%</i>				
Expected return	14.45%	13.09%	18.13%	18.02%
Standard deviation	10.70%	8.00%	11.92%	13.62%
Skewness	0.72	–0.11	–1.05	–1.03
Kurtosis	3.62	4.16	6.97	6.81
Average utility	1.00%	0.96%	0.19%	–0.06%
<i>S-shaped at 0%</i>				
Expected return	14.33%	12.99%	12.64%	13.30%
Standard deviation	3.89%	3.35%	3.07%	4.07%
Skewness	–0.62	–0.50	–0.60	–0.61
Kurtosis	4.38	4.12	3.85	3.11
Average utility	119.59%	114.98%	115.28%	112.65%
<i>S-shaped at 0.5%</i>				
Expected return	15.67%	13.84%	17.39%	16.95%
Standard deviation	8.65%	9.41%	10.57%	9.27%
Skewness	–0.12	0.00	–1.24	–1.13
Kurtosis	4.26	3.90	7.00	7.27
Average utility	48.47%	35.72%	61.39%	58.75%

Table 4 Frequency (%) of F-S utility>M-V utility

	Out-of-sample periods	
	1998–2003	1994–1998
Kinked at –1%	61	67
Kinked at –5%	56	85
S-shaped at 0%	78	93
S-shaped at 0.5%	91	98

the average outperformance of full-scale optimisation is only marginal in some cases. This combination of results suggests that although mean-variance analysis usually does worse than full-scale optimisation, it occasionally does much better, which raises its average performance.

Part V: Summary and conclusion

Mean-variance analysis suffers from two sources of error: approximation error and estimation error. Because returns are not

precisely normally distributed and investors do not have quadratic utility, mean-variance solutions only approximate the true utility-maximising portfolio. Moreover, the estimated means, variances, and covariances are unlikely to match the realised means, variances, and covariances out of sample.

Full-scale optimisation serves as an alternative to mean-variance analysis. It relies on sophisticated search algorithms to identify the in-sample utility-maximising portfolio given any empirical or theoretical return distribution and any description of investor utility. Therefore, it is not subject to approximation error as long as the search algorithm is sufficiently effective. Yet it does suffer from estimation error, because the in-sample distribution of returns will not prevail precisely out of sample.

We bootstrap returns from out-of-sample periods to generate thousands of alternative histories. We then apply the in-sample

weights of the full-scale and mean–variance portfolios to the bootstrapped samples to evaluate their out-of-sample robustness. Our analysis shows that full-scale optimisation offers a better in-sample solution than mean–variance analysis based on a representative set of hedge funds for investors with kinked or S-shaped preferences. We also show that this result prevails out of sample on average and most of the time, although as should be expected, with some attenuation.

We feel obliged to add two cautionary comments. First, we chose a sample of returns that is significantly non-normal, and we assumed a set of utility functions that, while perhaps realistic, are unconventional. If instead we used returns that were more normally distributed or based our analysis on variations of power utility, we would unlikely observe significant in-sample approximation error from mean–variance analysis.

Second, our out-of-sample results are specific to the sample we used for comparing full-scale optimisation to mean–variance analysis. If we were to use other samples for which the higher moments were not as persistent, full-scale optimisation might not dominate mean–variance analysis out of sample. It is therefore important to understand qualitatively why we should expect a particular distribution from an asset or fund. In the case of hedge funds, there are compelling reasons to expect non-normal higher moments. Consider merger arbitrage funds, for example. These funds attempt to profit by acquiring the stock of companies that are takeover targets and by selling the stock of the acquiring companies. The return of such a strategy depends on a binary event. Either the merger occurs or it does not. If it does, the arbitrage generates an unusually high return, and if it does not, the arbitrage produces an unusually low return. It makes sense, therefore, to expect a higher concentration of returns in the tails of the distribution than near the centre of the distribution from this strategy.

In conclusion, we suggest that if you optimise among assets that are likely to have persistent non-normal higher moments and you care about thresholds or view gains and losses differently, you should consider full-scale optimisation as an alternative to mean–variance analysis.

Acknowledgement

We thank the participants of QWAFEFW (April 2006) for helpful comments.

Notes

1. Mean–variance analysis does not strictly require normality. It is suitable for a broader class of distributions called elliptical distributions. See, for example, Embrechts *et al.* (2002). However, normality does not always justify mean–variance analysis. Even though all information about a normal distribution is captured by the first two moments, this information is not linear in variance. Kurtosis, for example, is proportional to variance squared; hence mean–variance analysis may fail to maximise utility for functions that are relatively more sensitive to kurtosis than to variance. We thank Jarrod Wilcox for pointing this out to us.
2. There is a large literature on estimation error in portfolio construction. See, for example, Bawa *et al.* (1979), Jobson *et al.* (1979), Jobson and Korkie (1980), and Jorion (1992).
3. When γ equals -1 , utility is expressed as $1-W^{-1}$.
4. See, for example, Levy and Markowitz (1979).
5. S-shaped utility parameters A , B , γ_1 , γ_2 are the same as illustrated in Figure 3.
6. See, for example, Cremers *et al.* (2005).
7. Note that some of the out-of-sample results have higher utility than their corresponding in-sample results. This outcome arises because we scaled the in-sample returns to produce better behaved means, but we did not scale the out-of-sample returns.

References

- Alexiev, J. (2004) ‘The Impact of Higher Moments on Hedge Fund Risk Exposure’, *Journal of Alternative Investments*, Spring 2005, 50–65.
- Bawa, V., Brown, S. and Klein, R. (1979) *Estimation Risk and Optimal Portfolio Choice*, Amsterdam, North Holland Publishing Co.
- Cremers, J.-H., Kritzman, M. and Page, S. (2005) ‘Optimal Hedge Fund Allocations: Do Higher Moments matter?’ *The Journal of Portfolio Management*, 31(3), 70–81.
- Davies, R. J., Kat, H. M. and Lu, S. (2003) ‘Higher Moment Portfolio Analysis with Hedge Funds,’ Unpublished working paper.
- Embrechts, P. A., McNeil, J. and Straumann, D. (2002) ‘Correlation and Dependence in Risk Management: Properties and Pitfalls’, In M. Dempster (ed.), *Risk*

- Management: Value at Risk and Beyond*, Cambridge, Cambridge University Press.
- Fung, W. and Hsieh, D. (2000) 'Performance Characteristics of Hedge Funds and Commodity Funds: Natural versus Spurious Biases', *Journal of Financial and Quantitative Analysis*, 35, 291–307.
- Gregoriou, G. and Gueyie, J. -P. (2003) 'Risk-Adjusted Performance of Funds of Hedge Funds Using a Modified Sharpe Ratio', *Journal of Wealth Management*, 6(3), 77–84.
- Jobson, J. D. and Korkie, B. (1980) 'Estimation of Markowitz Efficient Portfolios', *Journal of the American Statistical Association*, 75, 544–554.
- Jobson, J. D., Korkie, B. and Ratti, V. (1979) 'Improved Estimation for Markowitz Portfolios Using James–Stein Type Estimators', *Proceedings of the American Statistical Association, Business and Economic Statistics Section*, Washington D.C. August 13–16, 279–284.
- Jorion, P. (1992) 'Portfolio Optimization in Practice', *Financial Analysts Journal*, 48(1), 68–74.
- Kahnemann, D. and Tversky, A. (1979) 'Prospect Theory: An Analysis of Decision Under Risk', *Econometrica*, 47, 263–290.
- Kat, H. M. and Lu, S. (2002) 'An Excursion into the Statistical Properties of Hedge Fund Returns', *Alternative Investment Research Centre Working Paper Series*, Working Paper #0016 (May 1, 2002).
- Levy, H. and Markowitz, H. M. (1979) 'Approximating Expected Utility by a Function of Mean and Variance', *American Economic Review*, 69(3), 1.
- Lo, A. W. (2001) 'Risk Management for Hedge Funds: Introduction and Overview', *Financial Analysts Journal*, 57, 16–33.
- Lo, A. W. (2005) *The Dynamics of the Hedge Fund Industry*, The Research Foundation of the CFA Institute, Amazon, August.
- McFall Lamm, R. Jr. (2003) 'Asymmetric Returns and Optimal Hedge Fund Portfolios', *Journal of Alternative Investments*, 6(2), Fall, 9–21.