The Revere Street Group Working Paper Series
Financial Economics
No. 272-1

Optimal Portfolios in Good and Bad Times

February 10, 1999

DRAFT

George Chow
Windham Capital Management Boston
806 Nantucket Drive
Redwood City, CA 94065
650 592-7360
GCHOW@FMG-STATESTREET.COM

Eric Jacquier
Carroll School of Management, Boston College
Fulton Hall
140 Commonwealth Avenue
Chestnut Hill, MA 02467-3808
617 552-2943
JACQUIER@JACQUIER.BC.EDU

Mark Kritzman
Windham Capital Management Boston
5 Revere Street
Cambridge, MA 02138
617 576-7360
WCMBOS@WORLD.STD.COM

Kenneth Lowry
State Street Bank and Trust Company
225 Franklin Street
Boston, MA 02110
617 664-3807
CLOWRY@FMG-STATESTREET.COM

Optimal Portfolios in Good and Bad Times[1]

Introduction

Harry Markowitz [1952] introduced an efficient process for selecting portfolios. His landmark innovation, known as mean-variance optimization, requires financial analysts to estimate expected returns, standard deviations, and correlations. With this information, Markowitz showed how to combine assets optimally so that for a particular level of expected return the resulting portfolio would offer the lowest possible level of expected risk, usually measured as standard deviation or variance. A continuum of these portfolios plotted in dimensions of expected return and standard deviation is called the efficient frontier. In 1990 Markowitz was awarded the Nobel Prize in economics for his pioneering work in portfolio selection.

Perhaps the greatest challenge to the implementation of portfolio selection is estimation of the requisite parameters. Financial analysts typically estimate expected returns by combining historical precedent with their views about current conditions. They are more agnostic about variances and correlations, however, and rely almost exclusively on equally weighted historical data from what they perceive to be a representative sample. Recent experience with failed hedge funds has focused attention on a serious limitation of this approach. Although this approach may provide reasonable risk estimates for the full investment horizon, it probably understates variances

[1] We thank Harry Markowitz for his helpful comments and, in particular, his suggestion to adapt Chow's [1995] multivariate objective function to blend covariance matrices. We also thank an anonymous referee, Stephen Brown, Kenneth Froot, and Alan Marcus for their comments as well as Edward Ladd, Jay Light, Jack Meyer, Michael Pradko, David Salem, Larry Siegel, David Swensen, and Richard Zeckhauser for their comments with respect to the application of this methodology to the Harvard Endowment Fund.

and correlations during periods of stress or financial crisis. During these periods, asset returns tend to become more volatile and more highly correlated. Thus, the diversification that characterizes the sample, on average, disappears when it is most needed.

For example, Siconolfi, Raghavan, and Pacelle [1998] wrote in the Wall Street Journal:

> *In mid-August, Russia abruptly defaulted on part of its debt and let the ruble fall, triggering a flight by investors from all types of risk into safe investments. That devastated some of LTCM's bets, leading to the huge losses of Aug. 21.*

Another way to think about this issue is to distinguish time measured observations from event measured observations. Consider, for example, a sample of monthly returns. In some months, it is quite likely that there were no significant events to cause prices to change, while in other months there may have been several important events. The typical estimation of variances and correlations assigns as much importance to the months with no significant events as it does to the event filled months. As an alternative, it may be informative to calibrate returns as a function of events and then to estimate risk parameters from these event-measured observations. This approach may provide a better representation of a portfolio's likely performance during stressful times.

We introduce two innovations to portfolio optimization in order to address the issue that risk parameters are unstable. First, we present a procedure for estimating risk parameters from multivariate outliers based on the rationale that outliers are more likely to be associated with stress related events than with noise. Keep in mind, however, that we define stress as periods that are unusual and not necessarily as periods characterized only by low or negative returns.

Second, we show how to construct portfolios that simultaneously balance risk parameters estimated from “quiet” or inside observations with those estimated from outlying observations. This latter innovation, which was suggested to us by Harry Markowitz, is an adaptation of Chow’s [1995] optimization algorithm for combining absolute and relative performance. In addition, we provide empirical results to demonstrate both procedures.

Multivariate Outliers

An outlier, given a return series for a single asset, is straightforward to identify. It is simply a return that falls outside a chosen confidence interval around the expected return. For example, we may decide to define an outlier as a return that falls within the tails that comprise 25% of the distribution, 12.5% on either side. Thus, if continuous returns are normally distributed, an outlier for a single return series is any continuous return that is greater than m+1.175 s or less than m-1.175 s, given that m equals the expected continuous return and s equals the standard deviation of continuous returns. For a normal distribution, 75% of the returns are likely to fall within 1.175s of the expected return m.

A multivariate outlier is more difficult to identify. It represents a set of returns for a particular period that is collectively unusual for one or more reasons. Perhaps one of the returns is sufficiently far from its mean to qualify the collection of returns for that period as an outlier.

Alternatively, a pair of returns that are usually highly correlated may exhibit a sufficient difference in their returns to render the period unusual. Thus a multivariate outlier may result from the unusual performance of an individual asset or from the unusual interaction of a collection of assets, none of which are necessarily unusual in isolation.

How do we determine explicitly whether to classify an observation as “usual” or as an “outlier?” Figure 1 shows two independent return series with equal variances presented as a scatter plot.

Figure 1

Scatter Plot of Independent Return Series with Equal Variances

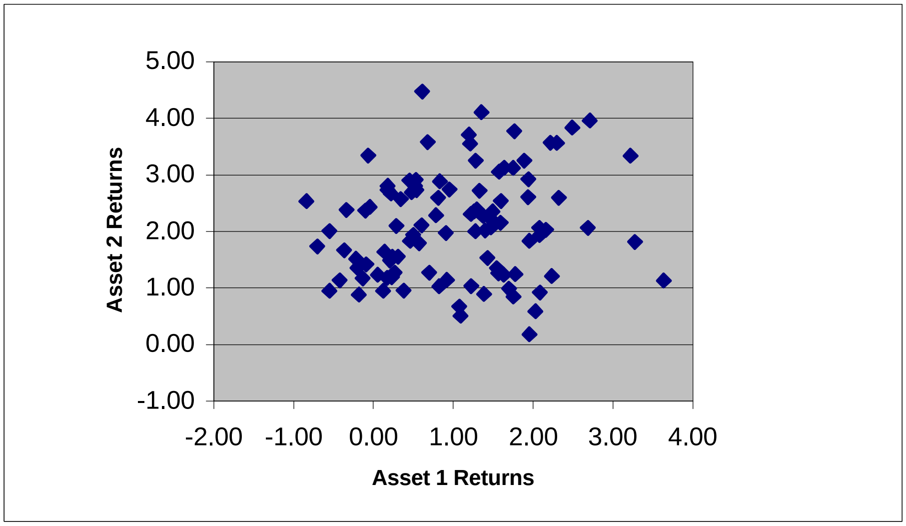

In order to identify outliers, we begin by drawing a circle around the mean of the data, as shown in Figure 2. The radius we choose for the circle is our tolerance for outliers and takes into account the variances of the two return series. The inside circle is our boundary for defining outliers. To determine which observations are outliers, we calculate the equation of a circle for each observation with its center located at the mean of the data and its perimeter passing through a given observation. If the radius of this calculated circle is greater than our "tolerance radius," we define that observation as an outlier. For example, in Figure 2, to determine if the observation designated by the star is an outlier, we calculate the radius of the circle with its center at the mean of the data and its perimeter passing through the star. Because the radius of the calculated circle is greater than the radius of our tolerance circle, the star is an outlier.

Figure 2

Tolerance Circle and Outlier Circle

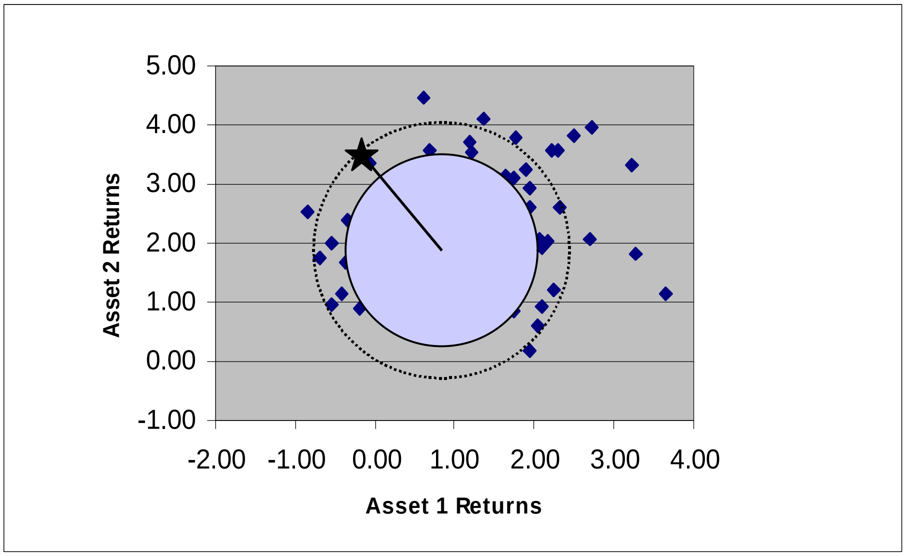

Intuitively and visually, this approach seems reasonable for a given sample of returns if the data are uncorrelated and have the same variance. When the return series have different variances, a circle is no longer appropriate for identifying outliers, as illustrated by Figure 3.

Figure 3

Scatter Plot of Uncorrelated Asset Returns with *Unequal* Variances

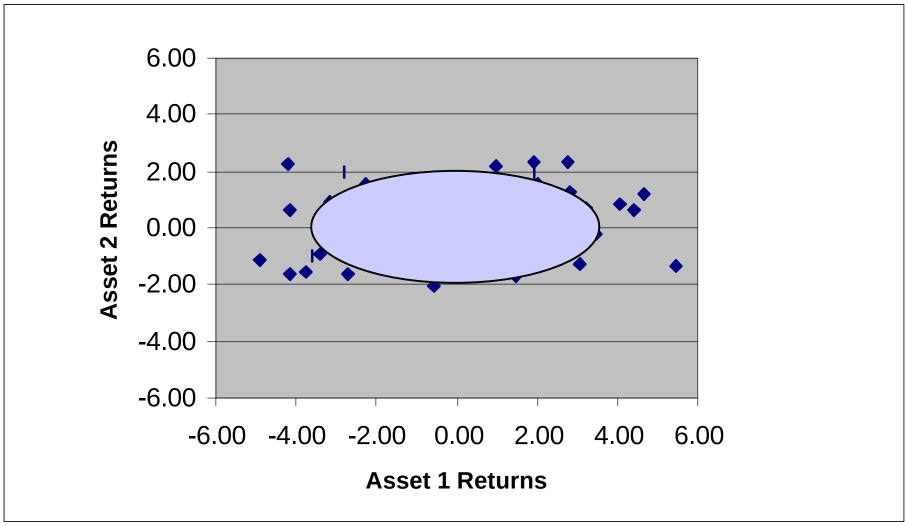

Figure 3 shows a scatter plot of two uncorrelated return series that have unequal variances. Under this condition an ellipse is the appropriate shape for defining the outlier boundary[2]. As before we start with our "tolerance ellipse" and for each point calculate an ellipse with a parallel perimeter. Then we compare their boundaries.

In Figure 4, we relax the remaining assumption and allow the correlation to be non-zero. That the ellipse is positively sloped implies that the return series are positively correlated.

[2] If we were to consider three return series, the outlier boundary would be an ellipsoid, a form that looks something like a football. We are unable to visualize an outlier boundary for samples that include more than three return series - at least not without the assistance of a controlled substance.

Figure 4

Scatter Plot of *Correlated* Asset Returns with Unequal Variances

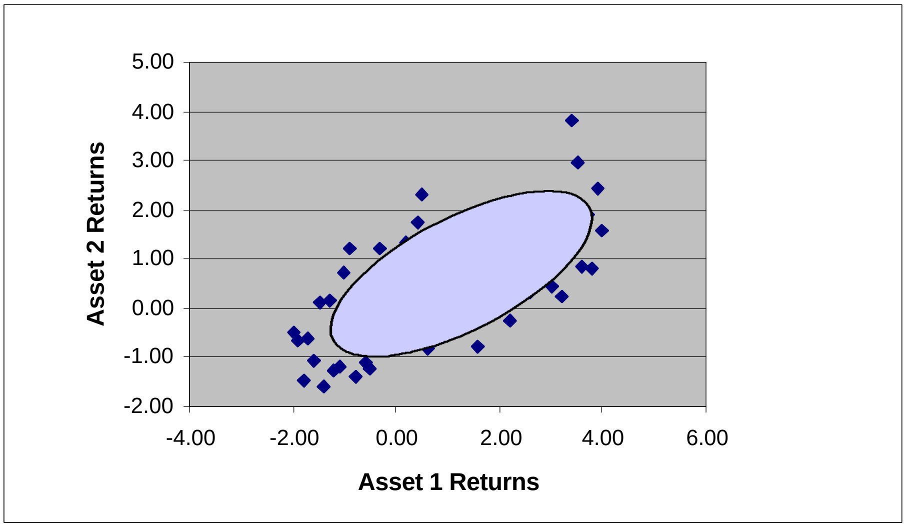

With correlated returns, we generate tolerance ellipses whose axes are rotated. Our basic intuition for identifying an outlier remains unchanged. However, when the return series are correlated or when we expand our sample to include more than three return series, we must employ matrix algebra for the exact computation of an outlier. It is given by Equation (1).

$$d_t = (y_t - \mu)\Sigma^{-1}(y_t - \mu)' \qquad (1)$$

where,

d_t = vector distance from multivariate average

y_t = return series

μ = mean vector of return series y_t

Σ = covariance matrix of return series y_t

The return series y_t is assumed to be normally distributed with a mean vector μ and a covariance matrix □. If we have six return series, for example, an individual observation of y_t would be the set of the six asset returns for a specific measurement interval. We choose our tolerance "distance" and examine the distance, d_t, for each vector in the series. If the observed d_t is greater than our tolerance distance, we define that vector as an outlier.

For two uncorrelated return series, Equation (1) simplifies. It becomes:

$$d_t = \frac{(y - m_y)^2}{s_y^2} + \frac{(x - m_x)^2}{s_x^2} \qquad (2)$$

This is the equation of an ellipse with horizontal and vertical axes. If the variances of the return series are equal, Equation (2) further simplifies into a circle.

For the general n-return normal series case, d_t is distributed as a Chi Square distribution with n degrees of freedom. Under this assumption, if we define an outlier as falling beyond the outer 25% of the distribution and we have 12 return series, our tolerance boundary is a Chi-Square score of 14.84. Using equation (1), we calculate the Chi-Square score for each vector in our series. If the observed score is greater than 14.84, that vector is an outlier.

Optimal Portfolios with Event Varying Covariance Matrices

In the previous section we showed how to identify multivariate outliers from which to estimate risk parameters. These observations are essentially representative of the stressful periods of high volatility that sometimes occur. They are characterized by higher correlations between assets and higher volatility for each asset. Even if investors care only about long term performance, they may not survive long enough to generate long term performance if their strategies are not robust to exceptional periods of stress. This would suggest using the outlying covariance matrix Σ_o, rather than the full sample covariance matrix Σ, for portfolio optimization. However, it would be short sighted to focus only on stressful periods. Such myopia could lead investors to hold unduly conservative portfolios. They would likely fail to achieve their long term objectives. We argue that investors care simultaneously about "quiet" risk and risk during stressful periods. We address this dual focus in two ways. First we allow investors to specify different degrees of risk aversion toward the two regimes. Second, we allow investors to assign different probabilities to the two regimes[3].

We start by augmenting the standard mean-variance objective function to include both measures of risk. The vector of returns has a mean μ and a covariance matrix Σ. It is convenient to think of the returns as originating from a mixture of two distributions, an *"inside"* distribution and an *"outlying"* one. The inside distribution occurs with probability p.[4] Thus, the covariance of returns Σ can be written as:

$$\Sigma = p\, \Sigma_i + (1-p)\, \Sigma_o, \tag{3}$$

[3] It is important to remember that regime, in this context, refers to the inside and outlier samples and not necessarily a period that is chronologically contiguous.

[4] We now are modeling asset returns as if they came from a discrete mixture of two different normal distributions. To derive the chi-square test, we assumed that all returns came from one same normal distribution. This minor contradiction is similar to that occurring in standard statistics where one uses a test valid under a null hypothesis to reject that null hypothesis. In a research in progress, we are studying a – theoretically – more efficient outlier test that makes full use of the model in equation (3). It is unclear whether it will differ much from the simple chi-square test in practice.

where (1-p) is the probability of falling within the outlying sample.[5] Substituting the blended covariance matrix in the standard equation for the expected utility of a portfolio P with a weight vector w, yields:

$$EU(R_p) = w'\mu - \lambda\,(p\, w'\Sigma_i\, w \;+ (1-p)\, w'\Sigma_o\, w), \tag{4}$$

where λ equals aversion toward full sample risk.

Next we allow the investor to vary risk aversion toward the two regimes. We begin by specifying our aversion to inside risk (λ_i) and to outlying risk (λ_o). Then we re-scale them to sum to two. For example, suppose our aversion to inside risk equals 2 and our aversion to outlying risk equals 3. We re-scale our inside risk aversion to equal 0.80 and our outlying risk aversion to equal 1.20, as shown below.

$$\lambda_i' = 2\lambda_i / (\lambda_i + \lambda_o) \tag{5}$$

$$\lambda_o' = 2\lambda_o / (\lambda_i + \lambda_o) \tag{6}$$

We then multiply the probability-weighted inside and outlying covariance matrices by their respective re-scaled risk aversions, as shown in Equation (7)

$$EU(R_p) = w'\mu - \lambda\,(\lambda_i' p\, w'\Sigma_i\, w \;+ \lambda_o'(1-p)\, w'\Sigma_o\, w\,) \tag{7}$$

[5] Note that this equality applies exactly to the estimates of these quantities as well, provided the estimation is consistent with the model. That is the number of observations from the inside and outside scenarios are T_i and T_o such that $T_i/T = (T-T_o)/T = p$.

These substitutions allow us to incorporate, in an intuitive manner, two very different pieces of information into the objective function. First, we are able to express our view as to the likelihood of each distribution for the upcoming investment period. Second, we are able to assign different degrees of risk aversion to the two regimes.

Although Equation (7) has the virtue of transparency, it is somewhat cumbersome. We can simplify it by defining a grand covariance matrix to equal:

$$\Sigma^{*} = (\lambda_{i}' p \Sigma_{i} + \lambda_{o}'(1-p)\, \Sigma_{o}) \tag{8}$$

This definition allows us to recast the objective function as:

$$EU(R_p) = w'\mu - \lambda\, (w'\Sigma^{*} w) \tag{9}$$

We can, of course, express a view on the relative likelihood of the two regimes and, at the same time, assign different risk aversions to them. It is important, however, to separate these two parameters. One is a forecast while the other is a behavioral parameter. Too often, our forecast of risk is confused with our attitude toward risk.

Finally, the internal consistency of our framework is confirmed by the following fact. Equation (7) collapses to the full sample objective function if: (1) we set p equal to the empirical frequency of the inside observations, and (2) we are equally averse to both regimes. Our framework nests the original Markowitz mean variance model in an intuitive manner.

Empirical Results

We use the historical monthly returns of eight asset classes beginning in January 1988 and continuing through September 1998 to demonstrate these procedures. These return series are eight of the asset class benchmarks used by the Harvard Management Company to evaluate the policy portfolio of Harvard's endowment fund. They are:

Domestic Equity	Foreign Bonds
Foreign Equity	High Yield Bonds
Emerging Market Equity	Commodities
Domestic Bonds	Cash Equivalents

The shaded months in Table 1 qualify as outliers for the period January 1998 through September 1998 based on Equation (1) and an outer boundary for which 25% of the multivariate distribution is excluded. This criterion corresponds to a Chi Square score of 10.22.

Table 1
Outliers for the Period 1/98 – 10/98
25% Boundary

	Domestic Equity	Foreign Equity	EM Equity	Domestic Bonds	Foreign Bonds	High Yield Bonds	Commodity	Cash
1/31/98	0.59%	4.29%	-5.70%	1.99%	0.63%	1.21%	-1.61%	0.47%
2/28/98	7.35%	6.88%	10.30%	-0.62%	1.44%	1.49%	-6.60%	0.47%
3/31/98	5.01%	3.77%	3.22%	0.20%	-1.41%	-1.08%	-0.52%	0.47%
4/30/98	1.14%	0.87%	-0.80%	0.40%	2.24%	0.71%	-2.73%	0.47%
5/31/98	-2.26%	-0.17%	-11.91%	1.59%	0.03%	2.09%	-4.72%	0.47%
6/30/98	3.45%	0.73%	-10.30%	1.85%	-0.30%	0.55%	-3.55%	0.47%
7/31/98	-1.66%	1.17%	2.87%	-0.18%	0.33%	-0.41%	-7.86%	0.47%
8/31/98	-15.15%	-12.43%	-25.42%	4.06%	2.73%	-11.80%	-5.90%	0.47%
9/30/98	6.86%	-3.03%	4.89%	3.57%	6.67%	-0.87%	10.22%	0.46%

The two bordered months are among 27 of the 129 months in the full sample that had a Chi Square score in excess of 10.22, which equals 20.9% of the months. That fewer than 25% of the months were selected implies that the multivariate distribution of these returns is slightly platykurtic near its center; more of the observations are clustered near the mean than theory would predict.

It is apparent that the equity returns in August 1998 were extraordinarily low. August 1998 was the month during which Russia defaulted on its sovereign debt, triggering a flight from risky assets. Even commodities, which typically diversify financial assets, generated significant losses. By contrast, the returns during September 1998 were not unusual in their magnitude. However, foreign equities, which typically move in tandem with domestic and emerging market equities, generated a loss while these asset classes produced significant gains. Also, commodities, which typically move inversely with equities, generated a gain during this period. February produced high returns, yet according to our outlier definition, these results were not so strange. When we consider more than three asset classes, it is not always clear why a particular set of returns qualifies or fails to qualify as an outlier because, as mentioned earlier, we cannot visualize the observations that lie outside the boundary. We must rely on mathematical techniques to identify these higher dimension outliers.

Table 2 shows the annualized standard deviations and the correlations estimated from the full sample of 129 months.

Table 2

Full Sample Risk Parameters

	Standard Deviation	Domestic Equity	Foreign Equity	EM Equity	Domestic Fixed	Foreign Fixed	High Yield Fixed	Commodity	Cash
Domestic Equity	12.99	1.00	0.50	0.39	0.39	0.05	0.51	-0.09	0.06
Foreign Equity	17.04	0.50	1.00	0.37	0.17	0.48	0.33	-0.09	-0.07
EM Equity	22.62	0.39	0.37	1.00	-0.08	-0.14	0.34	-0.05	-0.06
Domestic Fixed	6.79	0.39	0.17	-0.08	1.00	0.28	0.13	-0.08	0.10
Foreign Fixed	9.57	0.05	0.48	-0.14	0.28	1.00	-0.03	0.03	-0.04
High Yield Fixed	8.00	0.51	0.33	0.34	0.13	-0.03	1.00	-0.14	-0.13
Commodities	15.94	-0.09	-0.09	-0.05	-0.08	0.03	-0.14	1.00	0.21
Cash	0.44	0.06	-0.07	-0.06	0.10	-0.04	-0.13	0.21	1.00

The average standard deviation of these asset classes based on the full sample is 11.67% while the average correlation is 11.94%. Table 3 presents the same information for the sample of 27 outliers rather than the full sample.

Table 3

Outlier Sample Risk Parameters

	Standard Deviation	Domestic Equity	Foreign Equity	EM Equity	Domestic Bonds	Foreign Bonds	High Yield Bonds	Commodity	Cash
Domestic Equity	20.55	1.00	0.57	0.49	0.43	0.14	0.73	-0.23	0.00
Foreign Equity	27.35	0.57	1.00	0.40	0.27	0.49	0.42	-0.29	0.00
EM Equity	34.29	0.49	0.40	1.00	0.08	-0.18	0.56	-0.25	-0.20
Domestic Fixed	9.74	0.43	0.27	0.08	1.00	0.32	0.13	-0.14	0.20
Foreign Fixed	14.09	0.14	0.49	-0.18	0.32	1.00	-0.12	-0.09	0.12
High Yield Fixed	14.85	0.73	0.42	0.56	0.13	-0.12	1.00	-0.18	-0.01
Commodities	24.83	-0.23	-0.29	-0.25	-0.14	-0.09	-0.18	1.00	0.20
Cash	0.47	0.00	0.00	-0.20	0.20	0.12	-0.01	0.20	1.00

Note that we use 25% as the outlying area of the distribution. There is of course a temptation to use this method to explore more extreme outliers. This of course causes the number of outliers selected to drop. This in turn may cause serious problems in the estimation of the correlation matrix in Table 3. There should always be far more observations than securities for the covariance matrix estimator to be reasonably precise. Here we have 27 observations for an 8 by 8 covariance matrix which includes 8*9/2 different parameters. As a rule of thumb, one would want at least twice as many observations as assets.[6]

The average standard deviation for the outlier sample is 18.27%, a 57% increase over the full sample standard deviation. An increase is to be expected by the very nature of the outlier selection process. The average correlation also rises for the outlier sample, but only by 15% to 13.72%. However, the change in the average correlation obscures important details about the diversification properties of the various asset classes. For example, excluding commodities, the average correlation increases 36% from 16.86% to 22.96%.[7] By contrast, the average correlation of commodities declines by a factor of 5, from –2.81% to –13.99%. Hence the full sample correlations belie the weaker diversification properties of financial assets during times of stress, yet understate the diversification benefits of commodities when markets are in turmoil. Figure 5 summarizes the differences between the full sample risk parameters and the outlier sample risk parameters.

---

[6] In the limit the estimator is meaningless if there are fewer observations than assets. For large covariance matrices, the only way out of this problem may be the use of a factor model to restrict the number of parameters in the covariance matrix.

[7] It can be shown that correlations between assets are higher for a subsample of outliers than for the whole sample, when the returns come from the same distribution. An interesting research in progress is to determine what fraction of the higher correlations observed is due to this, and what fraction is due to the fact that outliers do come

Figure 5

Full Sample and Outlier Sample Differences

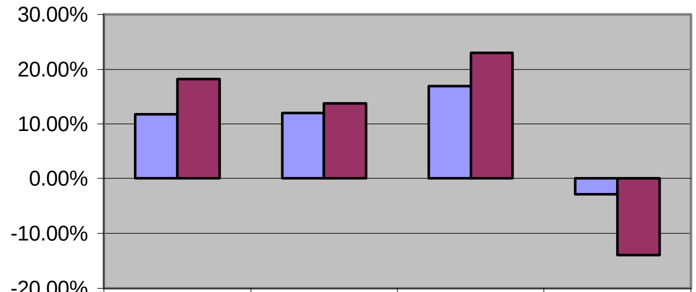

	Risk	Corr	Corr Fin	Corr Com
Full Sample	11.67%	11.94%	16.86%	-2.81%
Outlier Sample	18.27%	13.72%	22.96%	-13.99%

Next we show how the composition of the optimal portfolio shifts as we switch from the full sample risk parameters to the outlier sample risk parameters. In both optimizations we assume the following means for the asset class returns; thus the changes reflect only the differences in risk parameters[8].

Domestic Equity	10.00%
Foreign Equity	10.25%
Emerging Market Equity	12.00%
Domestic Fixed Income	7.00%
Foreign Fixed Income	7.25%
High Yield Fixed Income	8.00%
Commodities	6.00%
Cash Equivalents	5.00%

from a different distribution. In any case, these higher correlations are those faced by the investor when the returns are more volatile.

Table 4

Full Sample and Outlier Sample Optimal Portfolios

	Full Sample Optimal Mix	Outlier Sample Optimal Mix
Domestic Equity	25%	2%
Foreign Equity	0%	0%
EM Equity	16%	11%
Domestic Bonds	8%	28%
Foreign Bonds	26%	25%
High Yield Bonds	22%	16%
Commodities	3%	12%
Cash	0%	7%
Normal Environment		
Expected Return	8.82%	7.58%
Standard Deviation	7.27%	4.37%
Stressful Environment		
Expected Return	8.82%	7.58%
Standard Deviation	12.32%	7.20%

Table 4 shows the optimal portfolio weights assuming that our risk aversion equals 2.5; that is, we are willing to sacrifice 2.5 units of expected return in order to lower our portfolio's variance by one unit. Given this degree of risk aversion, together with the covariance matrix estimated from the full sample of monthly returns, we should allocate most of our portfolio to equities, foreign assets, and high yield bonds. Very little of the portfolio should be allocated to domestic bonds and commodities. This allocation offers an expected return of 8.82% with a standard deviation of only 7.27%, assuming the full sample covariance matrix characterizes the risk of our investment horizon.

[8] These expected returns are chosen for their reasonableness and are not related to either the full sample means or the outlier sample means, nor do they represent the views of the Harvard Management Company.

However, when we experience an environment that is better represented by the risk parameters associated with the outlier sample, then this portfolio will experience much greater volatility, as shown by the almost 70% increase in its standard deviation from 7.27% to 12.32%.

If we were to structure an optimal portfolio for an outlier environment, we should reduce the equity component to 13%, and increase commodities to 12% and fixed income investments to about 75% of the portfolio. This portfolio is optimal for periods that are characterized by the covariance matrix estimated from the sample's outliers, assuming that our risk aversion remains constant. This shift reduces the portfolio's volatility during such periods from 12.32% to 7.20%, and it also has lower volatility should the normal environment prevail. Such a shift, however, is not without cost. These changes in allocation reduce expected return from 8.82% to 7.58%.

Herein lies the problem. If we optimize based on the full sample covariance matrix, the portfolio will be significantly sub-optimal when we encounter periods of financial stress, and indeed may not survive such a period without unpropitious adjustments. However, if we optimize based on the outlier covariance matrix, the portfolio's expected return for the full horizon is lower than desired. As with many choices, the best solution is to compromise. Table 5 presents optimal portfolios based on blended covariance matrices along with their expected performance during quiet environments (inside sample), normal environments (full sample), and stressful environments (outlying sample).

Table 5

Optimal Portfolios from Blended Covariance Matrices

	Empirical Probability Higher Outlier Aversion	Equal Probability Equal Aversion	Equal Probability Higher Outlier Aversion
D Equity	21%	12%	9%
F Equity	0%	0%	0%
EM Equity	16%	14%	13%
D Bonds	14%	23%	25%
F Bonds	26%	26%	27%
HY Bonds	17%	14%	15%
Commodity	6%	10%	11%
Cash	0%	0%	0%
Quiet Environment			
Expected Return	8.60%	8.18%	8.05%
Standard Deviation	4.91%	4.50%	4.38%
Normal Environment			
Expected Return	8.60%	8.18%	8.05%
Standard Deviation	6.70%	5.81%	5.58%
Stress Environment			
Expected Return	8.60%	8.18%	8.05%
Standard Deviation	11.03%	9.13%	8.64%

The first portfolio shown in Table 5 assumes that outliers will occur with the same frequency as they occurred empirically, but that we are 1.5 times as averse to outlier risk as we are to risk during quiet periods. As expected, this relatively higher aversion to outlier risk, when compared to the full sample optimal portfolio, shifts the portfolio away from equities and toward commodities and fixed income assets.

The next column shows the optimal portfolio assuming that we are equally averse to outlier risk and quiet period risk, but that we believe outlier events will occur 50% of the time rather than their empirical frequency of 20.9% of the months. Again, the optimal portfolio is more conservative than the full sample portfolio because, even though we are equally averse to both environments, we expect stressful periods to occur more frequently than they did in the full sample.

The final column shows the optimal portfolio assuming that outlier events will occur 50% of the time rather than their empirical frequency and that we are 1.5 times as averse to outlier risk as we are to risk during quiet periods. Thus we emphasize the outlier sample in two ways – by assigning it a greater probability than its empirical frequency and by raising our relative aversion to outlier risk. This dual emphasis results in an optimal portfolio that closely resembles the portfolio estimated solely from the outlier sample.

Conclusion

We introduce a methodology to address the stylized fact that risk is unstable. Specifically, we identify multivariate outliers and use these outliers to estimate a new covariance matrix. We argue that a covariance matrix estimated from outliers provides a better representation of a portfolio's riskiness during periods of financial stress than a covariance matrix estimated from the full sample of observations.

We also introduce a procedure for blending an inside sample covariance matrix with an outlier sample covariance matrix. This procedure enables investors to express their views about the likelihood of each regime and to differentiate their aversion to them.

We present empirical results that support the view that volatility and correlations estimated from outliers differ significantly from full sample estimates. In addition, we identify optimal portfolios from both covariance matrices, holding constant expected returns and risk aversion. Given our sample, the volatility of the optimal portfolio estimated from the full sample covariance matrix nearly doubles when it is subjected to the outlier sample covariance matrix. As expected, the outlier sample covariance matrix produces a much more conservative optimal mix but with a concomitantly lower expected return.

Finally, we present results based on covariance matrices that are blended from quiet and stressful periods. These results show the sensitivity of portfolio weights to variations in our views about the relative likelihood of quiet and stressful times and our relative aversion to each risk regime.

References

Chow, George, 1995, "Portfolio Selection Based on Return, Risk, and Relative Performance," *Financial Analysts Journal*, vol. 51, no.2 (May/June): 54-60.

Markowitz, Harry, 1952, "Portfolio Selection," *The Journal of Finance*, vol. 7, no. 1 (March 1952): 77-91.

Michaud, Richard, "Efficient Asset Management: A Practical Guide to Stock Portfolio Optimization and Asset Allocation," *Harvard Business School Press*, (June 1998).

Siconolfi, Michael, Anita Raghavan, and Mitchell Pacelle, "All Bets Are Off: How the Salesmanship and Brainpower Failed at Long-Term Capital," *Wall Street Journal*, (November 16, 1998): A1, A18-19.