COLEMAN FUNG
RISK MANAGEMENT
RESEARCH CENTER

Working Paper # 2011 -01

Minimizing Shortfall

Lisa Goldberg, University of California Berkeley,MSCI
Michael Y. Hayes, MSCI
Ola Mahmoud, MSCI

2011

**University of California
Berkeley**

Minimizing Shortfall

Lisa R. Goldberg
Department of Statistics, University of California at Berkeley
MSCI
2100 Milvia Street
Berkeley, CA 94704
USA
lisa.goldberg@msci.com

Michael Y. Hayes
MSCI
2100 Milvia Street
Berkeley, CA 94704
USA
michael.hayes@msci.com

Ola Mahmoud
MSCI
Rue de la Confederation 8-10, 3rd Floor
CH 1204 Geneva
Switzerland
ola.mahmoud@msci.com

1 Introduction

Despite the increasing sophistication of Finance in the past 30 years, quantitative tools for building portfolios remain entrenched in the paradigm proposed by Markowitz in 1952; these tools offer investors a trade-off between mean return and variance. However, Markowitz himself was not satisfied with variance, which penalizes gains and losses equally. Instead, he preferred semi-deviation that only penalizes losses.

Recent work has made downside risk optimization practical, but there has been no reliable non-Gaussian risk model. This void is filled with Barra Extreme Risk (BxR), an empirical, fundamental factor-based model that captures features of return beyond variance. BxR reflects persistent characteristics, such as the higher asymmetry and downside risk of high-yield bonds compared to government bonds, or Growth stocks compared to Value stocks.

This paper describes an empirical study of shortfall optimization with Barra Extreme Risk. We compare minimum shortfall to minimum variance portfolios in the US, UK, and Japanese equity markets using Barra Style Factors (Value, Growth, Momentum, etc.). We show that minimizing shortfall generally improves performance over minimizing variance, especially during down-markets, over the period 1985-2010. The outperformance of shortfall is due to intuitive tilts towards protective factors like Value, and away from aggressive factors like Growth and Momentum. The outperformance is largest for the shortfall that measures overall asymmetry rather than the extreme losses.

2 Background

In this section we review the definition and motivation for shortfall, the formulation of variance and shortfall optimization problems, and the Barra Extreme Risk model.

2.1 Volatility and Shortfall as Risk Measures

Volatility, or the square root of variance, measures the average dispersion over the entire distribution of portfolio gains and losses. Volatility is the central concept in many standard statistics such as risk contribution, beta, and correlation (Goldberg, et al, 2010). Its usefulness stems from its empirical and mathematical properties. Empirically, volatility is persistent from one period to another; realized volatility in one period is generally highly correlated with realized volatility in the next. Mathematically, volatility is a convex risk measure, amenable to the tools of convex optimization and analysis (Goldberg and Hayes, 2010). A minimum of a convex risk measure is unique, so once a minimum is found it is guaranteed to be a global minimum.

Although useful, volatility does not describe every aspect of risk. Even as he proposed volatility as a risk measure, Markowitz (1952) pointed out that a better risk measure would only penalize losses, and he proposed semi-deviation as a desirable alternative. An alternative risk measure that has gained attention in recent years is *shortfall*, which is the average (or expected) value of the largest losses. The shortfall confidence level specifies the magnitude of these largest losses. For example, the 95% shortfall is the average over the 5% largest losses. Given N possible portfolio outcomes, shortfall is formally defined as

$$s_p = \frac{1}{N(1-p)} \sum_{i=1}^{N(1-p)} r_{(i)} \qquad (1)$$

where $r_{(i)}$ are the ordered return scenarios and p is the confidence level.[1]

2.2 Variance and Shortfall Optimization

Like variance, shortfall is a convex risk measure and can be efficiently minimized. For asset weights w, expected returns α, risk aversion λ, and covariance matrix Σ, the standard mean-variance problem is:

$$\max_w \; w'\alpha - \lambda w'\Sigma w \tag{2}$$

Similarly, the mean-shortfall problem is:

$$\max_w \; w'\alpha - \lambda s_p(w) \tag{3}$$

where $s_p(w)$ is the empirical shortfall estimator at confidence level p. Rockafellar and Uryasev (2000, 2002), Krokhmal, et al (2002), and Bertsimas, et al (2004) show how to formulate shortfall optimization as a linear program (LP) amenable to standard optimization algorithms. In Appendix C we review the formulation of variance, shortfall, and combined variance-shortfall optimization in more detail.

2.3 The Barra Extreme Risk Model

Although the promise of alternative risk measures has been recognized since variance was first introduced, the missing link has been a reliable forecast of these risk measures. One approach used for downside risk has been a Normal (Gaussian) model. However, the Normal model reduces to a variance model for linear instruments such as equities. This is because any risk measure[2] is a fixed multiple of volatility when returns are normally distributed. Therefore, Normal models of risk do not add additional insight for linear instruments, even when looking at alternative risk measures. To make use of alternative risk measures, a non-normal risk model is needed.

A non-parametric approach uses historical returns as forecast return scenarios (known in the context of Value at Risk as Historical VaR). While avoiding any distributional assumptions, Historical VaR assumes that returns drawn from history are representative of the future. It is widely accepted that volatility (and covariances) change over time, which constitutes an argument against Historical VaR. Moreover, Historical VaR generally uses a relatively short return history (e.g., 1 year). Still, certain assets may not have sufficient history (e.g., newly issued equities), or their history may be irrelevant (companies that change from growth to value, small cap to large cap, or from one industry to another).

These obstacles to historical estimation are addressed by Barra Extreme Risk (BxR). BxR generates forecast scenarios using Barra factor return history, which is uniformly available over a long period, and represents return characteristics based on fundamental company characteristics. Furthermore, in the reduced dimensionality of Barra factors, it is possible to account for the difference between the current

[1] Here we assume that N(1-p) is an integer. The general formula for continuous outcomes is conceptually similar; see Acerbi and Tasche (2001) for a detailed discussion.

[2] More specifically, any risk measure that is a function of the single-period return distribution.

and historical covariance. By accounting for this difference, history is made relevant to the current risk climate, and risk forecasts respond quickly to changing market conditions. The BxR approach is reviewed in Appendix A, and Dubikovsky, et al (2010) present broad out-of-sample tests that show how the BxR model is more consistent with market behavior than the conditional normal model.

3 Uncertainty in Shortfall Forecasts

Before endeavoring to construct portfolios based on alternative measures, a basic question must be answered: can the risk measure be estimated with enough certainty to be useful? Two aspects of forecast uncertainty are precision, measured by estimation error, and predictability, measured by persistence (i.e., do historical returns predict future risk?).

3.1 Estimation Error

Optimized weights are subject to error, even for a perfect risk model, because risk is always estimated with a finite sample. Kondor, et al (2007) explain that estimation error increases with the ratio N/T, where N is the number of assets and T is the sample length. Naturally, estimation error plays a larger role in shortfall than in volatility, because a large amount of the input data is only used in aggregate to define the largest losses.

We briefly study the effect of estimation error using simulated data. Because the true distribution of the simulated variables is known, the true optimal portfolio is also known. Estimation error is measured in two ways. First, we compute the average risk of the optimized portfolio divided by the true minimum risk (risk error). Second, we compute the average angle between the optimized weight vector and the true optimal weight vector (weight error). Both of these measure the proximity of the optimized portfolio to the true optimum. The risk error measures the average amount of extra risk that is taken in optimized portfolios due to random fluctuations. Because it is not obvious how much risk error is acceptable, we also compute the weight error that has a concrete acceptable upper bound. This upper bound is defined by the average weight error of a randomly chosen set of positive weights. If the average weight error is greater than this upper bound, an investor is better off guessing at a random set of weights rather than trying to compute an optimal portfolio. This upper bound is around 35 degrees; we compute an analytical formula for the upper bound in Appendix B. The concept of weight error is illustrated for a two-asset portfolio in Figure 1.

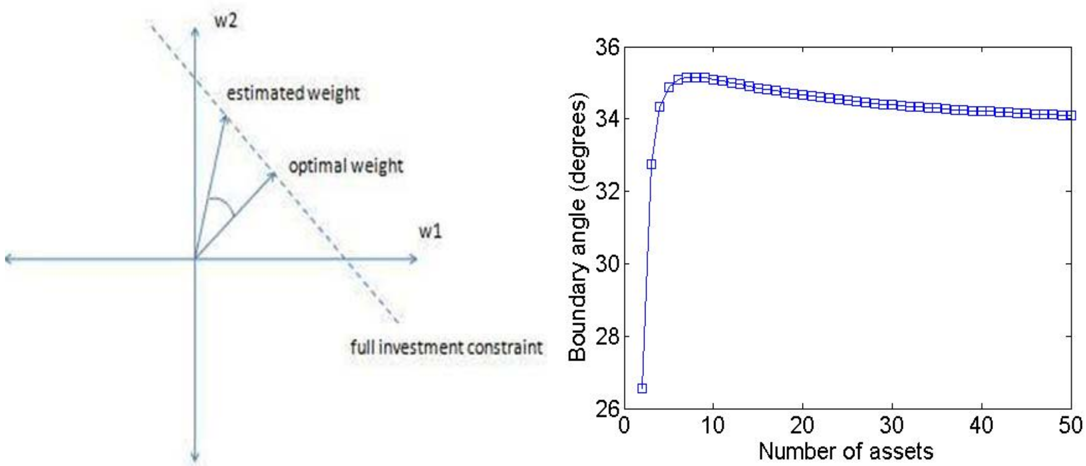

Figure 1: *Illustration of weight-error angle for a two-asset portfolio (left) and its upper bound (right).*

We measure these effects using simulated standard normal random variables applied to 10 assets with sample lengths of 1000, 3000, 5000, and 7000. For each sample length, we simulate returns and minimize shortfall at several confidence levels using the full, equal-weighted sample. We repeat this process 100 times and take the average across the optimized portfolios. The results are shown in Figure 2.

We see that for our parameters, in this simplified setting, estimation error is negligible for all confidence levels tested. While this does not rule out large estimation error for all possible distributions, our parameters satisfy the baseline normal criterion. We are thus able to control estimation error in this study by considering a low-dimensional risk space (fundamental factors) and a long history of factor returns from 1981.

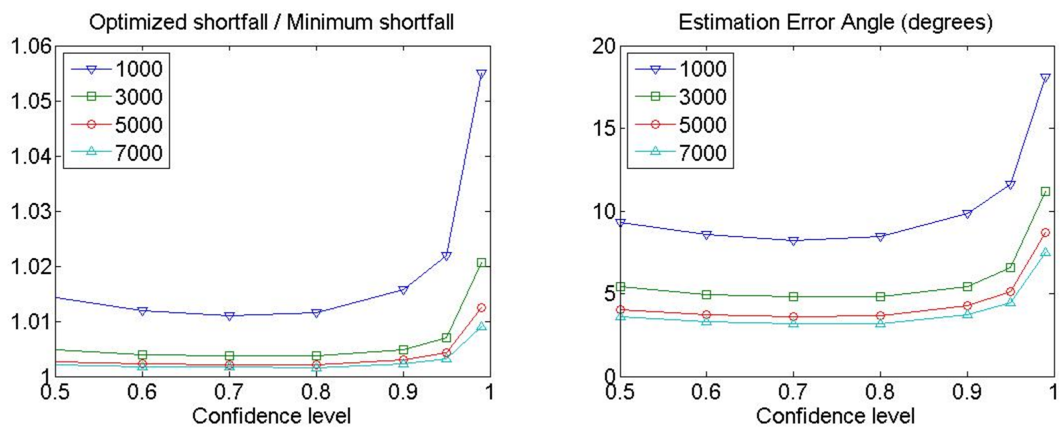

Figure 2: *Measures of estimation error using simulated, standard normal random variables.*

3.2 Persistence

The complementary question to estimation error is persistence. Even if we were able to exactly measure the risk in one period, and it does not persist in the next period, then it is of no use in risk forecasting or portfolio construction. This is especially relevant for the BxR model, which uses up to 30 years of daily return history to build its forecasts. We postulate that Barra fundamental factors (Size, Growth, Value, etc.) reflect characteristics that are "fundamental" to stock behavior, and are thus persistent across long periods of time. To test this hypothesis, we compute BxR forecasts using disjoint 15-year samples. Because BxR allows for volatility to evolve over time, we focus on a *non-normality* (NN) statistic that is independent of volatility.[3] The NN statistic is a measure of percent deviation from normality, formally defined as the percent difference between BxR shortfall (xShortfall) and a Normal shortfall forecast:

$$\mathrm{NN} = \frac{\mathrm{xShortfall}}{\mathrm{Normal\ Shortfall}} - 1.$$

[3] More specifically, NN is independent of volatility when Normal Shortfall is estimated using the same half-life as that used to normalize the BxR returns. Mathematically, NN of a factor is the same as NN of a factor times a positive constant.

Positive NN means that the BxR shortfall forecast exceeds the normal estimate; negative NN means that the normal estimate exceeds the BxR forecast. Zero NN implies that the BxR and Normal estimates coincide. We compute the NN statistics of Barra USE3 Style factors, on the gain and loss tails, and test the null hypothesis that the NN statistics persist from one period to the next.[4] The results are shown in Figure 3, along with confidence intervals on the difference. If the confidence interval crosses the dashed diagonal, we cannot reject the null hypothesis that the NN statistics persist across periods. Out of 36 factors, we find only a handful of significant outliers: Size at three confidence levels, and Earnings Yield at the 60% confidence level. This result shows that Barra fundamental factors display persistent gain/loss tail features, even over long time periods, supporting the proposition that suitably modified long histories of fundamental factors are a useful input to portfolio construction. Furthermore, the long history provides a large number of forecast scenarios, allowing us to control estimation error as described in the previous section.

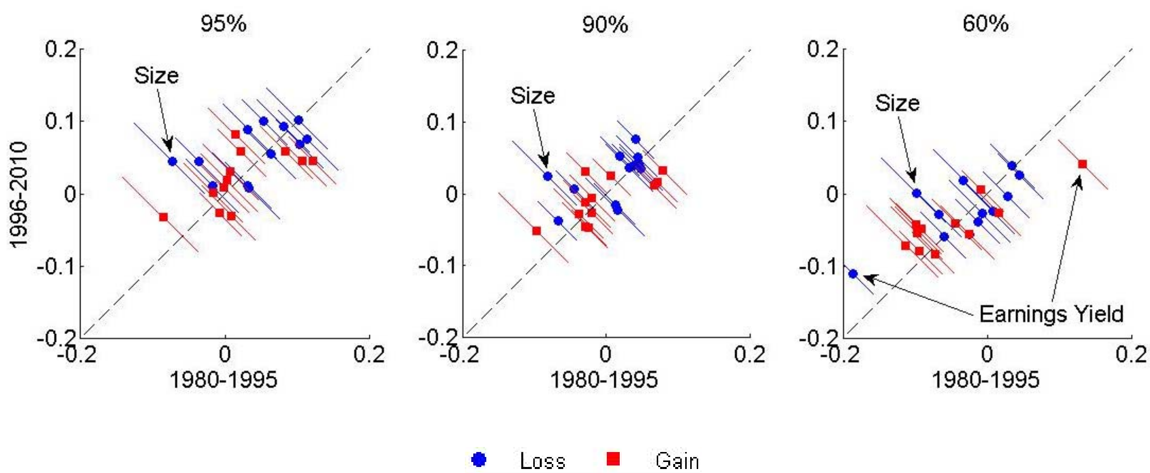

Figure 3: *Persistence of non-normality (NN) of loss/gain tails of Style factors in the Barra US factor model (USE3).*

4 Optimization Framework

For our empirical study, we consider the active management problem in terms of fundamental factors. We examine the performance of an active strategy that is long minimum shortfall and short minimum variance. Alternatively, this performance measures the value added by minimizing shortfall instead of minimizing variance.

To simplify the analysis, we do not include expected returns (alpha), and seek only to minimize risk. We further simplify the analysis by considering a small optimization universe consisting of a market index and a small number of Barra Style Factor portfolios (Size, Growth, Value, etc.; see Table 4).[5] We carry out the study in three markets: US, UK, and Japan, during the period 1985-2010.

Each minimum shortfall and minimum variance portfolio is constrained as follows: the weight of the index is set to 100% to reflect full investment in the market portfolio. Consequently, the index weight of

[4] Confidence intervals are computed by bootstrapping the BxR scaled returns in each period.

[5] A factor portfolio return is equivalent to the Barra factor return; some Barra Style factor portfolios are also listed as MSCI indices.

the active strategy is zero. We constrain each individual Style factor exposure to the range [-2, 2]. The sum of Style factor exposures is set to zero, so that the active bets are dollar-neutral. This enforces a reasonable level of Style exposure that can be achieved using a moderately sized investment universe.

Inputs to the shortfall optimizer are daily returns prior to the analysis date. The return history begins in 1981 and the backtesting period starts in 1985, so the shortfall optimization is informed by a minimum history of 4 years. These time-series are adjusted using the BxR methodology as explained in Appendix A, representing forecast return scenarios in the shortfall objective function. Covariance forecasts are made using an exponentially weighted moving average (EWMA) of trailing factor returns.

We tested a range of parameters, including: the shortfall confidence level (60% to 99%); the correlation and volatility half-lives used for BxR covariance rescaling and for the forecast covariance matrix (21, 90, 180 days); and the rebalancing frequency (daily, weekly, monthly, quarterly). With the exception of confidence level, we find little sensitivity to the optimization parameters. Here we focus on a 21-day half-life and a monthly rebalancing frequency.

US	**UK**	**Japan**
(MSCI USA)	(MSCI UK)	(MSCI Japan)
Volatility	Size	Volatility
Momentum	Momentum	Size
Size	Volatility	Momentum
Trading Activity	Trading Activity	Trading Activity
Growth	Leverage	Value
Earnings Yield	Value	Interest Rate Sensitivity
Value	Yield	Growth
Earnings Variability	Foreign Sensitivity	Leverage
Leverage	Growth	Foreign Sensitivity
Currency Sensitivity		
Yield		

***Table 1**: Equity Style factors and "market factors" (MSCI USA, UK, Japan) used in optimization.*

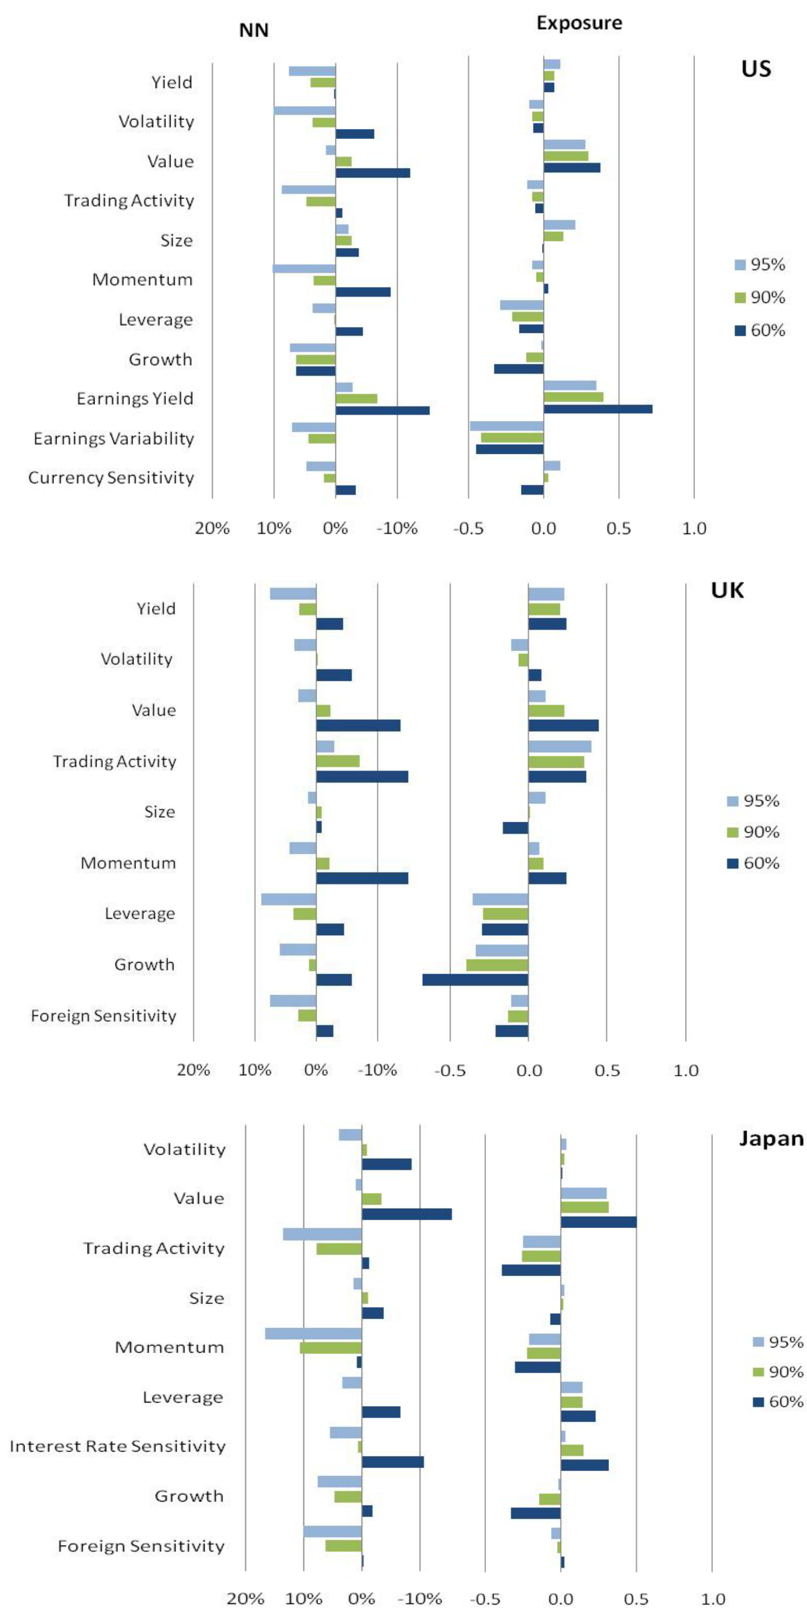

Figure 4: *Non-normality (NN) of Style factors (left) and average monthly excess exposure of minimum shortfall over minimum variance (right).*

5 Optimization Results

5.1 Optimal Exposures

Style factor exposures of the active portfolio are shown in Figure 5. The active strategy tilts consistently towards Value in the UK and Japan, and towards the related Earnings Yield factor in the US. It tilts away from Leverage, Growth, and Earnings Variability in the US, away from Leverage and Growth in the UK, and away from Trading Activity, Leverage, and Growth in Japan.

The excess exposures can be partly understood in terms of the NN statistics, also shown in Figures 5 and 6. The factors favored by minimum shortfall have smaller (more negative) NN statistics, and those avoided have larger (more positive) NN statistics. The largest visible exception to this trend is 60% Growth in the UK, which is under-weighted in spite of its negative NN. This can be partly explained by the fact that Growth is the fourth riskiest factor in the UK market (as measured by NN). Other exceptions include Currency Sensitivity in the US and the related Foreign Sensitivity in the UK and Japan, which can also be explained by their NN relative to the other available factors.

The Volatility, Size, and Trading Activity factors are highly correlated with the market, and are an attractive hedge of market risk in both minimum variance and minimum shortfall. Consequently, their exposures consistently approach or coincide with the lower bound of -2. This means that their exposure in the active portfolio is close to zero. Two exceptions are Trading Activity in Japan and the UK. In Japan, Trading Activity is one of the riskiest factors (by NN) and therefore favored by the variance optimizer. In the UK, by contrast, Trading Activity is one of the least risky factors (by NN), so it is favored by the shortfall optimizer.

BxR captures not only the extremes, but also the overall asymmetry of the distribution. The NN statistics show that at high confidence levels most factors are riskier than Normal (*fat-tailed*). However, at the 60% level, where shortfall examines nearly the entire loss side of the distribution, many factors are *less* risky than Normal (*positively skewed*). For the Barra Style factors, the strongest non-Normal risk signal comes from overall asymmetry rather than the extremes. This leads to larger active bets at lower confidence levels.

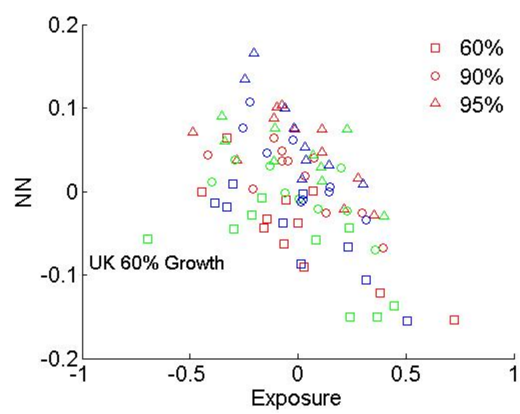

Figure 5: *Relationship of active factor exposure to non-normality (NN); (US – red, UK – green, Japan – blue).*

5.2 Optimal Portfolios

Performance

Figure 7 shows the cumulative returns of the active portfolios. The active strategy shows consistently strong performance during the entire back-testing period, for all confidence levels, and in all three equity markets. In other words, minimum shortfall consistently outperforms minimum variance. We also see that the lower the shortfall confidence level, the larger the outperformance. It is worth noting that in the 1-2 years leading up to a financial crisis (e.g., 1986, 1998-99), minimum variance outperforms minimum shortfall. This is followed by a large improvement in the minimum shortfall portfolios during the subsequent turmoil.

When comparing these returns to the market returns, we see that (especially in the US) they are almost mirror images of one another. When a crisis hits, the active portfolios remain unaffected and even show gains. This observation suggests that the active portfolios (long minimum shortfall, short minimum variance) can be used for downside protection that outperforms the market and limits losses during turbulent times.

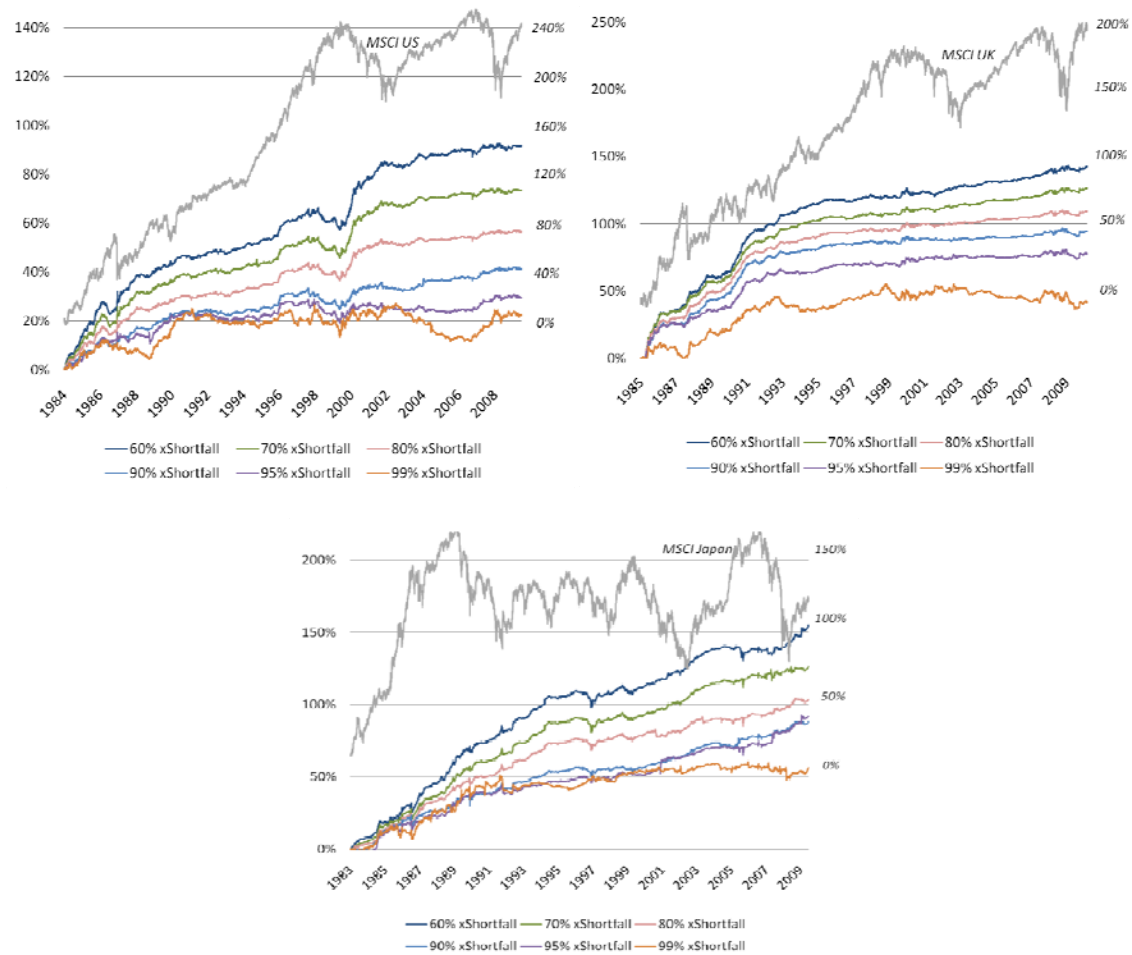

__Figure 6:__ Daily cumulative returns of xShortfall-optimal portfolios in excess of variance-optimal portfolio for the US (top left), UK (top right) and Japan (bottom). The corresponding cumulative returns of the market index are shown in grey, with their scale given by the axis on the right.

The outperformance and downside protection that the minimum shortfall portfolios offer is apparent when comparing their absolute cumulative returns to those of the market (see Appendix E). In all three markets, the 60% xShortfall optimal portfolio constantly floats above the market. On the other hand, the variance optimal portfolio either exceeds the market (UK) or underperforms, especially in down-markets (US and Japan).

Return Attribution

To understand the outperformance of minimum shortfall, we perform a return attribution on the excess returns of minimum shortfall over minimum variance. Figure 8 shows the cumulative returns for each Style factor multiplied by their excess exposures (at the 60% shortfall confidence level). Ignoring compounding effects, the sum of the cumulative returns of each of these factors equals the excess returns of minimum shortfall over minimum variance.

Most of the outperformance of minimum shortfall is due to tilts towards Earnings Yield (US), Value (UK and Japan), and Trading Activity (UK), and tilts away from Growth (US, UK, Japan), Trading Activity (Japan), Leverage (Japan), and Momentum (Japan). Most of these tilts correspond to conventional wisdom about stock characteristics: Value is protective, and Growth, Leverage, and Momentum are aggressive. A notable exception is Trading Activity, which plays little role in the US, but is favored in the UK and avoided in Japan. This may be partly due to the different definition of Trading Activity in the two models,[6] but more likely reflects qualitative differences between the Japanese and UK equity markets. The Japanese equity market has been bearish for most of the test period, while the UK has followed the global business cycle.

In Figure 5 we showed that the excess exposures decrease for higher shortfall confidence levels. This is because the strongest signal of non-normality comes from examining the entire loss side of the distribution, and not just the tail. Strikingly, the return attribution shows that considering the core of the return distribution will limit losses more effectively than looking deeper into the loss tail.

Risk Analysis

Having looked at the performance of minimum shortfall portfolios, relative to their minimum variance counterparts, a remaining question is whether minimizing shortfall reduces extreme risk. To answer this, we compute the realized volatility, realized Sharpe ratio (realized return/realized volatility), and realized 95% shortfall (the average over the largest losses), for down- and up-markets (see Table 9). For down-markets we take the crisis periods of 1987-1988, 2000-2002, and 2007-2008, and for up-markets the remaining years in our backtesting period. When comparing the variance and shortfall optimal portfolios in absolute terms, we see that realized volatility is similar for all optimal portfolios, and consistently lower than that of the market. The Sharpe ratio is a little higher than that of the market, and highest for the 60% shortfall optimal portfolio. All optimized portfolios have significantly lower realized shortfall compared to the market. The active portfolios have negligible realized risk (volatility and shortfall) and higher Sharpe ratios, especially during down-markets. We finally see that in the framework of full-investment in the market, together with an active hedge using our optimal excess portfolios, realized risk would only decrease by 1%, but realized shortfall would be reduced during down-markets. Analysis of the realized risk of the optimal portfolios in the UK and Japanese markets shows similar results.

[6] Trading Activity in UKE7 is a weighted average of monthly, quarterly, and annual share turnover, while in JPE3 it also includes recent growth in trading volume; however, this descriptor is weighted by only 3%; the remainder of the factor is defined analogously to the UKE7 factor.

Figure 7: *Return attribution for the monthly rebalanced 60% xShortfall optimal portfolio. For each Style factor, its contribution to the overall return is obtained by multiplying its cumulative returns by the excess optimal exposures.*

	Realized Volatility		**Realized Sharpe Ratio**		**Realized 95% Shortfall**	
	Up-markets	*Down-markets*	*Up-markets*	*Down-markets*	*Up-markets*	*Down-markets*
MSCI USA	15%	27%	0.07	-0.02	2.16%	3.98%
Variance-optimal	11%	19%	0.08	-0.01	1.56%	3.13%
60% xShortfall optimal	11%	19%	0.10	0.01	1.57%	3.00%
95% xShortfall optimal	11%	20%	0.08	0.00	1.59%	3.15%
60% active portfolio	1%	2%	0.05	0.10	0.22%	0.26%
95% active portfolio	1%	2%	0.01	0.02	0.22%	0.25%
MSCI USA + 60% active	14%	26%	0.09	-0.02	1.94%	3.78%
MSCI USA + 95% active	14%	26%	0.09	-0.02	1.96%	3.83%

Table 2: *Comparison of average realized volatility, realized Shape ratio (realized return/realized volatility), and realized 95% shortfall.*

The variance- and shortfall-optimal portfolio beta is shown in Figure 10, using a two-year rolling window.[7] All optimized portfolios display similar betas over time, which means that the beta of the active portfolio is nearly zero. Beta is on average significantly smaller than one, going as low as 0.2, indicating that the returns of our optimized portfolios do not generally follow market returns.

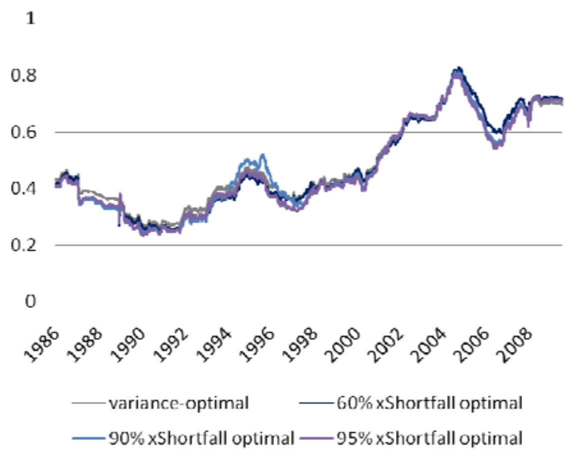

Figure 8: *Evolution of betas for optimal portfolios in the US market. Standard volatility beta with respect to the market is calculated using a two-year rolling window.*

6 Conclusion

Our empirical study has shown that shortfall optimization, combined with Barra Extreme Risk (BxR), can capture information beyond variance, and that this information can translate to downside protection and outperformance. The key element of BxR is the Barra fundamental factor model, which provides a consistent and uniform view of history. Variance has so far been the main risk measure to optimize against because of its simple quadratic definition and its empirical persistence. With the advent of Barra Extreme Risk, downside risk optimization has become a viable alternative to variance optimization.

[7] Beta is computed with respect to the market, represented by the MSCI USA Index.

Appendix A: Barra Extreme Risk Methodology

The Barra Extreme Risk (BxR) model assumes that Barra fundamental factor returns are stationary, but their covariance evolves with time. For a vector of factor returns f with covariance matrix Σ, we have

$$f = \Sigma^{1/2} g \qquad (4)$$

where the uncorrelated factor returns (g) are assumed to be identically distributed over time. The uncorrelated factor returns can be recovered with the transformation

$$g = \Sigma^{-1/2} f \qquad (5)$$

We estimate Σ using an exponential weighted moving average (EWMA) of trailing factor returns. Forecast scenarios $\tilde{f}_T$ on analysis date T are given by

$$\tilde{f}_T = \Sigma_T^{1/2} \Sigma_t^{-1/2} f_t \qquad (6)$$

where t is a historical date preceding T. These *covariance-normalized factor returns* are the inputs to shortfall optimization. Further details are given in Dubikovsky, et al (2010).

Appendix B: Estimation Error and Derivation of the Weight-Error Angle

We start by mathematically encapsulating the range of all feasible weight vectors **w** that satisfies the given optimization constraints. Consider the two-asset portfolio setting of Figure 1, where the set of all possible weights lies on the intersection of the dashed line with the first quadrant. We define the *feasibility range* in this two-dimensional setting to be the length of this line, which is √2. In three dimensions, feasible weights lie on the area of an equilateral triangle with edge length equal to √2; in four dimensions, feasibility is the volume of a tetrahedron with edge length equal to √2. In general, given n+1 assets and their weights, the feasibility range is the n-dimensional volume of the n-dimensional simplex with equal-sided edges equal to √2. This volume is equal to

$$V_{feasible} = \frac{\sqrt{n+1}}{n!} \qquad (7)$$

Since the assumption is that optimal weights are given by equal weights, the optimal weight vector $\mathbf{w}_{op}$ for n+1 assets has coordinates $\mathbf{w}_{op}$ = [1/(n+1), …, 1/(n+1)], and its length is its norm |**w**|.

Finding the boundary angle β around the optimal weight vector (beyond which one would be better off guessing a random set of weights rather than computing the optimal weights) reduces to finding the n-dimensional object with sub-volume $V_{boundary}$ that has half the volume of the total feasibility range (i.e., $V_{boundary}$ = ½ $V_{feasible}$). Since we seek an angle rotating around the optimal vector $\mathbf{w}_{op}$, the object with volume $V_{boundary}$ must be an n-dimensional hypersphere centered at the coordinates of $\mathbf{w}_{op}$. Its radius is unknown, but can be backed out from its volume:

For n even, $r = \left[\frac{\sqrt{n+1}}{2(2\pi)^{n/2}(n!!!)}\right]^{1/n}$ and for n odd, $r = \left[\frac{\sqrt{n+1}}{4(2\pi)^{n-1/2}(n!!)}\right]^{1/n}$,

where n!!! is the product of all positive odd integers less than or equal to n, and n!! is the product of all positive even integers less than or equal to n. With some higher-dimensional imagination, one may now be able to see that the boundary angle β is adjacent to the optimal vector $\mathbf{w}_{op}$ and has an opposite edge given by the radius r of the n-dimensional hypersphere (with the right angle being between the vector and the radius). Therefore, the boundary angle β is given by

$$\beta = \mathrm{atan}(r/|w_{op}|) \qquad (8)$$

and can be generated for any number of assets n+1 (dimension n) since it is dependent only on n.

Appendix C: Variance and Shortfall Optimization

In this appendix we review the standard formulation of mean-variance optimization as a quadratic program (QP) and mean-shortfall optimization as a linear program (LP). Both QPs and LPs can be solved using standard optimization algorithms.

Variance Optimization. Given a vector of weights w, covariance matrix Σ, vector of expected returns α, and risk aversion parameter λ, the mean-variance optimization problem is:

$$\max_w \; w'\alpha - \lambda w'\Sigma w \tag{9}$$

subject to any set of linear equality or inequality constraints (long-only, full investment, etc.).

Shortfall Optimization. Given vectors $r_1,\ldots,r_T$ of forecast return scenarios, weight vector w, empirical shortfall estimator $s_p(w) = -\frac{1}{K}\sum_{i=1}^{K} w'r_{(i)}$ for confidence level p, $K = \lfloor T(1-p) \rfloor$, and risk aversion parameter Λ, we seek weights w minimizing $s_p(w)$:

$$\max_w \; w'\alpha - \Lambda s_p(w) \tag{10}$$

The shortfall estimator $s_p(w)$ is an average of order statistics, and it is not obvious how to solve the shortfall optimization. However, Rockafellar and Uryasev (2000) and others have shown how to formulate this as an equivalent linear program with T+1 additional variables and 2T additional constraints:

$$\begin{aligned} &\max_{w,z,t} w'\alpha + \Lambda\left(t - \frac{1}{K}\sum_{i=1}^{T} z_i \right) \\ &\text{s.t.} \qquad z \geq 0, \quad z_i > t - w'r_i \end{aligned} \tag{11}$$

subject to any set of linear equality or inequality constraints (long-only, full investment, etc.). In the next section (Appendix D), we sketch how to convert the optimization problem (8) into its LP equivalent formulation.

Variance Shortfall Optimization. Clearly the variance and shortfall terms can be combined into a single objective function to give another standard quadratic program:

$$\begin{aligned} &\max_{w,z,t} \ w'\alpha + \Lambda\left(t - \frac{1}{K}\sum_{i=1}^{T} z_i\right) - \lambda w'\Sigma w \\ &\text{s.t.} \qquad z \geq 0, \quad z_i > t - w'r_i \end{aligned} \tag{12}$$

subject to any set of linear equality or inequality constraints (long-only, full investment, etc.).

Appendix D: Linearization of Shortfall Optimization

In this appendix, we begin by considering this shortfall minimization problem:

$$\min_w \quad s_p(w) \tag{13}$$

Let $\mathbf{r}_1,\ldots,\ \mathbf{r}_T$ be T vectors of forecast return scenarios for N assets, and $\mathbf{w}$ be the weight vector. The portfolio return at time t is given by $\mathbf{p}_t = \mathbf{w'r}_t$. If the sorted portfolio returns are written in increasing order as $(\mathbf{w'r})_{(1)} \leq (\mathbf{w'r})_{(2)} \leq \cdots \leq (\mathbf{w'r})_{(T)}$, the empirical shortfall estimator is

$$s_p = -\frac{1}{K}\sum_{i=1}^{K} w'r_{(i)} \tag{14}$$

where $K = \lfloor T(1-p) \rfloor$. The linearization of this optimization problem is based on two crucial observations that convert the order statistic into a linear sum that is subject to linear constraints.

Observation I. The sum of the t smallest portfolio returns is always smaller than or equal to the sum of any other combination of t returns. Formally, for any t < T, we have

$$\sum_{i=1}^{t} w'r_{(i)} \leq \sum_{i\in S} w'r_i \tag{15}$$

where we are indexing on the right hand side over all possible sets S that contain t portfolio returns (i.e., |S|=t, for t=1,…, T). The shortfall optimization problem (13) can therefore be rewritten as

$$\begin{aligned} &\max_{w,p} \quad -\frac{1}{K}\sum_{i=1}^{K} p_{(i)} \\ &\text{s.t.} \qquad \sum_{i=1}^{t} p_{(i)} \leq \sum_{i\in S} w'r_i \end{aligned} \tag{16}$$

Observation II. The sum of the K smallest portfolio returns $\sum_{i=1}^{K} p_{(i)}$ as it appears in the objective function of (16) is in fact the value of the linear optimization problem (17)

$$\begin{aligned} &\min_x \ \sum_{i=1}^{T} x_i p_i \\ &\text{s.t.} \quad \sum_{i=1}^{T} x_i = K, \quad 0 \le x_i \le 1 \end{aligned} \tag{17}$$

This can be proven formally by induction on the number K. However, to help understand why this is true, notice that for K=1 we get the minimum value of combination of portfolio returns if we assign all the weight to the smallest return $p_{(1)}$, since any other fraction of weight assigned to a larger return will yield a larger total value. Similarly, for K=2 we get the minimum combination of portfolio returns if we assign all weights to the smallest returns. Since each x_i cannot be larger than 1, we pick the two smallest returns. By strong duality, optimization problem (17) is equivalent to

$$\begin{aligned} &\max_{t,z} \ Kt + \sum_{i=1}^{T} z_i \\ &\text{s.t.} \quad t + z_i \le p_i, \quad z_i \le 0, \quad i = 1,\dots,T \end{aligned} \tag{18}$$

We use this observation to rewrite optimization (16) as

$$\begin{aligned} &\min_w \quad -\frac{1}{K} \max_{t,z} \ Kt + \sum_{i=1}^{T} z_i \\ &\text{s.t.} \quad t + z_i \le w' r_i, \quad z_i \le 0, \quad i = 1,\dots,T \end{aligned} \tag{19}$$

Using the fact that max(f) = -min(-f), we finally convert optimization problem (18) into the following linear optimization, which is equivalent to (13):

$$\begin{aligned} &\min_{w,t,z} \quad -t - \frac{1}{K} \sum_{i=1}^{T} z_i \\ &\text{s.t.} \quad t + z_i \le w' r_i, \quad z_i \le 0, \quad i = 1,\dots,T \end{aligned} \tag{20}$$

Appendix E: Absolute Cumulative Returns

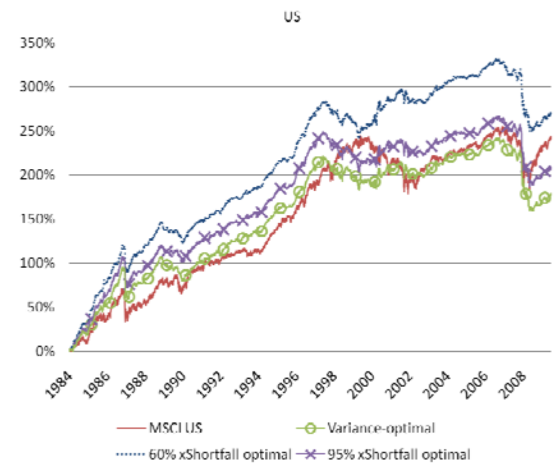

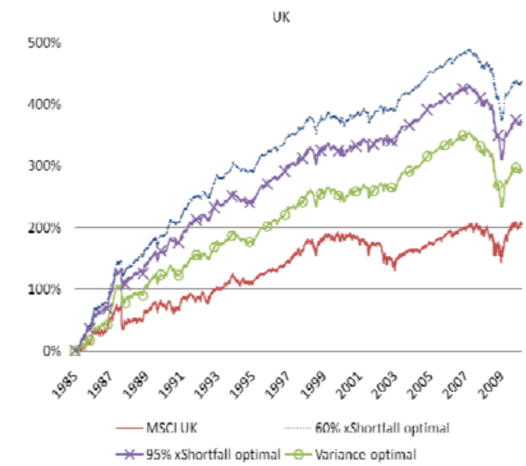

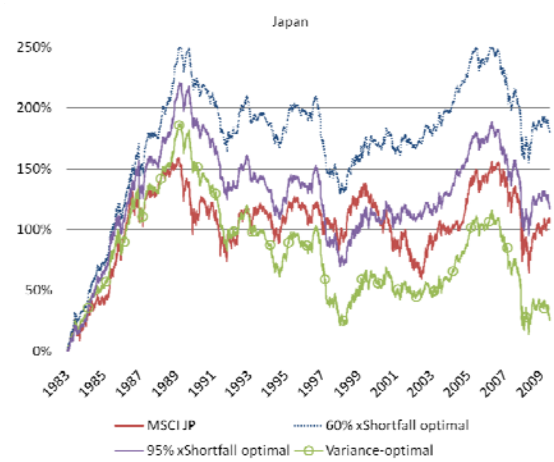

References

Carlo Acerbi and Dirk Tasche. “Expected Shortfall: A Natural Coherent Alternative to Value at Risk.” Economic Notes, 31: 379–388. doi: 10.1111/1468-0300.00091, 2002.

Angelo Barbieri, Vladislav Dubikovsky, Alexei Gladkevich, Lisa R. Goldberg and Michael Y. Hayes. “Central Limits and Financial Risk.” Quantitative Finance, 10(10):1091–1097, 2010.

Jennifer Bender, Jyh-Huei Lee and Dan Stefek. “Constraining Shortfall.” MSCI Barra Research Insight, April, 2010.

D. Bertsimas, G.J. Lauprete and A. Samarov. “Shortfall as a Risk Measure: Properties, Optimization and Applications.” Journal of Economic Dynamics and Control, 28:1353–1381, 2004.

Vladislav Dubikovsky, Michael Y. Hayes, Lisa R. Goldberg and Ming Liu. “Forecasting Extreme Risk of Equity Portfolios with Fundamental Factors.” In Arthur Berd, editor, *Lessons from the Credit Crisis*, Chapter 13. Risk Books, 2010.

Lisa R. Goldberg and Michael Y. Hayes. “The Long View of Financial Risk.” Journal of Investment Management, 8(1):39–48, 2010.

Lisa R. Goldberg, Michael Y. Hayes, Jose Menchero and Indrajit Mitra. “Extreme Risk Analysis.” Journal of Performance Measurement, Spring:17–30, 2010.

Imre Kondor, Szilárd Pafka and Gábor Nagy. “Noise sensitivity of portfolio selection under various risk measures.” Journal of Banking and Finance, 31:1545–1573, 2007.

P. Krokhmal, J. Palmquist and S. Uryasev. “Portfolio optimization with conditional value-at-risk objective and constraints.” Journal of Risk, 4(2):11–27, 2002.

Harry Markowitz. “Portfolio Selection.” Journal of Finance, VII(1):77–91, 1952.

R. Tyrrell Rockafellar and Stanislav Uryasev. “Optimization of conditional value-at-risk.” Journal of Risk, 2:493–517, 2000.

R. Tyrrell Rockafellar and Stanislav Uryasev. “Conditional Value-at-Risk for General Loss Distributions.” Journal of Banking and Finance, 26:1443–1471, 2002.