Downside Risk Management in Emerging Markets

By Issam S. Strub, PhD, and Edward D. Baker

The recent financial crisis and the corresponding market crash have had a tremendous impact on investors with significant exposures to equity markets but without appropriate risk control tools. Indeed, within just a few weeks a number of pension funds, university endowments, and mutual funds suffered catastrophic capital losses, in some cases registering a drop in net asset value in excess of 50 percent. This has renewed investment professionals' interest in techniques for downside risk mitigation, which requires both accurate risk measurement tools and effective risk management strategies.

In this article we develop and evaluate a number of such tools and strategies in the specific context of an emerging markets equity long-only product. While the tools developed in this article can be applied to other asset classes, we chose to focus on emerging markets due to their growing importance among global equity markets (Lee 2010). For example, the fraction of global equity markets capitalization originating from emerging countries as measured by the MSCI All Country World Index (ACWI) increased from 4 percent in 2003 to almost 13 percent in 2010. Such a change has led to substantial inflows to emerging market funds ($83.2 billion in 2009 and $68.5 billion through the first ten months of 2010 according to EPFR Global.[1])

Additionally, risk management in emerging markets equities is of particular interest because these markets tend to be highly speculative and volatile in nature; stock prices can change suddenly due to a number of fundamental and technical factors as well as changes in risk aversion. This often is compounded by liquidity issues and clustering, increasing the importance of risk-control algorithms. We start by presenting tools available to measure downside risk; then we develop risk-adjusted strategies and apply them to asset allocation between emerging markets equities and cash and at a later stage between emerging markets equities and U.S. bonds. This allows us to evaluate the effectiveness of each technique in different settings. We demonstrate that it is possible to significantly reduce both volatility and maximum drawdown without a notable decrease in returns by adjusting the allocation to equities according to risk levels. In particular, when capital is allocated between equities and bonds rather than simply modifying the cash level, risk-adjusted return (measured by the Sharpe ratio), more than doubles.

Downside Risk Measurement

Expected Shortfall

A commonly used measure of downside risk is Value at Risk (VaR), which is defined as the minimum loss a portfolio of assets will experience over a given time horizon with a given probability. VaR can be computed easily from the historical daily return distribution by ordering the daily returns and selecting the quantile corresponding to the confidence level chosen (for example, 95 percent). Unfortunately, VaR is concerned only with the number of losses that exceed the VaR confidence level and not the size of these losses; obviously, investment professionals would prefer a risk measure that accounts for the magnitude of large losses as well the probability of occurrence.

To obtain a more complete measure of large losses, one should examine the entire shape of the tail of the distribution of losses beyond the VaR, which is called expected shortfall (ES) (Christoffersen 2003, McNeil et al. 2005), also referred to as Conditional Value at Risk (CVaR) or Tail VaR.

Put in mathematical terms, the ES for a daily return distribution F at a given confidence level α is given by

$$\mathrm{ES}_\alpha = -\mathrm{E}\{X|X \leq -\mathrm{VaR}_\alpha\} \tag{1}$$

where

$$\mathrm{VaR}_\alpha = -F^{-1}(1-\alpha). \tag{2}$$

Historical Volatility-Based Expected Shortfall

While ES presents a number of interesting properties, such as being a coherent risk measure and a convex function of the portfolio weights (unlike VaR), which makes it suitable for portfolio optimization, its computation requires an explicit expression of the portfolio return distribution function F, which is usually unknown in practice. To circumvent this problem, one may fit a given distribution to the left tail of historical daily returns and then use this distribution to compute the ES. For a normal distribution, computing ES is fairly simple once the standard deviation (or volatility) σ of the return distribution is known. For example, at the 95-percent confidence level, standard deviation, VaR, and ES are related by

$$\mathrm{VaR} \simeq 1.645 \times \sigma \text{ and } \mathrm{ES} \simeq 1.26 \times \mathrm{VaR} \tag{3}$$

leading to

$$\mathrm{ES} \simeq 2.07 \times \sigma\,. \tag{4}$$

Extreme Value Theory-Based Expected Shortfall

The generalized Pareto distribution (GPD) also can be used to model downside risk and compute ES. GPD originated from extreme value theory (EVT) (Balkema and de Haan 1974, Pickands 1975), a branch of statistics dedicated to modeling extreme events. EVT's central result states that the extreme tail of a wide range of distributions can be approximately described by the GPD $G_{\zeta,\beta}$ given by

$$G_{\zeta,\beta}(y) = \begin{cases} 1 - \left(1 + \frac{\zeta y}{\beta}\right)^{-\frac{1}{\zeta}}, \mathit{for}\ \zeta \neq 0 \\ 1 - \exp\left(-\frac{y}{\beta}\right), \mathit{for}\ \zeta = 0 \end{cases} \tag{5}$$

where $\beta > 0$, and the support of $G_{\zeta,\beta}$ is $y \geq 0$ when $\zeta \geq 0$ and $0 \leq y \leq -\frac{\beta}{\zeta}$ when $\zeta < 0$.

The shape and scale parameters ζ and β can be estimated using maximum likelihood estimation (MLE) by fitting a GPD distribution to the tail of the return distribution after a given threshold u. Once this is done, the expected shortfall can be computed by

$$ES_\alpha = \frac{VaR_\alpha + \beta - \zeta u}{1 - \zeta} \tag{6}$$

where the VaR for a GPD can be estimated by

$$VaR_\alpha = u + \frac{\beta}{\zeta}\left(\left(\frac{\alpha N}{N_u}\right)^{-\zeta} - 1\right) \tag{7}$$

with N equal to the total number of observations and N_u equal to the number of observations exceeding the threshold u.

The preceding result requires that observations be independent and identically distributed, which often is not the case for daily returns because they present some level of autocorrelation. Therefore, we start by filtering the daily returns and then applying EVT to the standardized residuals (see McNeil and Frey 2000, Nystrom and Skoglund 2005), then fit a GPD to the tails through maximum likelihood estimation. Once this is done, we obtain the shape and scale parameters and replace these values in equation 6 to compute the ES at the level of confidence desired. EVT has been used during the past decade for risk management in finance, and there has been a notable increase in the number of publications on the subject since the recent financial crisis (e.g., Cascon and Shadwick 2009; Ghorbel and Trabelsi 2008, 2009; Goldberg et al. 2008, 2009; McNeil and Frey 2000; McNeil et al. 2005; and Nystrom and Skoglund 2005).

Practical Implementation

We illustrate the difference between the volatility/normal distribution and EVT approaches to ES computation by applying the two algorithms to daily historical returns from a stock (Mongolia Energy, ticker code 276:HK) for which the downside risk obviously is not modeled accurately by a normal distribution. We consider 1,000 daily returns for this stock and fit both a normal distribution and a GPD to the left tail of the daily returns. The results are presented in figure 1. We can see that while both techniques yield similar VaR values at the 95-percent confidence level (in this case 7.8 percent), the 95-percent ES, which can be visually identified as the area under a given distribution curve left of the 95-percent VaR threshold, is significantly higher when computed using the GPD than when using volatility and a normal distribution assumption. Note that this is a pathological case selected as an example because the difference between the two methods for ES computation is particularly pronounced for this stock. In the following, we study an equity index containing several thousand stocks for which the distinction is much less obvious because averaging among a large number of stocks tends to reduce this phenomenon.

FIGURE 1: COMPARISON OF GENERALIZED PARETO AND NORMAL DISTRIBUTION

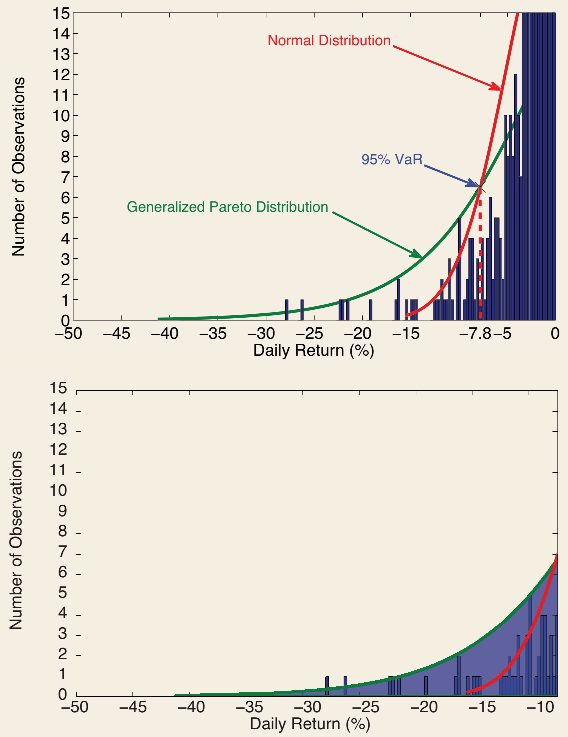

Top: Comparison of generalized Pareto and normal distribution. Note that the GPD models the left tail of the daily returns much more accurately than the normal distribution.

Bottom: 95-percent ES for each distribution. The ES is represented by the shaded area under the green (GPD) or red (normal distribution) curves. Here it is apparent that the ES computed using a normal distribution underestimates the downside risk when compared with a GPD.

Downside Risk Management

Asset Allocation between Equities and Cash

We consider a long-only portfolio of emerging market equities and assume that the portfolio manager can adjust the asset allocation between equities and cash. The emerging market equities are represented by the MSCI Emerging Markets Investable Market Index (EM IMI) USD Net Index (ticker code MIMUEMRN) and the equity allocation is computed using

$$\text{Equity Alloc.} = \max\left(\min\left(\frac{\text{Max. ES}}{\text{Current ES}}, \text{Max. Equity}\right), 1 - \text{Max. Cash}\right) \quad (8)$$

where Max. ES is the chosen maximum daily ES threshold; Max. Equity is the maximum equity allocation; and Max. Cash is the maximum cash allocation.

The derivation of this formula is straightforward: We start by computing the ratio of the maximum acceptable ES over the current ES and bound this value from above by the maximum equity allocation (typically set at 1 if no leverage is allowed); this number is then bounded from below by 1 minus the maximum cash level to obtain the final equity allocation.

In the following section, the equity allocation is adjusted with the objective of a maximum daily ES at the 95-percent confidence level of 2 percent (Max. ES = 2 percent in equation 8). Of course, any other value can be chosen as the ES threshold; the current value of 2 percent is selected as a middle ground between more conservative or aggressive risk levels. The Current ES is computed from historical returns using either the volatility-based method (equation 4) or EVT (equation 6). While daily returns are used for this computation, the process is done on a weekly basis and the cash level is adjusted at the beginning of each week using the ES computed at the end of the previous week; a weekly rather than daily adjustment frequency is selected both for the sake of realism because changes in equity holdings from one day to the next are limited by liquidity constraints, and also because daily portfolio rebalancing would incur significant transaction costs. For the volatility-based computation we use thirty-day historical volatility; we take 500 days of historical returns into account for the EVT-based computation.

We study the daily MSCI Emerging Markets IMI USD Net Index prices from January 2000 until the end of August 2010 and normalize the values of each portfolio to a base of 100 at the beginning of the period considered. We examine the case in which the cash level can vary between 0 percent and 100 percent (Max. Cash = 1) as well as the case in which the cash level is constrained between 0 percent and 20 percent (Max. Cash = 0.2), a typical situation in long-only equity products. In both cases the maximum equity allocation is set at 100 percent (Max. Equity = 1) so that no leverage is employed, again a common feature of long-only equity funds. The results are summarized in table 1; the historical performance of each portfolio is presented in figure 2, and the historical evolution of the equity allocation level is plotted in figure 3.

e notice by examining the data in table 1 that the two risk-adjusted portfolios have both a significantly lower volatility and maximum drawdown than the MSCI EM IMI USD Net index; the annualized volatility is reduced from an original value of 20.96 percent down to 14.71 percent for the volatility-adjusted portfolio and 12.93 percent for the EVT-adjusted portfolio. Similarly, the maximum drawdown, which for the period we consider occurred during the 2008 financial crisis, is 65.44 percent for the index compared to 51.54 percent and 49.79 percent, respectively, for the volatility and EVT-adjusted portfolios. While this is not entirely surprising given that we allow the equity allocation to vary between 0 percent and 100 percent, the annualized returns obtained by the risk-adjusted portfolios are not reduced by the same factor as the volatility or maximum drawdown. Annualized return for the index is 8.59 percent while it is 9.05 percent for the volatility-adjusted portfolio and 7.14 percent for the EVT-adjusted portfolio. Note that the volatility-adjusted portfolio has a higher return than the index, which is remarkable considering that the volatility has been reduced by almost 30 percent and the maximum drawdown is lower by more than 20 percent. This improved risk-adjusted return for the two risk-adjusted portfolios can be seen in the Sharpe ratios that are 0.62 and 0.55, respectively, for the volatility and EVT-adjusted portfolios compared to 0.41 for the index. We can therefore see the substantial benefits coming from a simple strategy that limits exposure to equity markets in times of heightened downside risk, measured in terms of ES.

The analysis of the equity allocation evolution in figure 3 shows that the EVT-adjusted allocation changes more smoothly than the volatility-adjusted allocation, but it also shows that the return to higher equity allocation levels after the 2008 crisis is slower when EVT is used. This should not come as a surprise because the EVT algorithm used to compute the ES requires 500 days of historical returns and as such is less responsive than the volatility algorithm that is based on only the previous thirty days. The 500-day requirement for the

TABLE 1: PERFORMANCE FOR THE MSCI EM IMI USD NET INDEX AND THE DIFFERENT RISK-ADJUSTED PORTFOLIOS

	Annual Return	Annual Volatility	Maximum Drawdown	Sharpe Ratio
MSCI EM IMI USD Net	8.59%	20.96%	65.44%	0.41
Volatility-adjusted portfolio	9.05%	14.71%	51.54%	0.62
EVT-adjusted portfolio	7.14%	12.93%	49.79%	0.55
Volatility-adjusted portfolio (max cash 20%)	8.81%	18.01%	57.91%	0.49
EVT-adjusted portfolio (max cash 20%)	7.57%	17.51%	56.84%	0.43

FIGURE 2: PERFORMANCE OF THE RISK-ADJUSTED PORTFOLIOS

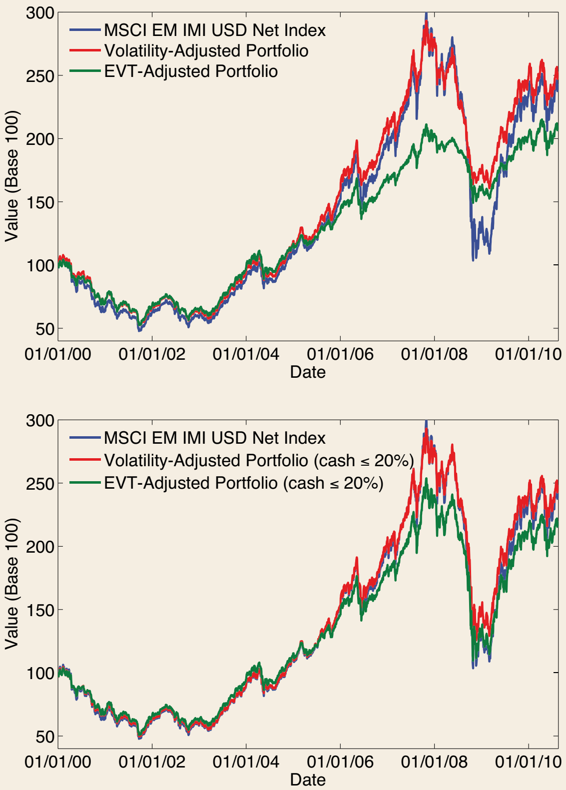

Top: Historical performance of the MSCI EM IMI USD Net Index and the two risk-adjusted EM portfolios.

Bottom: Historical performance of the MSCI EM IMI USD Net Index and the two risk-adjusted portfolios in the case of a constrained cash level (20-percent cash).

FIGURE 3: EVOLUTION OF EQUITY ALLOCATIONS FOR THE RISK-ADJUSTED PORTFOLIOS

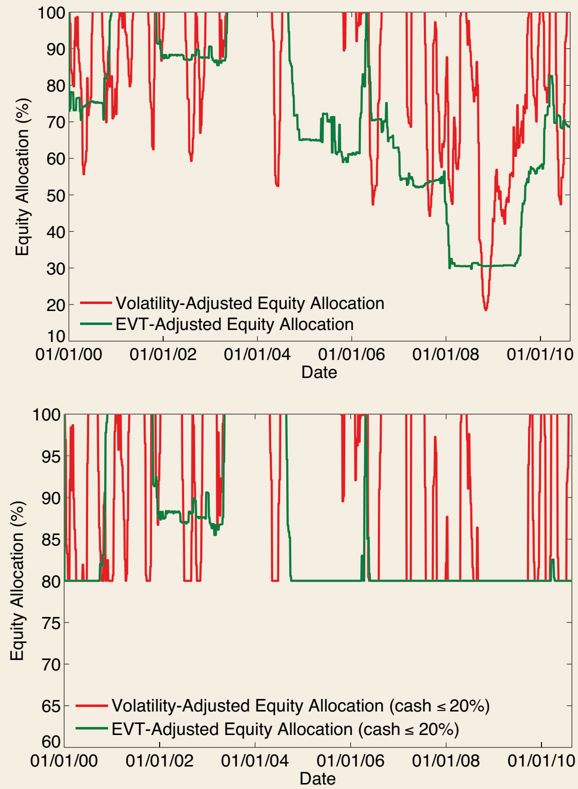

Top: Historical evolution of the equity allocation level for each risk-adjusted portfolio.

Bottom: Historical evolution of the equity allocation level for each risk-adjusted portfolio in the case of a constrained cash level (20-percent cash).

EVT algorithm comes from the need to have enough data in the left tail of the return distribution beyond the confidence level selected (95 percent in our case) to be able to fit a GPD. For the volatility algorithm, one also could use a longer (e.g., sixty or ninety days) or shorter (e.g., ten days) time period from which historical volatility is extracted, depending on the desired sensitivity of the algorithm to market changes because longer (shorter) time periods will result in a smoother (more abrupt) allocation level evolution.

Note also that the downside protection comes with an occasional significant decrease in equity allocation. For example, during the 2008 crisis, the equity allocation fell to 18.5 percent and 29.7 percent, respectively, for the volatility and EVT-based strategies.

The aforementioned situation in which cash levels can exceed 70 percent rarely is seen in practice because portfolio managers usually are required by their mandates to keep the cash level under a certain limit at all times. We now examine the performance of both risk-adjusted portfolios when the cash level is constrained between 0 percent and 20 percent. As can be seen from the historical performance of the portfolios presented in figure 2, the difference between the two portfolios and the index is much smaller than in the unconstrained case. The annualized return for the volatility-based strategy is slightly higher than for the index, while the annualized return of the EVT-based strategy is lower than the index by about 100 basis points; the volatilities of both strategies are lower than the volatility of the index and the Sharpe ratios are moderately improved at 0.49 and 0.43, respectively, from 0.41 for the index. The improvement in terms of maximum drawdown, while not as significant as in the unconstrained case, still is noticeable with a maximum drawdown of 57.91 percent and 56.84 percent, respectively, which is lower than the 65.44-percent drawdown registered by the index. One conclusion of this comparison is that by limiting a portfolio manager's ability to raise the cash level above a certain threshold, investors also will hamper their potential protection against downside risk; as such, downside risk management techniques are somewhat ineffective as well as impractical when such constraints are imposed.

Asset Allocation between Equities and Bonds

We now consider a situation in which the portfolio manager can adjust the asset allocation between equities and fixed income instruments. The equity market will be represented, as earlier, by the MSCI EM IMI USD Net index and the bond market by the J.P. Morgan U.S. Aggregate Bond index (ticker code JGAGUSUS). This index was preferred to an emerging market bond index because its correlation with the MCSI EM IMI USD Net is 0.75 compared to 0.83 for the J.P. Morgan Emerging Markets Bond Index, meaning that the U.S. Bond index provides greater diversification. We use as a benchmark an equal-weighted composite of the MSCI EM IMI USD Net and the J.P. Morgan U.S. Aggregate Bond indexes and compute the asset allocation to bonds or stocks weekly using a similar algorithm as before; the 95-percent ES for each index is computed on a historical basis using either the volatility or EVT-based method and then the allocation is adjusted to a 95-percent ES of 2.0 percent for equities and 0.5 percent for bonds. We select a lower ES target for bonds because they typically experience a significantly lower volatility than equities; therefore, if we selected the same ES target for bonds and equities, the result would be a very limited allocation to equities at all times (at most 20 percent in the present case). This generates equity and bond risk factors given by

$$\text{Equity Risk Factor} = \frac{\text{Max. Equity Expected Shortfall}}{\text{Current Equity Expected Shortfall}} \quad (9)$$

and

$$\text{Bond Risk Factor} = \frac{\text{Max. Bond Expected Shortfall}}{\text{Current Bond Expected Shortfall}} \quad (10)$$

which then are normalized to obtain the following equity and bond asset allocations:

$$\text{Equity Allocation} = \frac{\text{Equity Risk Factor}}{\text{Equity Risk Factor} + \text{Bond Risk Factor}} \quad (11)$$

$$\text{Bond Allocation} = \frac{\text{Bond Risk Factor}}{\text{Equity Risk Factor} + \text{Bond Risk Factor}} \quad (12)$$

Historical performance of the equity index, equity/bond composite, and two portfolios is plotted in figure 4; evolution of the equity asset allocation for each portfolio is presented in figure 5. Table 2 shows the corresponding return and risk numbers. The first apparent aspect is that the volatility and maximum drawdown are noticeably lower for the equity/bond composite and for both risk-adjusted portfolios compared to the equity index. Volatility is divided by a factor of two or more from the 20.96 percent of the MSCI EM IMI USD Net to 10.47 percent, 8.94 percent, and 8.44 percent for the equity/bond composite and the volatility and EVT-based strategies, respectively. Similarly the maximum drawdown falls from 65.44 percent for the equity index to 39.76 percent for the equity/bond composite, 34.55 percent for the volatility-based strategy, and 28.15 percent for the EVT-based strategy. As in the equity and cash case discussed earlier, the EVT-adjusted portfolio presents the largest reduction in downside risk measured both in terms of volatility and maximum drawdown. While both risk-adjusted strategies offer significant downside risk protection, the annualized return of the volatility-adjusted portfolio at 8.84 percent is higher than the returns of both the pure equity and equity/bond benchmarks; for the EVT-based strategy, the annualized return is lower at 7.35 percent, which is 100 basis points lower than the equity/bond composite. This is due to a lower equity allocation obtained when using EVT to compute ES (figure 5). In terms of risk-adjusted performance, the risk-adjusted portfolios outperform the equity/bond equal-weighted composite with Sharpe ratios of 0.99 and 0.80, respectively, for the volatility and EVT-based strategies compared to 0.80 for the composite. The Sharpe ratios of the risk-adjusted portfolios are more than double the 0.41 value registered by the MSCI EM IMI USD Net index. This is due to the combination of downside risk protection and diversification obtained by using ES to allocate between equities and bonds.

FIGURE 4: HISTORICAL PERFORMANCE OF EQUAL-WEIGHTED EM EQUITY/U.S. BOND COMPOSITE INDEX, MSCI EM IMI USD NET INDEX, AND TWO RISK-ADJUSTED PORTFOLIOS

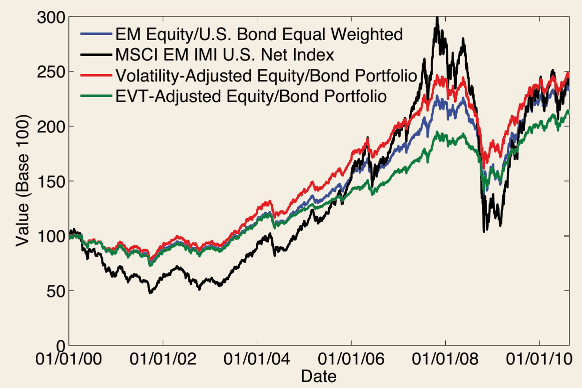

Figure 5 shows the proportion invested in equities across time for each portfolio. Once again, the historical evolution of the equity allocation for the EVT-based strategy is much smoother than for the volatility-based strategy. Note that the EVT-adjusted portfolio has a significantly lower equity allocation than its volatility-adjusted counterpart in 2008 and 2009; this explains both the lower drawdown experienced during the financial crisis and the lower performance during the subsequent market rebound.

Conclusion

In this article we developed and compared two algorithms for downside risk management and applied them to an emerging markets equities long-only fund. First, we considered a situation in which a portfolio manager could adjust the asset allocation between an emerging markets equity index and cash on a weekly basis both with and without constraints on the

TABLE 2: PERFORMANCE FOR THE MSCI EM IMI USD NET INDEX, EQUAL-WEIGHTED EQUITY/BOND COMPOSITE, AND TWO RISK-ADJUSTED PORTFOLIOS

	Annual Return	Annual Volatility	Maximum Drawdown	Sharpe Ratio
MSCI EM IMI USD Net	8.59%	20.96%	65.44%	0.41
EM Equity/U.S. Bond Equal Weighted	8.35%	10.47%	39.76%	0.80
Volatility-Adjusted Equity/Bond Portfolio	8.84%	8.94%	34.55%	0.99
EVT-Adjusted Equity/Bond Portfolio	7.35%	8.44%	28.15%	0.87

FIGURE 5: HISTORICAL EVOLUTION OF EQUITY ALLOCATION LEVEL FOR EACH RISK-ADJUSTED PORTFOLIO

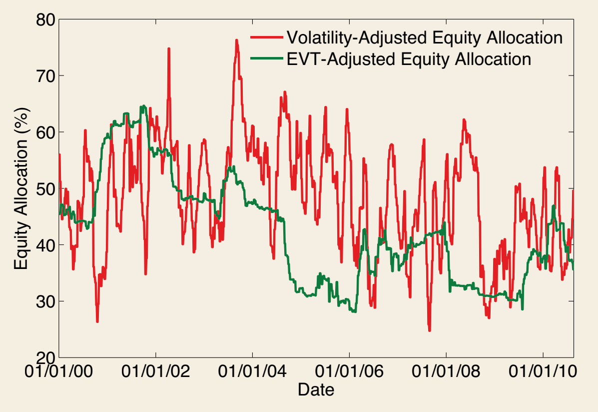

maximum cash level. We showed that a significant downside risk reduction could be obtained from both a volatility-based algorithm and an EVT-based algorithm when the cash level is unrestricted. This is attained by having a cash level at or above 70 percent in extreme market conditions such as the 2008 financial crisis. The benefits are apparent both in terms of volatility and maximum drawdown reduction as well as in terms of risk-adjusted return. Indeed, the Sharpe ratios registered by these portfolios are up to 50 percent higher than the original index Sharpe ratio. However, most long-only equity funds will have restrictions on the maximum cash level allowable; 20 percent is a typical upper bound. In this situation, our study demonstrated that while an improvement still was obtained with respect to the index, it was much less significant than in the unconstrained situation; and the risk-adjusted portfolios in the constrained case both experienced maximum drawdowns of more than 55 percent. Thus, by limiting the portfolio manager's ability to increase the cash level, investors are in effect depriving themselves of downside risk protection; contrary to a common opinion, holding significant levels of cash in a portfolio will not always have a negative impact on performance as long as this allocation is adjusted judiciously. For example, our volatility-adjusted portfolio had an annualized return during 2000–2010 that was higher than the index's return while the maximum drawdown and volatility were noticeably reduced.

We then considered a mixed equity/bond portfolio for which the asset allocation was determined by computing ES using both volatility and EVT. The resulting portfolios had volatilities less than half that of the equity benchmark, and maximum drawdowns also were reduced in the same proportions without a significant decrease in annualized returns. As a result, Sharpe ratios more than doubled. This shows that using bonds is more effective than simply increasing cash levels in terms of downside risk protection.

This study also concludes that while using volatility rather than EVT might lead to underestimating the ES, in particular for certain stocks that present unusually fat left tails, this is less of an issue when considering a broad index containing thousands of securities due to the averaging effect. Volatility-adjusted portfolios consistently outperformed their EVT-based counterparts both in terms of annualized returns and Sharpe ratio. This is mostly due to the increased responsiveness when using volatility as a risk measure compared to EVT, which requires using a significant number of daily returns (500 days in our case) in order to generate meaningful results. When it comes to reducing maximum drawdown, however, the EVT-based algorithm performs better than the volatility-based method. For example, during the 2008 financial crisis, the EVT-based portfolios suffered smaller losses than the volatility-adjusted portfolios, which is not surprising given that EVT's objective is to model extreme losses such as those observed during market crashes. That said, by definition these extreme events are relatively rare and using EVT during the ten-year period we studied would result in lower equity exposure levels than for the volatility-based strategy and a certain underperformance, particularly during the market rebounds occurring immediately after a crash. One could envision a situation in which volatility is used most of the time to monitor downside risk with a switch to EVT when risk aversion reaches high levels, which could be identified by volatility exceeding a certain threshold; once volatility has decreased below this threshold, the portfolio manager would return to the volatility-based approach.

In any case, risk adjusting a portfolio presents obvious advantages compared to a more passive risk management approach. Note also that while we considered an equity index for this study, in practice the portfolio manager can adjust his equity allocation and also use judicious stock selection to move to more-defensive positions when market conditions require it. This would result in an even greater outperformance with respect to the benchmark.

Issam S. Strub, PhD, is a research scientist with The Cambridge Strategy in London, UK. Contact him at issam.strub@thecambridgestrategy.com.

Edward D. Baker is executive chairman of The Cambridge Strategy in London, UK. Contact him at edward.baker@thecambridgestrategy.com.

Endnote

1 EPFR Fund Flows (http://www.epfr.com/fundflows.aspx).

References

Balkema, A., and L. de Haan. 1974. Residual Lifetime at Great Age. *Annals of Probability* 2, no. 5: 792–804.

Cascon, A., and W. F. Shadwick. 2009. A New Approach to Tail Risk. *Journal of Investment Consulting* 10, no. 1: 33–48.

Christoffersen, P. F. 2003. *Elements of Financial Risk Management* (San Diego: Academic Press).

Ghorbel, A., and A. Trabelsi. 2008. Predictive performance of conditional Extreme Value Theory in Value-at-Risk Estimation. *International Journal of Monetary Economics and Finance* 1, no. 2: 121–148.

———. 2009. Measure of financial risk using conditional extreme value copulas with EVT margins. *Journal of Risk 11*, no. 4 (summer): 51–85.

Goldberg, L. R., G. Miller, and J. Weinstein. 2008. Beyond value-at-risk: forecasting portfolio loss at multiple horizons. *Journal of Investment Management* 6, no. 2: 73–98.

Goldberg, L. R., M. Y. Hayes, J. Menchero, and I. Mitra. 2009. Extreme Risk Analysis. MSCI Barra Working Paper. Available at http://www.msci.com/resources/research_papers/extreme_risk_analysis_1.html.

Lee, J. 2010. Changing Spaces: Emerging Markets Investing. AQR Capital Management. Available at http://www.aqr.com/Download.aspx?link_id=1.

McNeil A. J., and R. Frey. 2000. Estimation of tail-related risk measures for heteroscedastic financial time series: an extreme value approach. *Journal of Empirical Finance* 7, no. 3–4 (November): 271–300.

McNeil A. J., R. Frey, and P. Embrechts. 2005. *Quantitative Risk Management: Concepts, Techniques, and Tools* (Princeton, NJ: Princeton University Press).

Nystrom, K., and J. Skoglund. 2005. Efficient Filtering of Financial Time Series and Extreme Value Theory, *Journal of Risk* 7, no. 2 (winter 2004/2005): 63–84.

Pickands, J. 1975. Statistical inference using extreme order statistics. *Annals of Statistics* 3: 119–131.