CONSTRUCTING LOW VOLATILITY STRATEGIES

Understanding Factor Investing

Mehdi Alighanbari, Stuart Doole, Lokesh Mrig, Durga Shankar

January 2016

CONTENTS

EXECUTIVE SUMMARY

Low volatility is one of the few factors that have historically performed well in turbulent markets. Moreover, over long periods of time, this defensive strategy has produced a premium over the market, contravening one of the most basic theories in finance — that one should not be rewarded with greater returns for taking less than market risk. Since the global financial crisis hit in 2008, low volatility has garnered increased attention from institutional investors.

Extensive research has investigated the low volatility anomaly, but the purpose of this paper is to discuss the practicalities of implementing a low volatility strategy. A low volatility strategy can be constructed in two key ways, using purely ranking-based (heuristic) approaches or optimization-based methods. While purely ranking-based approaches are simpler to understand, we find that optimization-based methods offer greater flexibility in constructing low volatility strategies. In addition, some purely ranking-based approaches provide unintended exposures to factors other than low volatility, which can affect the risk/return profile significantly. Optimization strategies can have shortcomings of their own; however, constraints can be used to fine-tune the construction methodology.

Using the MSCI World Minimum Volatility Index as an example, we demonstrate how a well-designed approach can benefit from the advantages and flexibility of an optimization-based methodology, while incorporating constraints that address the shortcomings of an unconstrained optimization such as high turnover and large active and unwanted sector and country bets.

The first step towards an effective minimum volatility index is a robust covariance matrix. Using a factor model and a fundamental factor model in particular can help reduce the number of parameters to be estimated and make the resulting covariance matrix more robust.

We reviewed the behavior of the MSCI World Minimum Volatility Index during various market regimes since its launch in 2008. The index reduced overall volatility by 30%, holding up better than the market during downturns. Over the long term, the index outperformed the market by 20 percentage points as the market itself gained 40%.

INTRODUCTION

The low volatility factor, while dating to the 1970s, has experienced renewed interest since the global financial crisis hit in 2008 as well as due to the growing adoption of factor indexes (also known as “smart beta”).

Historically, the factor’s performance has declined less than the market during times of crises and market downturns. When embedded in portfolios, the defensive characteristics of the factor have tended to protect capital during turbulent markets.

In addition, an extensive body of research shows that low volatility portfolios have outperformed the market over long periods of time; this outperformance has been persistent across time and regions. The low volatility factor’s performance is a puzzle because it is apparently at odds with one of the most basic principles in finance: that higher volatility is associated with higher returns. According to the Capital Asset Pricing Model (CAPM), one should not expect a long-term premium for taking *less risk* than the market as a whole. The historical return premium has mainly been explained using behavioral finance arguments, which we summarize in the next section.

Low volatility investing is a broad topic and a vast body of research has been dedicated to this subject. In this paper, we intend to address the practicalities of constructing low volatility strategies, responding to common questions that investors raise when evaluating these strategies.

This paper is the fifth in a series exploring each of the six key factors that have historically offered long-term excess returns: value, quality, momentum, yield, low volatility and low size.

LOW VOLATILITY INVESTING

While apparently contradicting one of the main principles of finance that higher risk is associated with higher returns, several empirical studies have demonstrated that lower volatility stocks have outperformed the market.[1]

Mostly behavioral arguments have been offered to explain the low volatility premium. Here are some of the leading explanations:

- **Lottery effect.** Some observers argue that buying a volatile stock is similar to buying a lottery ticket where the customer pays a small fee in the hope of winning a large amount of money — albeit at a very low probability. Therefore, investors often overpay for high volatility stocks and underpay for low volatility stocks due to the "irrational" preference for volatile stocks.
- **Representativeness.** Investors tend to overpay for "glamorous" high volatility stocks because of the well-publicized success of a handful of such stocks; the speculative nature of such stocks is often ignored by investors.
- **Overconfidence.** Investors are overconfident in their ability to forecast the future, and the extent of their differences in opinions is greater for stocks with more uncertain outcomes (high volatility stocks). Also, it is tougher for pessimistic investors to express a negative view via a short sell, resulting in optimists driving up the prices of high volatility stocks and hence lower expected future returns for high volatility stocks.
- **Agency issue.** Asset managers tend to avoid low volatility stocks because there is less research conducted by brokers and others on these less glamorous stocks.
- **Asymmetric behaviors.** When the market is on a declining trend, the dispersion of beta between low volatility and high volatility portfolios has tended to increase (i.e., low volatility stocks experienced a much lower beta, or risk, vis-à-vis the market). Therefore, the low volatility stocks have experienced smaller declines than their high volatility counterparts. When a bull market occurs, this dispersion has been smaller and thus low volatility stocks have underperformed only slightly. Net, the low volatility stocks have performed better over the long term.

[1] Black (1972), Haugen and Baker (1991), Chan, et al. (1999), Jagannathan and Ma (2003), Clark, et al.(2006), Ang et al. (2006), Blitz and Vliet(2007), Nielsen and Subramanian(2008), Sefton, et al. (2011), Baker and Haugen (2012), Frazzini and Pedersen (2014), Muijsson, et al. (2014), and Stambaugh, et al. (2015).

There are also theories that are not based on behavioral biases. For example, Baker, et al. (2011) observed that low volatility stocks often tend to have lower betas; overweighting them may lead to higher tracking errors for institutional investors. Such tracking errors need to be justified by sufficient excess returns (alpha). In other words, institutional investors cannot buy low volatility stocks wholesale. To a certain extent, the benchmark issue prohibits institutional investors from fully exploiting the low volatility anomaly.

Separately, Frazzini and Pedersen (2014) argue that the underperformance of higher beta assets is partly due to leverage constraints and margin requirements faced by many investors. According to CAPM, all investors invest in the highest Sharpe ratio portfolio and leverage or de-leverage this portfolio to meet their objectives. However, many investors are constrained in their use of leverage; instead, they increase their risky holdings. This increased demand for high beta assets may result in lower long-term risk-adjusted returns than for low beta assets.

Lastly, Muijsson, et al. (2014) explain the outperformance of low beta stocks based on interest-rate movements. Their analyses show that the low and high beta portfolios have been affected differently by changes in the risk-free rate. Returns on low beta portfolios have increased when the rate decreases and returns on high beta portfolios have risen when the rate increases. They conclude that the main factor behind low volatility anomaly likely stems from exogenous macroeconomic factors such as government monetary policies.

LOW VOLATILITY STRATEGIES: HEURISTIC VS OPTIMIZATION-BASED

Numerous methodologies have been developed over the years to implement low volatility strategies. The more recent phenomenon of “smart beta” indexes has sparked interest in creating investible low volatility indexes. All of these indexes and the underlying approaches can be categorized into two distinct groups: *heuristic* and *optimization-based*.

Heuristic approaches tend to be simple, purely ranking-based indexes. In comparison, optimization is a more sophisticated approach to creating low volatility indexes. This sophistication may make the process more complex but it provides significant flexibility that, if designed properly, can considerably improve the quality of the resulting index, improving its replicability and avoiding unintended exposures to styles, countries or sectors.

Let’s start with the formula to calculate the volatility of an index:

$$\sigma_p^2 = W^T.\ \Phi\ .W = \sum_{i=1}^{N} w_i^2 \sigma_i^2 + \sum_{i=1}^{N} \sum_{j=1, j<>i}^{N} \rho_{ij} w_i w_j \sigma_i \sigma_j$$

where σ_p is the volatility of the index returns, σ_i is the volatility of stock (asset) i in the index, w_i is the weight of stock i and ρ_{ij} is the correlation between stock i and stock j.

The objective is to find weights (w_i) that result in the lowest volatility for the index.[2] There are different methodologies to calculate these weights and create an optimal index portfolio but they all have two stages. The first stage is to estimate volatilities and correlations (the covariance matrix - Φ in the above equation). The second stage is to use these estimates to calculate the optimal weights. Therefore, the quality of a low volatility index depends on the accuracy of the estimations as well as how well the weights are calculated.

HEURISTIC APPROACHES

Many purely ranking-based approaches have been developed to create a low volatility index. The underlying principle of most of these approaches is to rank the universe of stocks based on the estimate of their volatility (total volatility, residual volatility or beta), select a subset of (or in some cases all) the constituents of the universe, and then apply different weighting schemes. Weighting can be determined by market capitalization, inverse of volatility, inverse of variance or various other methodologies.

Constraints may be applied to these heuristic approaches to ensure there are acceptable levels of liquidity and investability, controls for sector and country exposures and limits on stock weights. The MSCI Risk Weighted Indexes and Volatility Tilt Indexes fit into this category.

These approaches are simple and transparent and their weighting schemes enjoy a good degree of flexibility. However, they generally are based on the volatility of individual stocks and ignore the correlation between stock returns (the second term in the equation), which can have a major impact on strategy volatility when cross sectional variation between correlations is high.

Some heuristic approaches also fail to provide a pure exposure to low volatility, *implicitly* providing exposure to other factors. Such approaches derive some of their risk/return behaviors from these residual factors.

Some low volatility strategies *explicitly* incorporate other factors. For instance, one can sort and select securities based on their volatility and then weight them based on company valuations. This type of approach employs multiple factors; while there is a diversification benefit to combining factors, the risk/return profiles of such strategies can result from their exposure to, say, value, as much as to volatility.

[2] Currency risk is an important consideration in designing low volatility strategies. We discuss the effect of currency on these strategies in Appendix 2.

OPTIMIZATION-BASED APPROACHES

While heuristic approaches reflect the volatility of individual stocks, optimization-based approaches account for both volatility and correlation effects.[3] Optimization can be performed in various ways, but the differences usually stem from the covariance matrix estimation (Φ in the equation) and how constraints are applied.

The simplest approach to obtain the covariance matrix is to use the historical returns of individual stocks and calculate their historical volatilities and pairwise correlations.[4] There are two main issues with this approach: As the number of stocks in the universe increases, the size of covariance matrix and therefore the number of parameters to be estimated becomes very large, sometimes requiring estimation of millions of parameters. Also, stock volatility and the correlation between them can be very unstable; using historical levels may provide poor estimates of future values (Vangelisti, 1992).

A more common approach to optimization, especially for a large universe of stocks, is using a factor model, such as a simple statistical model that applies principal component analysis or a more elaborate fundamental factor model. These models effectively reduce the size of covariance matrix to be estimated, making calculations less complex and more stable: The size of covariance matrix remains constant for a fixed number of factors and does not change even if the number of stocks in the universe varies.

In addition, a fundamental factor model such as a simple 5-factor Carhart or commercial models take advantage of economic intuition to measure realistic and stable correlations across the investment universe. Fundamental factor models tend to use current stock characteristics, resulting in a timelier, stable and robust covariance matrix. The MSCI Minimum Volatility Indexes, which were introduced in 2008, currently use the Barra GEM2 factor model.[5]

[3] Please see Appendix 1 for a more detailed discussion regarding the effect of correlation and volatility on low volatility indexes.

[4] Correlation between returns of every pair of stocks in the selected universe.

[5] MSCI Minimum Volatility Indexes were launched using a previous version of the Barra equity model (GEM). With the introduction of the more advanced GEM2 Model, these indexes adopted the new model in 2009.

PRACTICALITIES OF OPTIMIZATION: USE OF CONSTRAINTS

A naively designed optimization can result in unwanted and extreme exposure to certain industries, countries or styles. In addition, a poorly designed optimization can result in high turnover at rebalancing periods. Adding constraints to the optimization, however, can mitigate these shortcomings without compromising the effectiveness of the optimization.

The design of the optimization and how constraints are incorporated are critical; manipulating the results by applying constraints after the optimization can undermine the entire exercise and make the resulting index suboptimal. For instance, optimizing the index without a sector constraint and later adjusting constituent weightings to impose sector constraints may impair the quality of the index. Similarly, running an unconstrained optimization might create an optimal long-short portfolio but the individual long and short legs can be far from optimal. Thus, constraints should be built into the optimization and the results should remain untouched.

Selecting constituents prior to optimization can be sensible. For instance, if we want to exclude certain stocks in the strategy, we can simply remove them before the optimization.

COMMON DRAWBACKS OF UNCONSTRAINED MINIMUM VOLATILITY

The three major issues with an unconstrained minimum volatility portfolio are: 1) biases towards certain sectors and countries, 2) unwanted style exposures and 3) high turnovers at rebalancing. In what follows we look at these potential shortcomings and demonstrate how they can be eradicated using well-designed constraints. We also mention some other ways to deal with these issues that can be less effective or in cases detrimental to the optimization process.

SECTOR, COUNTRY AND STYLE BIASES

Stocks from different sectors and countries exhibit varying levels of volatility. For instance, utility companies tend to be low volatility stocks while information technology stocks tend to be more volatile. An unconstrained minimum volatility portfolio would often have a positive and persistent bias towards low volatility sectors and negative exposure to higher volatility ones. Similarly, some country markets tend to be less volatile while others are more volatile. Though sometimes these biases are desirable, in many cases investors prefer to limit active exposures to sectors or countries.

An effective optimization framework incorporates constraints to limit unintended and unwanted biases. For example, the MSCI Minimum Volatility Indexes uses the Barra Open Optimizer, which allows for incorporating a wide range of constraints, to calculate the index.

Exhibit 1 and Exhibit 2[6] compare the sector and country exposures of a constrained minimum volatility index to those for an unconstrained index. We have used the MSCI World Index as of November 25, 2014 as the universe and we have constrained the optimization to be long only. The only additional constraint applied in this example is to keep sector and country exposures of the minimum volatility portfolio within 5 percentage points of the parent index weightings.

Without these constraints, we would see large biases in consumer staples and utilities as well as in Japanese and Hong Kong equities. Adding these constraints creates an optimized index while avoiding large and unwanted bets on any sector or country.

Exhibit 1: Active Sector Exposures Constrained vs. Unconstrained Min Vol Strategies

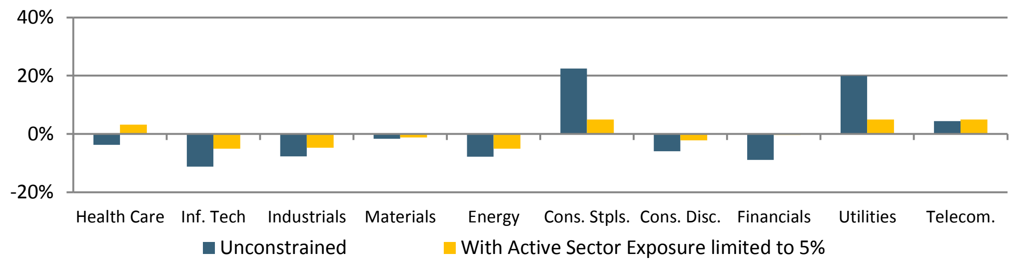

Data as of November 25, 2014

Exhibit 2: Active Country Exposures Constrained vs. Unconstrained Min Vol Strategies

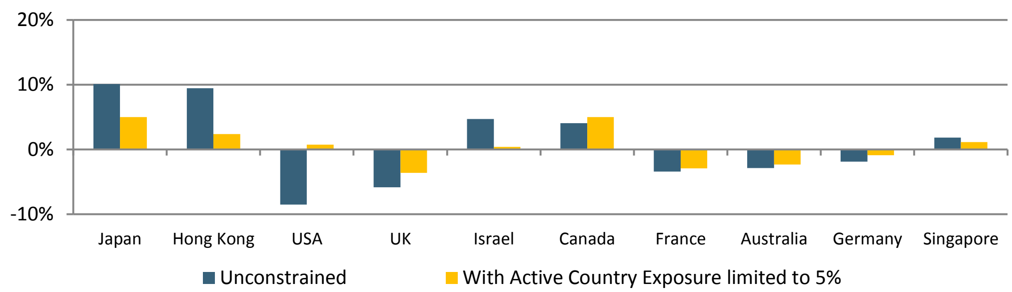

Data as of November 25, 2014

Style factors such as value, leverage and size are important in designing a strategy. Sometimes, exposures to these factors are intended to capture historical long-term premia and are explicitly part of the design and optimization process. But often these exposures emerge unknowingly and randomly. When creating a strategy, whether through an optimization-based or heuristic approach, the factor exposures will deviate from the market.

[6]Only countries with significant exposure have been included.

For instance, a low volatility portfolio may result in unintentional low size and high value exposures.

These unintended active style exposures may be tolerated to a certain extent, as a long-only strategy often includes some level of residual active style exposure. But with the right tool and design, these unintended style exposures can be restricted to specified levels. Fundamental-based models such as Barra factor models (used in conjunction with the Barra Open Optimizer) explicitly allow limits on unwanted style exposures.

In Exhibit 3, we contrast the effect of an optimized minimum volatility index with no style constraints with one that limits exposure to style factors. In the unconstrained index, while the large negative exposure to volatility is intended, the large active residual exposures to growth and size are accidental and may be unwanted. In the MSCI Minimum Volatility Index, we constrain all the style factors excluding volatility to within 0.25 standard deviation of their parent index. The constraints keep all the styles within range while having minimal impact on the desired volatility exposure.

The constraint on the value factor in the MSCI Minimum Volatility Indexes also implicitly prevents the index from over-weighting richly valued companies (crowded stocks). While low volatility stocks tend to be high quality stocks and show higher valuations, these constraints ensure that the resulting index limits exposure to over-valued stocks compared to the relevant equity market or the parent index.

Exhibit 3: Active Style Exposures – Constrained vs. Unconstrained Min Vol Strategies

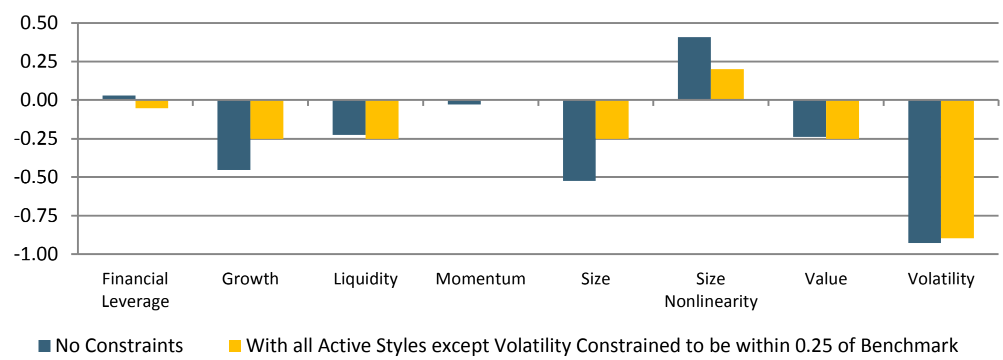

Adding any constraint to the minimum volatility optimization results in an index with greater expected volatility, but for a well-designed optimization the increase is minimal. Exhibit 4 illustrates the considerable reduction in expected volatility achieved by moving from the market cap index (the MSCI World Index) to a minimum volatility index, with varying levels of constraints. It also reveals that a small increase in expected volatility occurs when different constraints are added to the unconstrained optimization.

Note: Applying the constraints after the optimization may cause a considerable deviation from optimality and in many cases a feasible solution may not be possible.

Exhibit 4: Expected Volatility for Different Levels of Constraints

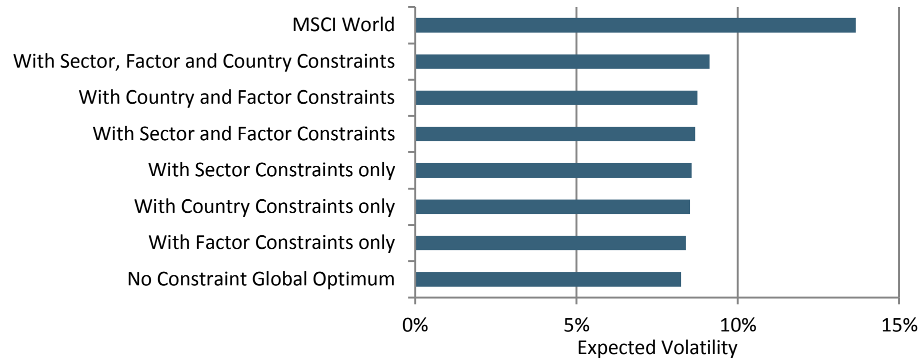

TURNOVER

A poorly designed minimum volatility index may also experience high turnover. One can reduce the rebalancing frequency to limit turnover to a certain extent while still maintaining the desired level of minimum volatility exposure. To further reduce the turnover to the desired level, we can explicitly apply turnover constraints to the optimization.

There is a clear trade-off between the level of turnover and the reduction in expected volatility. However, this relationship is not linear. The marginal benefit of incurring additional turnover to reduce volatility decreases as turnover increases.

To illustrate this point, first we examine the effect of going from a market cap index to a minimum volatility index. We run optimizations on the MSCI World Index while changing the constraint on turnover (illustrated by the blue line in Exhibit 5). Allowing 50% turnover relative to the market cap index would have reduced the volatility to 9.6 % from 13.7%. In comparison, the index with no turnover constraint would have achieved volatility of 9.1% with 76% turnover. This means 90% of possible risk reduction would have been achieved by allowing 50% turnover relative to the MSCI World Index.

More important is to examine turnover of the minimum volatility index at the rebalancing dates, when stocks are added to and subtracted from the index. In this example, starting from the MSCI Minimum Volatility Index (before rebalancing), the volatility level of 9.6% can be achieved by allowing only 10% turnover, also resulting in a 90% of possible risk reduction compared to the parent index. The yellow line illustrates the risk reduction that is achieved by different turnover constraints when we start from a minimum volatility index just before rebalancing.

For the MSCI Minimum Volatility Indexes, we have chosen to rebalance semi-annually and constrain the turnover to 10% on each rebalancing, resulting in 20% annual one-way turnover.

Exhibit 5: Effect of Turnover Constraint on Expected Volatility

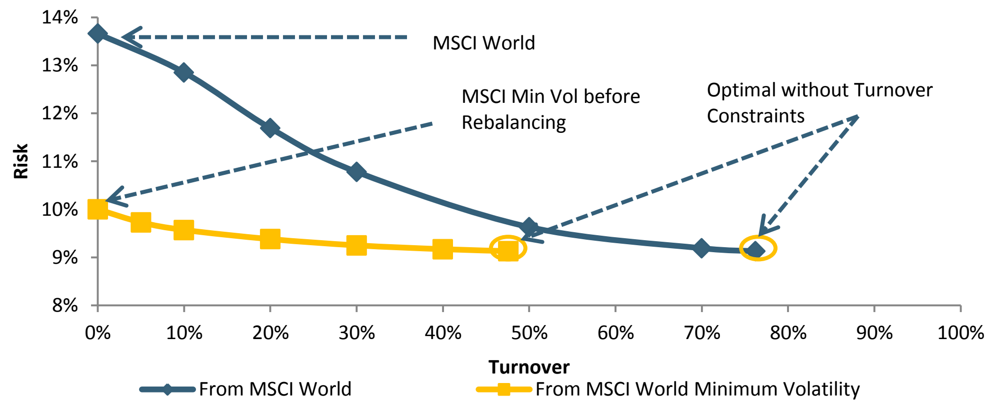

Data as of November 25, 2014

TURNOVER AND PATH DEPENDENCY

While a constraint is necessary to avoid excessive turnover and associated costs, it can create path dependency — that is, the index constituents and their weighting can depend on the index's launch date. This can become problematic when the index deviates significantly from the optimal index. Careful design can mitigate sub-optimality due to path dependency.

To see if path dependency has affected the MSCI Minimum Volatility Indexes,

Exhibit 6 and Exhibit 7 compare *ex-post* risk/return behavior of indexes that are constructed in exactly the same way, except for their starting dates. The figures show considerable *ex-post* risk reduction and increases in the returns of minimum volatility indexes compared to the market cap index. Moreover, the differences in annualized volatility between the five different minimum volatility indexes are negligible and there is no evidence of better performance for the newer indexes. Thus, the respective start dates have not affected either riskiness or returns.

We also looked at the difference between the MSCI Minimum Volatility Index and its turnover-relaxed counterpart over time. We look at two measurements, *ex-ante* risk spread and active share.[7]

Exhibit 8 shows the *ex-ante* risk for the MSCI World Index as well as two versions of the minimum volatility index, with and without turnover constraints. The risk for the two versions of minimum volatility indexes is almost identical; it is hard to see the difference. On the right axis, we measured this small difference (shaded line). Not only is the difference small (10-30 basis points), there does not seem to be any obvious upward trend that might indicate the turnover constraint effect is accumulating over time.

In Exhibit 9, we used active share to show how different the two indexes (with and without turnover constraints) are and whether they diverged over time. As we expected, the turnover constraint resulted in differences in the two indexes. While there was an increase in active share initially, it seems that this parameter stabilized, meaning that the constrained index does not vary much from the optimal approach over time.

An unconstrained optimized index is created by considering only the main objective — reducing its volatility — ignoring other important considerations such as capacity and concentration, liquidity, turnover and unintended exposures. However, alternative indexes can achieve very similar levels of volatility reduction while also accounting for these other considerations. Through use of constraints, the optimization process can create an alternative index that not only achieves the main objective of reducing volatility but also addresses these other investment considerations.

Exhibit 6: Analyzing Path Dependency: Ex-Post Volatility

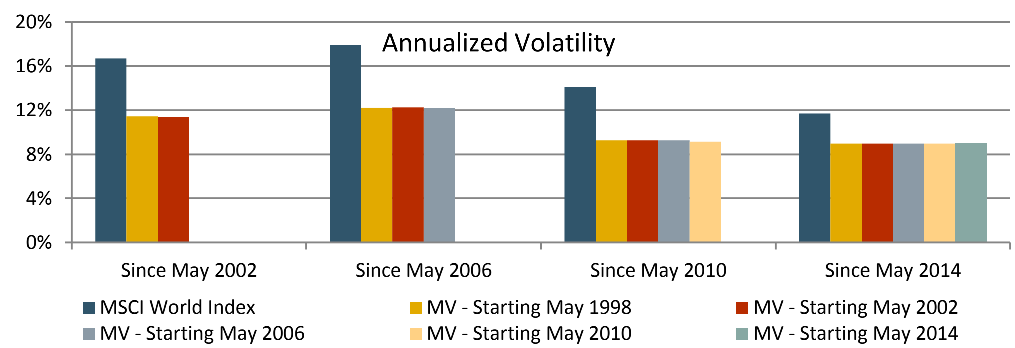

[7] Active share measures how two indexes (or portfolios) differ from each other by comparing the weight of each stock in the indexes. It is calculated as half of the sum of the absolute difference between the weights of each stock in the two indexes.

Exhibit 7: Analyzing Path Dependency: Ex-Post Return

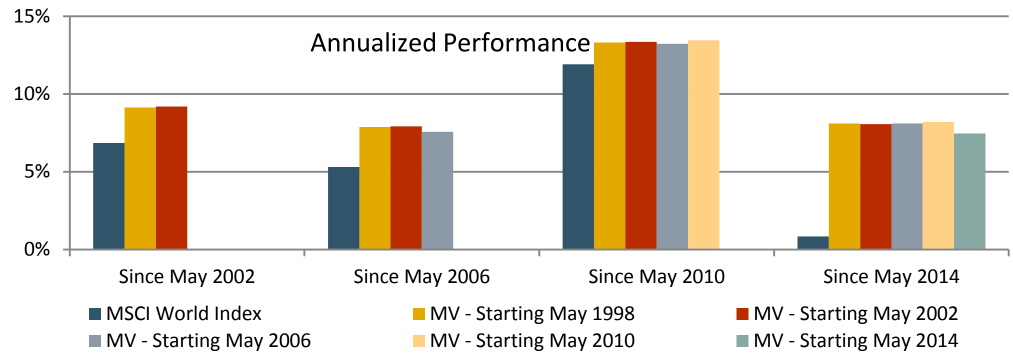

Exhibit 8: Analyzing Path Dependency: Ex-Ante Volatility

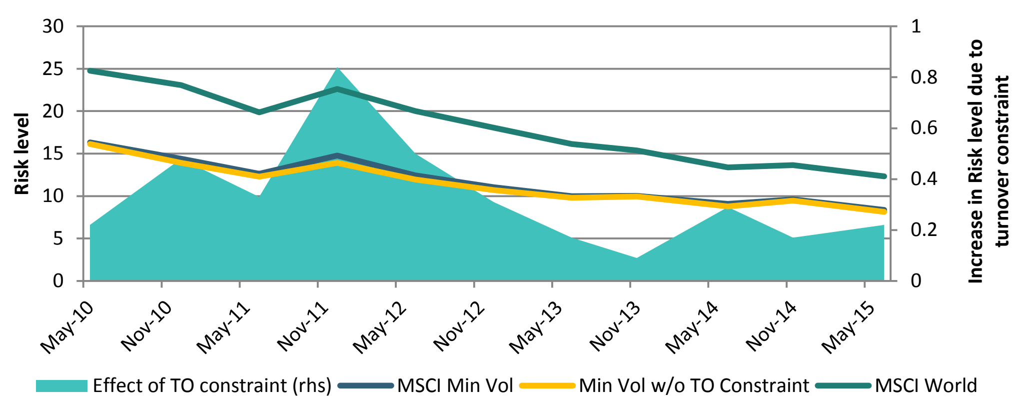

Exhibit 9: Analyzing Path Dependency: Active Share

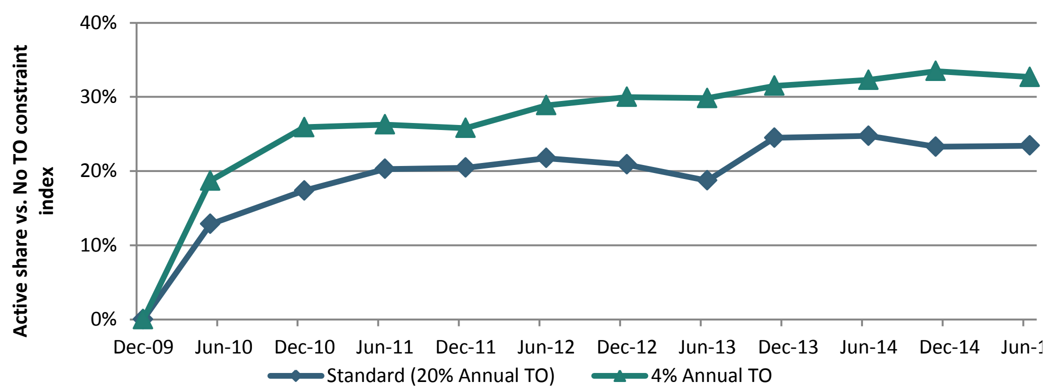

PERFORMANCE OVER MARKET CYCLES AND ECONOMIC REGIMES

While risk is often measured by the level of volatility, there are other measures of risk that are important, such as drawdowns during bear markets and market turmoil. In this section, we present several analyses to show how the MSCI Minimum Volatility Index has historically performed during different market regimes and during rising and falling markets.

PERFORMANCE IN BEAR MARKETS

We start by analyzing the performance of MSCI World Minimum Volatility Index during different market downturns (bear markets) over the past 27 years[8] (Alighanbari, et al. 2014). We define bear market as a decline of 20% or more in the MSCI World Index for a period lasting at least two months. There were four bear markets during this time frame (gray shaded areas in Exhibit 10). The MSCI World Minimum Volatility Index (blue line) outperformed the market (the MSCI World Index) across all four bear market periods, demonstrating its strong defensive characteristics.

Exhibit 10: Minimum Volatility Behavior in Bear Markets

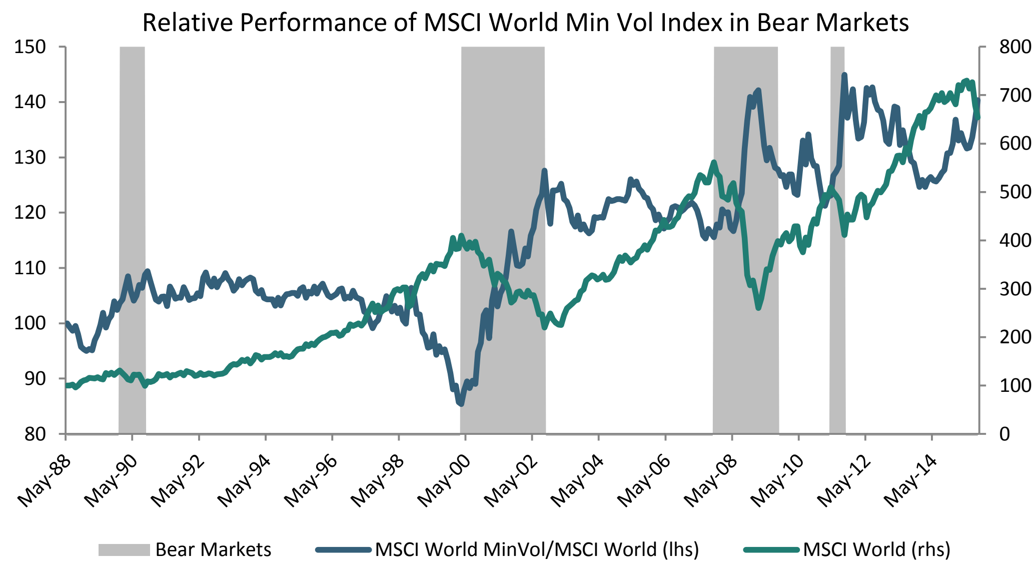

Exhibit 11 shows that the MSCI World Minimum Volatility Index experienced substantially smaller drawdowns during these market downturns when protecting wealth was most

[8] MSCI Minimum Volatility historical data starts at May 31, 1988. Please refer to the disclaimers at the end of this document regarding use of simulated or backtested data.

important. In addition, the Minimum Volatility Index also produced lower realized volatility during the turbulent periods.

Exhibit 11: MSCI Minimum Volatility Index Risk/Return over Bear Markets

	Absolute Returns (Gross USD)			Realized Volatility	
Bear Market Periods	MSCI World	MSCI World Minimum Volatility	Active Return	MSCI World	MSCI World Minimum Volatility
Dec 89 - Sep 90	-24.0%	-20.2%	3.8%	21.8%	19.7%
Mar 00 - Sep 02	-46.3%	-19.8%	26.5%	16.5%	11.0%
Oct 07 - Feb 09	-53.7%	-43.0%	10.6%	21.9%	17.1%
Apr 11 - Sep 11	-19.4%	-5.1%	14.3%	15.9%	8.9%

UPSIDE POTENTIAL

Not surprisingly, the MSCI World Minimum Volatility Index has outperformed the market when the market has declined overall. In Exhibit 12, we see that the index outperformed a negative market 88% of the time with an average outperformance of 8.8 percentage points based on one-year rolling periods. Where the index underperformed the market, the average shortfall was only 1.24 percentage points.

In the years where market return exceeded 10%, the Minimum Volatility Index underperformed; the level of underperformance increased as the market return rose. However, for moderate positive return periods (0%-10%), the index outperformed the market 67% of the time.

Exhibit 12: Performance Comparison: MSCI World Minimum Volatility Index vs Market

MSCI World Rolling 1-Year Return	<0	0-10%	10%-20%	20%-30%	>30%
Hit rate of Outperformance	88.6%	67.3%	39.8%	21.2%	0.0%
Average Outperformance	8.8%	4.5%	3.1%	1.3%	0.0%
Average Underperformance	-1.24%	-2.6%	-3.8%	-6.0%	-9.3%
No. of Observations	79	52	118	52	16

BEHAVIOR DURING DIFFERENT ECONOMIC REGIMES

The minimum volatility factor has performed well in defensive markets. Here we examine its behavior against various market/macro signals, extending previous MSCI research on how various equity factors behaved in changing economic environments (Gupta et al., 2014). Exhibit 13 extends the bivariate regime analysis introduced by Gupta et al. to include market risk indicators such as VIX and credit spreads. The blue bars in this exhibit show the average monthly active return of the MSCI World Minimum Volatility Index when the indicated market/macro measure is decreasing, while the yellow bars indicate when the measure is rising. On average, the Index outperformed the market during periods of economic contraction (defined as a decrease in the OECD's Composite Leading Indicators), high volatility (increasing VIX), widening of credit spread and rising inflation, as well as falling interest rates. The pattern of outperformance re-emphasizes the defensive behavior of low volatility strategies.

Exhibit 13: MSCI World Minimum Volatility Behavior during Market/Macro Regimes

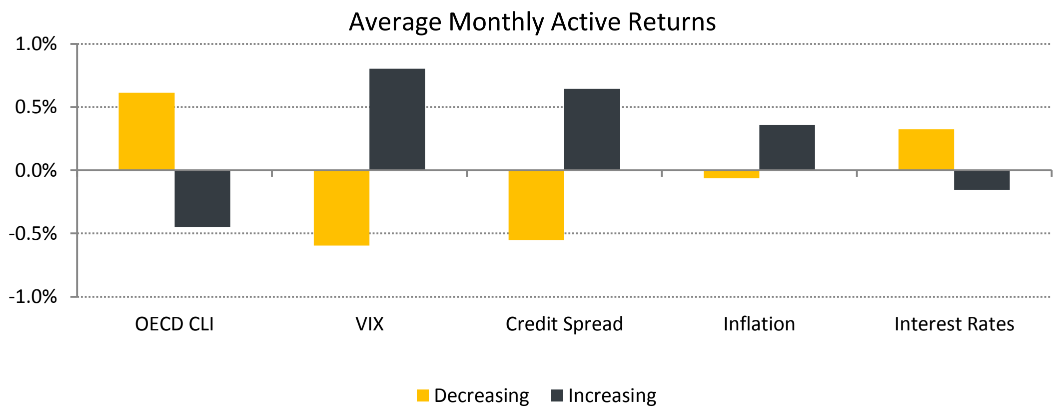

MSCI MINIMUM VOLATILITY INDEX: OUT OF SAMPLE PERFORMANCE

Since the MSCI World Minimum Volatility Index family was launched in 2008,[9] it has returned about 60% while the market has gained about 40% (Exhibit 14). The index's volatility over this out-of-sample period (calculated using monthly returns) has been 12.8% compared to 18% for the parent index, a roughly 30% reduction.

In Exhibit 15, the shaded area shows the performance of the MSCI World Index over this period and the yellow line the performance of the Minimum Volatility Index relative to the MSCI World Index. The MSCI Minimum Volatility Index was launched as the market plunged about 50% in 2008; the index outperformed the market by about 20 percentage points. Similarly, in 2011, the market declined about 20% and the Minimum Volatility Index avoided most of that drawdown. More recently, the Minimum Volatility Index held up well during market turbulence in August 2015.

The MSCI World Minimum Volatility Index has exhibited considerably lower volatility than the broad market index since its launch. It has also demonstrated strong defensive characteristics, significantly outperforming the market during market downturns. Finally, it has outperformed the market over the seven-year period, providing a long-term premium. This outperformance was achieved over a period where the market itself gained 40%.

[9] MSCI World Minimum Volatility (USD) was launched on April 14, 2008

Exhibit 14: MSCI World Minimum Volatility Index Performance since Launch

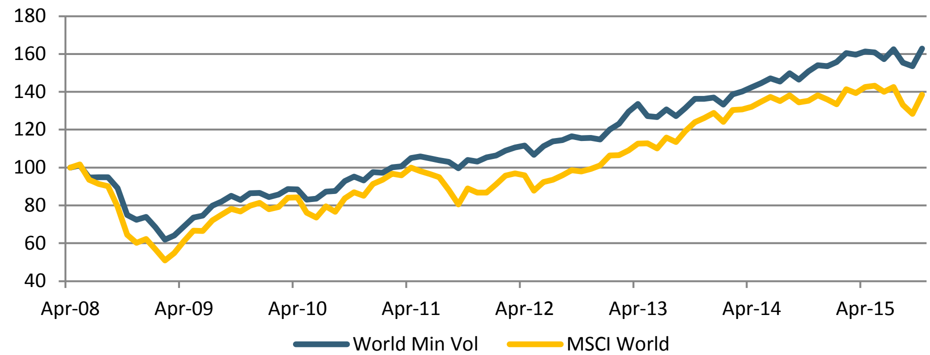

Exhibit 15: Defensive Behavior of MSCI World Minimum Volatility Index

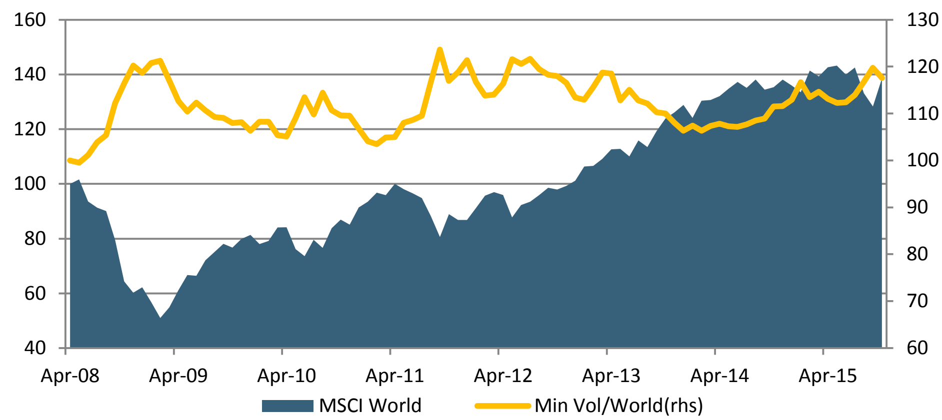

CONCLUSION

The minimum volatility factor is one of the few factors that have performed well during turbulent markets, providing capital preservation when it is needed most. Yet it remains an anomaly as it has produced better-than-market returns over long time periods while offering less risk.

In this paper, we looked at some of the practicalities of designing a low volatility strategy. There are two main ways to design these strategies: heuristic (purely ranking-based) and optimization-based approaches. While heuristic approaches tend to be simpler, optimization-based approaches provide a more flexible framework to incorporate different types of constraints. Moreover, only optimization-based approaches can take full advantage of the correlation between stocks, a key component in designing a low volatility strategy.

Low volatility indexes, whether created using a purely ranking-based approach or by optimization, can result in large unintended tilt towards other style factors. While combining multiple factors can be a sensible approach for diversification purposes, sometimes these residual factors can have more effect on the risk/return profile of the strategy than the volatility factor itself.

Optimization-based approaches have their pitfalls. Estimating the full covariance matrix can be cumbersome as the number of stocks increases. Use of a factor model (especially a fundamental factor model), can help reduce the number of parameters to be estimated and make the resulting covariance matrix more stable.

Constraints also are important in designing a minimum volatility index. Implementing constraints directly in the optimization can help achieve the desired level of sector, country and style exposures and limit turnover without compromising much in risk reduction. Applying constraints after optimization can defeat the whole purpose of the optimization and can result in an inferior index.

Finally, we examined the behavior of MSCI World Minimum Volatility Index during different market regimes to show the different characteristics of the index. Using 27 years of available data as well as the seven years of history since the index has been live, the index has produced considerably lower volatility than the market, has behaved defensively in market downturns and has outperformed the market during these periods. Over the long term, this index has generated superior performance, benefiting from the low volatility premium.

REFERENCES

M. Alighanbari,, R.A. Subramanian and P. Kulkarni. (2014). "Factor Indexes in Perspective: Insights from 40 Years of Data." MSCI Research Insight.

A. Ang, R. Hodrick, Y. Xing and X. Zhang. (2006). “The Cross-Section of Volatility and Expected Returns.” *Journal of Finance* 61(1), 259-299.

M. Baker, B. Bradley and J. Wurgler. (2011). “Benchmarks as Limits to Arbitrage: Understanding the Low-Volatility Anomaly.” *Financial Analysts Journal* 67(1), 40-54.

N.L. Baker and R.A. Haugen. (2012). “Low Risk Stocks Outperform within All Observable Markets of the World.” available at SSRN: http://ssrn.com/abstract=2055431.

F. Black. (1972). “Capital Market Equilibrium with Restricted Borrowing.” *Journal of Business* 45(3), 444-455.

D. Blitz and P. Van Vliet. (2007). “The Volatility Effect: Lower Risk without Lower Return.” *Journal of Portfolio Management* 34(1), 102-113.

L.K.C. Chan, J. Karceski and J. Lakonishok. (1999). “On Portfolio Optimization: Forecasting Covariances and Choosing the Risk Model.” *The Review of Financial Studies* 12(5), 937-974.

Y. Choueifaty and Y. Coignard. (2008). “Towards Maximum Diversification*.” Journal of Portfolio Management* 34(4), 40-51.

R. Clarke, H. De Silva and S. Thorley. (2006). “Minimum-Variance Portfolios in the US Equity Market.” *Journal of Portfolio Management* 33, 10-24.

A. Frazzini and L.H. Pedersen. (2014). "Betting Against Beta." *Journal of Financial Economics* 111(1), 1-25.

A. Gupta, A. Kassam, R. Suryanarayanan and K. Varga. (2014). “Index Performance in Changing Economic Environments.” MSCI Research Insight.

R. Haugen and N. Baker. (1991). “The Efficient Market Inefficiency of Capitalization-Weighted Stock Portfolios.” *Journal of Portfolio Management* (Spring 1991), 35-40.

R. Jagannathan and T. Ma. (2003). “Risk Reduction in Large Portfolios: Why Imposing the Wrong Constraints Helps.” *Journal of Finance* 58(4), 1651-1684.

C. Muijsson, E. Fishwick and S. Satchell. (2014). “Taking the Art out of Smart Beta.” http://sydney.edu.au/business/__data/assets/pdf_file/0003/214356/DP-2014-008_Taking_The_Art_out_of_Smart_Beta.pdf. Discussion paper.

F. Nielsen and R.A. Subramanian. (2008). “Far From the Madding Crowd – Volatility Efficient indexes.” MSCI Barra Research Insights.

J. Sefton, D. Jessop, G. De Rossi, C. Jones and H. Zhang. (2011). “Low-risk investing,” UBS Research.

R.F. Stambaugh, J. Yu and Y. Yuan. (2015). “Arbitrage Asymmetry and the Idiosyncratic Volatility Puzzle.” *Journal of Finance* 70(5), 1903-1948.

M. Vangelisti. (1992). “Minimum-Variance Strategies: Do They Work?” Barra Newsletter.

APPENDIX 1: CORRELATION AND VOLATILITY

CORRELATION MATTERS

An optimization-based strategy relies on both correlation and volatility to create the minimum volatility index. In this section, we look at the effect of correlation on the overall volatility reduction of the minimum volatility index.

While we know the reduction of volatility (minimum volatility index compared to the market cap index) comes from both selecting lower volatility stocks and selecting lowly correlated stocks, it is not easy to separate the two effects. Below, we create a proxy for the effect of selecting lower volatility stocks versus the correlation reduction.

For these analyses we use the MSCI World and MSCI World Minimum Volatility indexes as of November 26, 2014. The risk levels for stocks as well as indexes are taken from the Barra GEM2 model.

Exhibit 16 demonstrates the distribution of volatility of the stocks in each index. The MSCI World had a relatively symmetric distribution around 25% volatility (with a slight tail for the higher volatilities). The MSCI World Minimum Volatility Index, as expected, picked more of the lower volatility stocks and therefore is skewed towards the left side of the graph. (Note: We are ignoring the weight of each stock in the index; this graph shows only the percentage of the stocks that are in each volatility segment.)

Exhibit 16: Volatility Distribution: MSCI World Minimum Volatility Index vs Market Cap

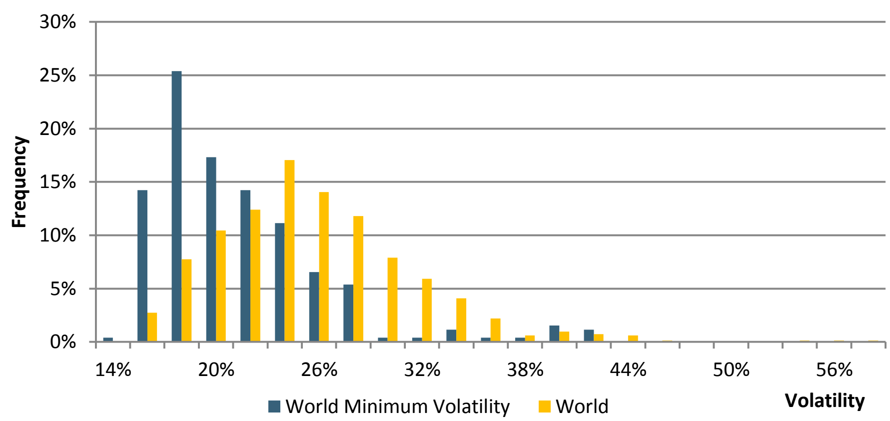

Clearly, some of the reduction that we see in the expected volatility of the MSCI World Minimum Volatility Index comes from selecting lower volatility stocks. This can be confirmed by looking at the average and weighted average of the constituents' volatility in the MSCI World Minimum Volatility Index versus its parent market cap index (Exhibit 17).

Exhibit 17: Volatility Reduction of MSCI World Minimum Volatility Index

	MSCI World	MSCI World Minimum Volatility
Volatility Estimate	13.66%	9.62%
Average Volatility of Constituents	24.69%	20.36%
Wt. Avg Volatility of Constituents	22.16%	18.83%

But the question remains how much of the volatility reduction comes from selecting lowly correlated stocks. To answer this question, let’s try separating the correlation effect. To do this, we need to make several assumptions and approximations.

For each index we have:

$$\sigma_p^2 = \sum w_i^2 \sigma_i^2 + \sum_{i<>j} \rho_{ij} w_i w_j \, \sigma_i \sigma_j$$

Now we assume that the correlation between all the stocks in the portfolio is equal to an average correlation:

$$\forall i, j; \quad \rho_{ij} = \rho_P$$

With this assumption, we can then calculate the average correlation level as:

$$\rho_p = \frac{\sigma_p^2 - \sum w_i^2 \sigma_i^2}{\sum_{i<>j} w_i w_j \, \sigma_i \sigma_j}$$

Applying this formula to the MSCI World Index and the MSCI Minimum Volatility Index, we find:

$$\rho_{World} = 0.38$$

$$\rho_{\mathrm{Min\ Vol}} = 0.26$$

The numbers clearly show that the MSCI Minimum Volatility Index is benefiting from a lower correlation between stocks selected.

To calculate a proxy for the effect of correlation on volatility reduction, we insert the correlation of the MSCI World Index (ρ_{World}) into the above equation for the MSCI World Minimum Volatility Index:

$$\hat{\sigma}_{\mathrm{Min\ Vol}}^2 = \sum w_i^2 \sigma_i^2 + \sum_{i<>j} \rho_{\mathrm{World}} w_i w_j \, \sigma_i \sigma_j = (11.6\%)^2$$

This means that if the correlation had stayed the same, the effect of selecting (and overweighting) lower volatility stocks in the minimum volatility index would have been a reduction in volatility from 13.7% to 11.6%. In other words, we can argue that from the 4.0% reduction of the volatility (Exhibit 17), approximately 2% comes from selecting lower volatility stocks and 2% comes from the correlation effect. While a number of assumptions are used in this calculation, the results demonstrate that selecting lower volatility stocks and giving them higher weights than the parent index as well as selecting stocks with lower correlations contributes to the reduction in volatility of the index.

A NOTE ON MAXIMUM DIVERSIFICATION

Discussing the maximum diversification approach to portfolio construction is out of the scope of this report. In this section, we briefly discuss the basic difference between the maximum diversification and minimum volatility approaches. The "*Most Diversified Portfolio"* is defined (Choueifaty and Coignard 2008) as the portfolio with the *maximum diversification ratio*, defined as:

$$D(P) = \frac{W'\emptyset}{\sqrt{\mathrm{W}'\Phi\mathrm{W}}}$$

Where W is the vector of weights of individual assets in the portfolio, Φ is the covariance matrix and Ø is the vector of asset volatilities. It is clear from the formula that the denominator is the volatility of the portfolio and the numerator is the weighted sum of the volatilities of the portfolio constituents. If we remove the numerator, the problem becomes a minimum variance problem. Including the numerator, maximization would try to minimize the volatility of the portfolio (the denominator) while selecting and overweighting higher volatility stocks. By doing so, the approach tries to achieve a reduction of portfolio volatility through selecting more volatile but lowly correlated assets. This process achieves the objective of having a highly diversified portfolio (assets with low correlation to each other) but at the same time targets high volatility stocks. In a sense, this approach is in line with CAPM theory and contrasts with the low volatility anomaly.

Comparing a maximum diversification portfolio with a purely ranking-based low volatility portfolio and minimum volatility portfolio, we find different treatments of volatility and correlations. In the purely ranking-based portfolio, the emphasis is on selecting and overweighting low volatility stocks while ignoring the correlation effect on reducing the overall volatility of the portfolio. This approach contrasts with maximum diversification, which selects high volatility stocks while trying to achieve portfolio volatility reduction through the correlation effect. The minimum volatility approach fits somewhere in between: It benefits from the correlation effect on portfolio volatility reduction, but at the same time implicitly selects and overweights lower volatility stocks.

ASSET ALLOCATION AND LOW VOLATILITY INDEXES

Even ignoring the long-term premium that has been achieved by low volatility indexes, their lower volatility puts them somewhere between equities and fixed income. Therefore, low volatility indexes can be used to construct a separate asset class in the allocation process. Incorporating low volatility indexes could have been especially helpful in times of low rates where investors often struggled to generate targeted rates of returns from their bond portfolios.

In Exhibit 18, we look at the effect of incorporating low volatility indexes into the asset allocation process. Replacing a market cap-weighted equity allocation with a low volatility index-based investment enhanced the return while reducing overall portfolio risk during the study period. Substituting a 60% allocation to market-cap-weighted equity with a low volatility index-based portfolio in a 60%/40% equity/fixed-income portfolio reduced risk by roughly 20% while enhancing return by an annualized 60 basis points during the study period.

In addition, incorporating the use of low volatility indexes in the asset mix in this way would have allowed investors to increase their allocation to equities without increasing risk. An 80%/20% allocation mix to low volatility equity and fixed income resulted in similar risk levels to those seen in the traditional 60%/40% allocation but with a higher return due to the premia from equities in general as well as from low volatility stocks.

Exhibit 18: Effect of Using Low Volatility Index Portfolios in Asset Allocation

	Equity Allocation					
Total Equity	MSCI World	Min Vol	Fixed Income Allocation*	Portfolio Return**	Portfolio Risk**	Risk Reduction***
60%	60%	0%	40%	6.69%	9.92%	
60%	40%	20%	40%	6.92%	9.11%	8.1%
60%	20%	40%	40%	7.12%	8.44%	14.9%
60%	0%	60%	40%	7.30%	7.94%	19.92%
80%	0%	80%	20%	7.63%	9.60%	3.18%

*Barclays Capital Global Aggregate

** Statistics determined over the period from Jan 1990 - Sep 2015

*** Compared to 60% MSCI World / 40% Fixed Income allocation

APPENDIX 2: CURRENCY HEDGING ISSUES

Currency treatment is an important consideration when it comes to low volatility indexes, regardless of whether constructed through optimization or heuristically. When designing a low volatility portfolio, it is important to specify the investor's local (reference) currency.

Let's suppose an investor wants to create a one-stock portfolio with the lowest historical volatility of daily returns from a universe of two Japanese stocks. A *Japanese* investor can simply pick the stock with lower historical volatility calculated using yen returns. For a *U.S. investor*, however, the answer may be different. The historical volatilities now should be based on the performance of each stock in USD. Depending on the correlation between the currency movement and stock movements, the lower volatility stock may be different for the Japanese investor than for the U.S. investor.

Similar considerations apply for the creation of low volatility indexes. For the MSCI Minimum Volatility Indexes, the base currency is defined in the equity risk model and within the optimization; different indexes can be generated that are *ex-ante* optimal for different investor reference currencies.

The volatility of an equity index depends on two things: 1) the volatility and correlation of the underlying stocks and 2) the behavior of the currencies of those securities and how the currency returns have related to stock returns. This equity-currency relationship is captured by the chosen global Barra Equity Model and hence in the optimization process where overall (forecast) index volatility is minimized. It is a little more complicated for a minimum volatility portfolio based on an international market that uses a different currency (for example, a U.S.-based investor looking at an MSCI Japan Minimum Volatility Index).

The choice of currency affects the composition of the low volatility index because the index will have some home bias for the investor as it reflects stock selections by the optimizer to reduce overall index risk. Is there a "correct" currency to use as a base for this optimization process? The theoretical answer depends on whether the index will be currency-hedged.

Suppose we have a Japanese investor interested in a minimum volatility index based on MSCI Japan. The investor and index currencies are identical. Yen is the base currency and currency plays no role. Effectively, we ignore the currency-related segments of the global Barra factor model (by setting currency exposures to zero). For a U.S. investor, the optimal (lowest forecast volatility) index is achieved with a base currency of the U.S. dollar. The sign of the correlation between stock returns and currency returns helps drive stock selection and weighting. To pick between two Japanese stocks with the same local currency volatility, we look for the one with the most negative (or least positive) correlation to the currency pair.

If the U.S. investor intends to hedge her equity currency exposure, then she would want to see the same returns in U.S. dollars as a Japanese investor sees in yen. Hence, the optimization currency is yen (local currency) and the USD-JPY FX hedge is applied to the resulting minimum volatility index.

LOCAL CURRENCY MINIMUM VOLATILITY INDEXES

When a currency exposure is “perfectly hedged” in an international portfolio, the investor sees (in her own local currency) the same return as a local investor in those same securities in their local currency. If the parent index has multiple currencies, the idea is the same for each currency in turn. The impact of the currency-related covariance in the risk model should be zero: in practice, we can set all the currency elements of the exposure matrix to zero. The optimizer will then to build an MSCI “local currency” minimum volatility index: the same for all investor perspectives.

Formally, of course, a minimum volatility index with *any* choice of base currency can be hedged to a target (investor) currency, but in the medium term we will be effectively (economically) over- or under-hedging with this “mixed” approach (so named as we are mixing hedging of currency risk via optimization with currency risk management with FX forwards). Local currency indexes offer a “cleaner” base for investors employing a hedging strategy with strong currency views. The view might be directional on the currency, or directional with respect to currency volatility, or even driven by non-market considerations. For example, the latter view might be driven by an accounting-based aversion to seeing currency returns influence profits in an investment portfolio subject to gain/loss constraints.

The local currency approach may also be convenient as the benchmark for active portfolio managers who run a single low volatility strategy portfolio but have investors or sub-funds with different FX hedging requirements. Other investors may stick with a mixed approach for operational simplicity (existing AUM, ease of establishing a new fund or share classes) after considering carefully the tracking error and hedge differences in different market conditions. The choice of a reserve currency outside the parent index’s currency region as the base currency is sometimes a surprisingly close proxy to the local currency solution (in a way that a commodity currency chosen as base currency is unlikely to be). Ultimately, the hedging strategy will depend on the investment horizon of the asset owner and her tolerance of currency-related impacts and capacity to directly hedge currency risk.

CASE STUDY: A GLOBAL “LOCAL CURRENCY” MINIMUM VOLATILITY INDEX

We now consider a USD-denominated investor looking for an MSCI ACWI Minimum Volatility index. We look at two cases: first, an index which takes account of all sources of risk; and second, the local currency version, which does not incorporate estimates of currency volatilities and correlations. In relation to risk minimization, the former approach does a better job of medium-term risk reduction for unhedged returns. However, we have already noted that the investor use cases for local currency construction place greater weight on other investment considerations.

We looked at the influence of the base currency on index characteristics over a full (12-year) simulation period and then, given the rise in FX volatility, over the five years to the end of 2014. For factor exposures, there are only minor differences between the approaches.

The sector weights of the local currency minimum volatility index relative to those in the standard index are more revealing. On average, in both simulations, the local currency index is overweight materials and the relative position in technology shows wide variation. In Exhibit 19, we see the “home bias” that favored U.S. information technology because of currency risk reduction, even if the sector is not associated with a low volatility strategy. The local currency version almost always has a maximum underweight in technology.

Exhibit 19: Active Sector Exposures over Five-Year Simulation

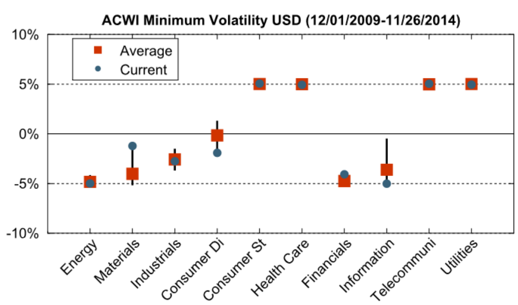

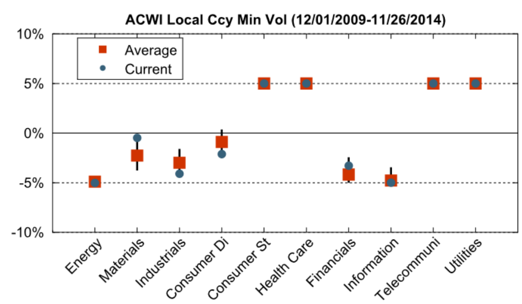

With a different parent index, the average sector difference can be more marked. When we compare an MSCI Japan Minimum Volatility index in yen to one optimized in U.S. dollars, the latter has a much higher *average* weight to consumer discretionary stocks. This overweight is plausible, given the sector’s sensitivity to dollar-yen currency correlations. Moreover, the level of sector differences can be dwarfed by stock-level differences in holdings: One-way turnover between minimum volatility indexes optimized in different currencies can be 20%-30%.

Relative country and regional exposures showed wider variation than sectors for both simulation periods. Without the impact of currency volatility and correlations, the country positioning for the local currency index is more stable, with active country constraints clearly binding and the “home bias” reversed.

Exhibit 20: Active Country Exposures – Over Five Years

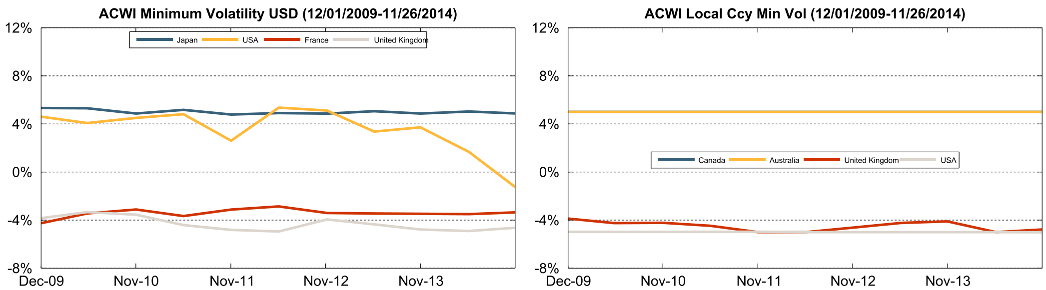

We also aggregated the percentage country weights into currency blocks and reviewed the time-series of exposures in turn for a range of reserve, developed and emerging market currencies (Exhibit 21). We see the home bias in the U.S. dollar exposure when the dollar is the base currency and the increase in yen exposure as the negative currency-equity correlation increases. Exposures for the euro (and sterling) are, however, similar. The commodity currencies generally see more stable currency positions in the local currency minimum volatility index.

Exhibit 21: Active Currency Exposures

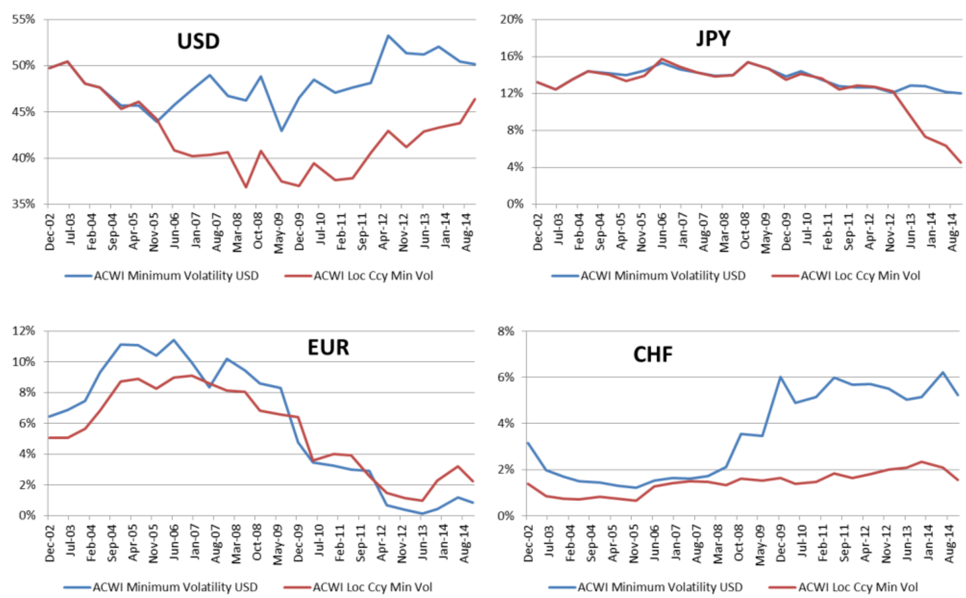

Simulation period: 11/29/2002 to 12/31/2014

The key performance characteristics of the minimum volatility indexes are those linked to the risk-driven index objectives, e.g., how do total risk, beta and tail-risk properties such as maximum drawdown change? The difference in sector, country and factor exposures (and for the unhedged index, the currency weights) already indicates some of the exposure to tail events. We also looked at the slippage between the returns on the two indexes (hedged and unhedged) because short-run drawdowns versus a benchmark or an alternative strategy can lead to investor regret. In Exhibit 22, we show the relative risk reduction view for U.S. dollar returns, while in Exhibit 23, we show the comparison for hedged returns (using the local currency return as a proxy for a "perfect hedge" overlay). To aid comparability, we show in each bar the percentage reduction in that risk measure compared to the parent ACWI index.

Exhibit 22: Percentage Risk Reduction in Simulation for Alternative MV Indexes

(Left-hand chart: 12-Year simulation; Right-hand chart: 5-year simulation)

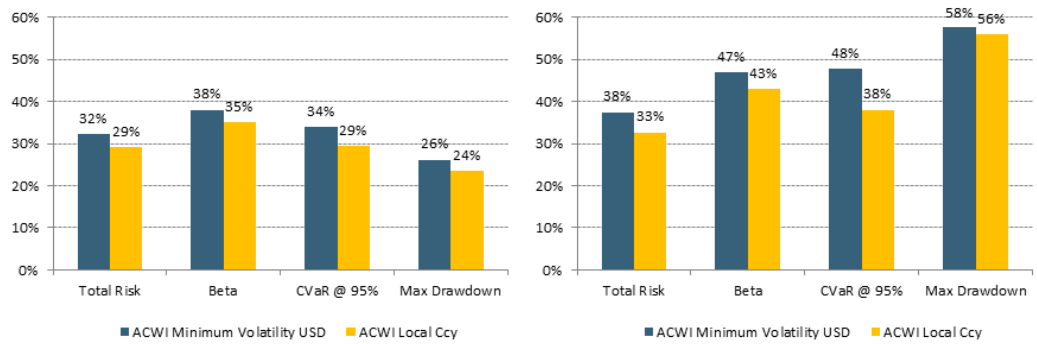

Exhibit 23: Percentage Risk Reduction in Simulation for Alternative Hedged MV Indexes

(Left-hand chart: 12-year simulation; Right-hand chart: 5-year simulation)

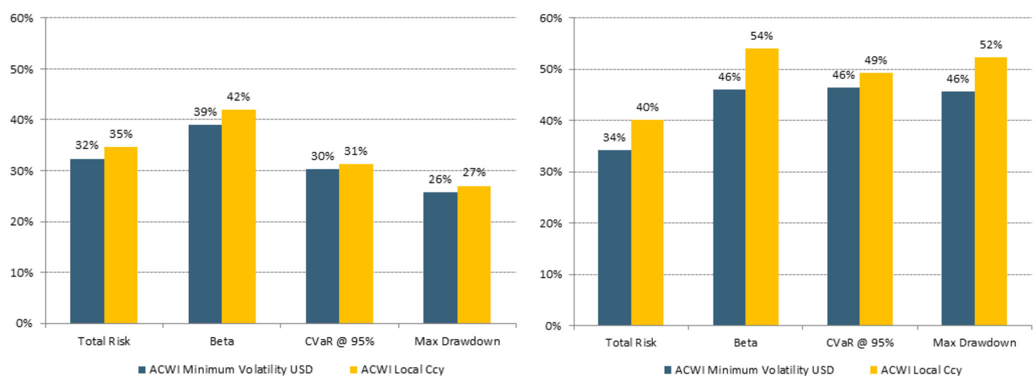

The performance difference coming from the varying exposures of a local currency approach can be episodic. Recently, as FX volatility has risen, the divergence has generally been at its greatest – the full-period (unhedged) tracking error of 1.4% is low compared with the most recent period and is a weak guide to performance dispersion (see Exhibit 24). While for unhedged returns the USD-optimized minimum volatility index offered the greater risk reduction, once hedged indexes are considered, it is the local currency optimized version with the hedge overlay that has shown greater relative risk reduction (especially over the last five years). But such differences are not so apparent in every parent index minimum volatilty simulation.

Exhibit 24: Tracking Risk and Return Characteristics for ACWI Local Currency MV Index

	ACWI Min Vol (Loc Ccy), 5y	ACWI Min Vol (Loc Ccy)	ACWI Min Vol (LC, hedged) 5y	ACWI Min Vol (LC, hedged)
Tracking Error (%)	1.8	1.4	1.5	1.2
Ann. act. retn (simulated)	-1.0	-0.2	-0.5	-0.1
Historical Beta	1.07	1.05	0.9	0.95
Max Active Drawdown (%)	6.7	6.7	3.7	4.0
Active Drawdown (mths)	44	44	38	69

CONTACT US

clientservice@msci.com

AMERICAS

Americas	1 888 588 4567 *
Atlanta	+ 1 404 551 3212
Boston	+ 1 617 532 0920
Chicago	+ 1 312 675 0545
Monterrey	+ 52 81 1253 4020
New York	+ 1 212 804 3901
San Francisco	+ 1 415 836 8800
Sao Paulo	+ 55 11 3706 1360
Toronto	+ 1 416 628 1007

EUROPE, MIDDLE EAST & AFRICA

Cape Town	+ 27 21 673 0100
Frankfurt	+ 49 69 133 859 00
Geneva	+ 41 22 817 9777
London	+ 44 20 7618 2222
Milan	+ 39 02 5849 0415
Paris	0800 91 59 17 *

ASIA PACIFIC

China North	10800 852 1032 *
China South	10800 152 1032 *
Hong Kong	+ 852 2844 9333
Mumbai	+ 91 22 6784 9160
Seoul	00798 8521 3392 *
Singapore	800 852 3749 *
Sydney	+ 61 2 9033 9333
Taipei	008 0112 7513 *
Thailand	0018 0015 6207 7181*
Tokyo	+ 81 3 5290 1555

* = toll free

ABOUT MSCI

For more than 40 years, MSCI's research-based indexes and analytics have helped the world's leading investors build and manage better portfolios. Clients rely on our offerings for deeper insights into the drivers of performance and risk in their portfolios, broad asset class coverage and innovative research.

Our line of products and services includes indexes, analytical models, data, real estate benchmarks and ESG research.

MSCI serves 98 of the top 100 largest money managers, according to the most recent P&I ranking.

For more information, visit us at www.msci.com.