Quantitative Strategy

Global Markets Research

Quantitative Strategy
North America United States

Deutsche Bank

9 February 2011

Portfolios Under Construction

Minimum Variance: Exposing the "magic"

Minimum Variance Portfolios
The minimum variance portfolio has been touted as the "magic" formula that generates superior risk-adjusted returns. We analyze and dissect the strategy to discover its secrets and vulnerabilities and propose an enhancement that works.

The Minimum Variance Portfolio Strategy

The minimum variance portfolio strategy is not a low-risk strategy
The common perception that the minimum variance portfolio strategy is a low-risk strategy is not entirely true. We find and show that the short-side of the portfolio can sometimes turn into a high-risk stock strategy.

Minimum variance portfolios are great diversifiers
We find some fascinating and useful properties about minimum variance portfolios. In particular, we show why they tend to possess very little correlation with traditional quant strategies.

Constraints on beta, styles and industries DO NOT work
Similar to other research, we find that adding constraints or neutralizing the portfolio against beta, style and industries does not work. However, we show why we find this is the case and show that neutralization could only serve to change the nature of the strategy.

An enhancement that works
We propose a slight and simple enhancement to the strategy, which offers protection for one of its vulnerabilities, while not veering the strategy from its intended purpose.

Adding it to an alpha-strategy
We show that adding the minimum variance strategy or our enhancement to an alpha-model can add significant value in terms of risk-adjusted performance. We attribute to the diversification benefit of the strategy.

Source: GettyImages.com

Team Contacts

Miguel-A Alvarez
Strategist
(+1) 212 250-8983
miguel-a.alvarez@db.com

Yin Luo, CFA
Strategist
(+1) 212 250-8983
yin.luo@db.com

Rochester Cahan, CFA
Strategist
(+1) 212 250-8983
rochester.cahan@db.com

Javed Jussa
Strategist
(+1) 212 250-4117
javed.jussa@db.com

Zongye Chen
Strategist
(+1) 212 250-2293
john.chen@db.com

Deutsche Bank Securities Inc.

Table of Contents

Letter to our readers

In a break with tradition, let us begin with what this report is not about. First, it is not a sales pitch. Our aim is not to convince the reader that the minimum variance portfolio is a great investment strategy in the meager attempt to sell an index or an investment product. In a similar spirit, we will not bore the reader with endless optimization results under different settings and ranging over mountains of statistically treated covariance matrices that have been tailored to work well in backtests. In addition, we will not "hard" or "soft" tilt toward a set of arbitrary generic alpha factors in search of the best combination and then justify the pairing after the fact.

Instead, our goal in this report is to dissect the strategy, break it up into its essential components, and understand its inner workings. This way we uncover its true nature, its intentions, its strengths and most importantly its vulnerabilities. In this quest, we will find a treasure trove of valuable insights about the strategy. In addition, we will uncover some useful and fascinating properties of the strategy that will help explain some of the puzzling results that are often encountered in the literature as well as in practice.

The report is quite extensive and is divided into three main parts. First, we begin by investigating the properties of the generic strategy under ideal circumstances; that is under no leverage target, trading costs, turnover limits or other typical constraints used by active managers. The goal is to identify the properties of the plain vanilla strategy; risk, return, beta, sector and style exposures, turnover and performance decay among others. These results will ground our expectations and reveal some interesting properties about the strategy. However, the empirical results can only get us so far. To really understand the strategy we have to study it analytically.

In part two, we put the strategy under a microscope and study it from an analytical perspective. The analytics will lead to various important insights and results. We start by using a simple and intuitive characterization of the strategy that will show a questionable alignment of the portfolio that arises from minimizing portfolio risk. We also find a bounty of analytical results that will reveal some useful properties and help explain many of the anomalies found in some of the empirical results. Last this analytical investigation will provide a foundation for improving the strategy.

In the last part of the report, we find out why many of the typical neutralization techniques only serve to deteriorate performance and end up positioning the strategy along the attribute it that was intended to be neutralized. We also show a slight enhancement to the strategy that improves performance without veering it too far from its original purpose. The enhancement is based an analytical finding and is not an attempt to exploit a result found in the empirical discovery. Last, we piggyback on the analytics to show the benefits and the effects of incorporating the minimum variance portfolio and its enhanced version into a traditional alpha strategy.

Readers who are familiar with the descriptive statistics of the strategy may start by reading the first section of the first part, then skip to the second and third parts of the report, which center more on the analytics and inner workings of the strategy.

As usual, comments or feedback are appreciated!

Yin, Rochester, Miguel, Javed and John.

Dissecting Minimum Variance

Why the interest?

The minimum variance portfolio (MVP) has been around for quite some time. It arises naturally as the left-most portfolio on the efficient frontier and in simple terms can be thought to be the fully-invested portfolio with minimum risk. Indeed, one of the appealing characteristics of this portfolio is that its construction does not depend on any stock return forecasts and therefore, it is sometimes perceived as a passive investment strategy[1]. In the end, the construction of the portfolio is solely dependent on the asset universe and its variance covariance matrix of returns.

Interest in the MVP strategy has increased over the years. Surprisingly, the popularity and acceptance of these strategies has extended beyond the quantitative community to the point that index and risk model providers are offering minimum variance indices[2] for domestic and international markets.

This sudden popularity is most likely attributable to the strategy's reported performance and perceived risk control. Backtests abound that show the risk-adjusted performance of the MVP to be superior to that of the market portfolio as well as many quant strategies. In addition, recent performance has been moderate, surpassing that of the market and a few heavily used quant factors that were darlings in the past, but have performed miserably in recent times. Figure 1 shows the cumulative market return for the Russell 3000 versus the MVP constructed using the Axioma's US medium horizon risk model and scaled at each point in time to have the same volatility of the market portfolio.[3]

Figure 1: Russell 3000 vs MVP[4] cumulative returns (Jan 1995 – Dec 2012)

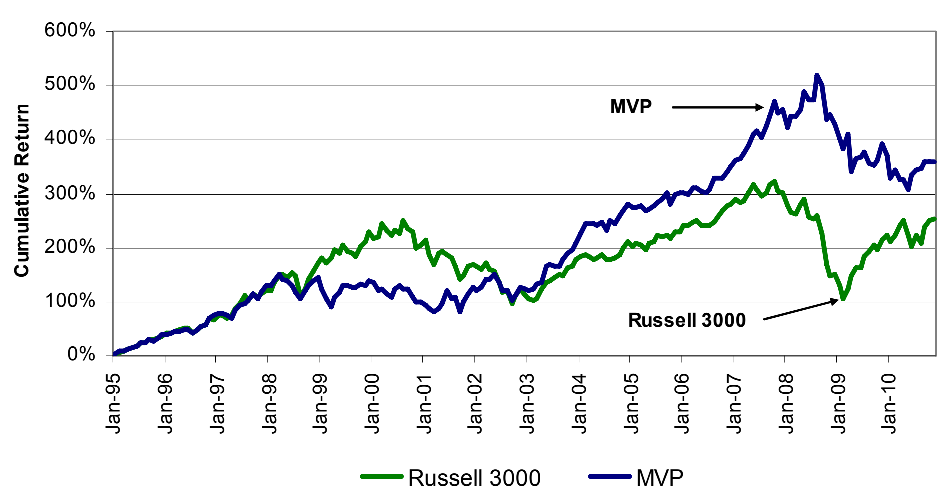

Source: Axioma,Compustat,Russell,Deutsche Bank

---

[1] As we will see below when we dissect the strategy this is far from the truth.

[2] MSCI has a set of minimum risk indices constructed from Barra risk models and Axioma is constructing a set of indices based on its own models.

[3] The MVP has much less risk than the market portfolio and so to effectively compare the two portfolios we put them on an equal ex-ante risk level. This gets at the question: "what would my performance be if I took the same risk in the MVP?" The volatility scaling is done using the 24-month trailing standard deviation of returns for each series.

[4] Minimum variance portfolio scaled to match Russell 3000 portfolio volatility using trailing 24-month standard deviation. The universe is taken to be our DB Quant Universe of stocks, which is roughly comparable to the Russell 3000.

In Figure 2, we show the risk-adjusted performance (information ratio) of the MVP, the market portfolio[5], and a set of mean-variance optimal (MVO) factor portfolios[6] over our DB Quant Universe (roughly the Russell 3000). In addition, we provide an equally-weighted quant strategy (Eqw Quant) that takes equal weight across the quant factors in the list. Note that when compared to some of these traditional quant factors, the MVP is actually quite a powerful strategy. In particular, during the entire period (Jan-1995 to Dec-2010) the MVP risk-adjusted performance ranked higher than the market and many portfolios constructed using conventional alpha-factors. This paper performance is quite an achievement given the strength of many of these quant factors experienced in the past.

As quants, we know that the world has changed; particularly after the summer of 2007. Conventional factors such as FY1 EPS revisions that once outperformed are now being classified as risk factors[7]. Because of this shift in the quantitative environment, we are always careful to check our empirical results over the more recent period (starting in Jan 2007). Figure 3 shows the same risk-adjusted results from Jan 2007 to the present. Note, FY1 Earnings Yield and Momentum[8] have lost their predictive power to the point where the extreme differences in factor portfolio IR, including that of the MVP. However, the MVP is still moderately positive and outperforms half of the traditional quant factors in our analysis.

Figure 2: IR: MVP, Market and Factor portfolios DB Quant Universe, (Jan 1995 – Dec 2010)

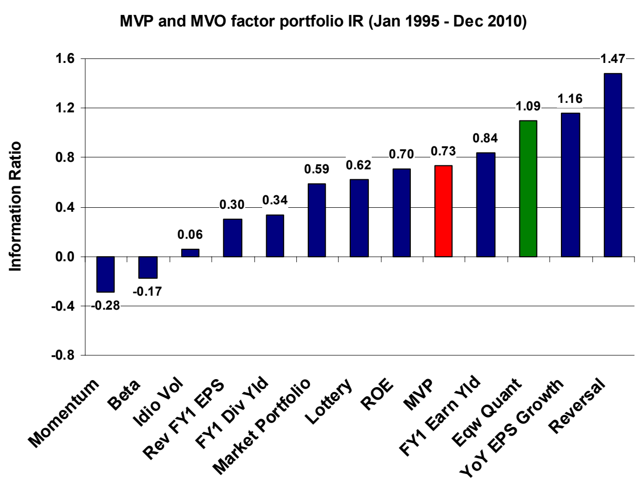

Source: Axioma, Comustat, IBES, Russell, S&P, Deutsche Bank

Figure 3: IR: MVP, Market and Factor Portfolios DB Quant Universe, (Jan 2007 – Dec 2010)

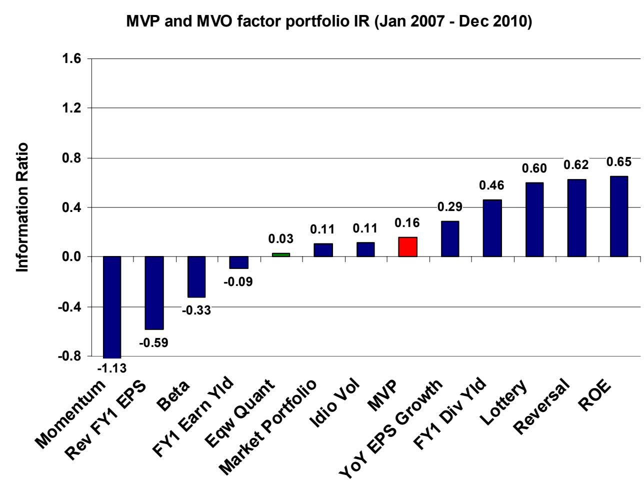

Source: Axioma, Comustat, IBES, Russell, S&P, Deutsche Bank

While the MVP results above look promising, we must remember that these statistics represent paper performance under ideal (unrealistic) conditions. The portfolios above are long/short constructed without applying any conventional constraints and more importantly, without taking into account the cost of borrowing and overall trading on performance. In the sections below, we will dissect the MVP strategy and look at some of its most important

[5] The market portfolio is the capitalization weighted portfolio of the universe.

[6] We use the characteristic portfolios for each factor (see Grinold & Kahn 1999). The characteristic portfolios have unit exposure to the alpha-factor and minimum risk. As we will see below the MVP is equivalent to the characteristic portfolio for the alpha-factor consisting of the ones vectors.

[7] Using advanced statistical techniques, our prior research (Luo *et al.*, 2010, "Launching Quantitative Strategy") shows that the unconditional mean of these factors has indeed trended to zero, while their volatility has not; implying that the only way to benefit from them is to time them correctly (see Luo *et al.*, 2010, "Style Rotation").

[8] In reality, it is unfair to show the mean-variance optimal portfolio of the Momentum factor. The reason is because the risk model has a very similar factor and the interaction between the two during the optimization processes deviates the exposure in a significant way. For more details see Lee and Stefek (2006). In fact, it can be shown that Momentum may become negatively aligned with reversal, which looking at the graphs is not a great idea.

properties. The analysis will allow us to gauge the extent to which this paper performance can be realized in the real world. In addition, we will look at the sensitivity of the strategy to different universes and to different risk models.

The second reason for the popularity of minimum variance portfolios is that that the strategy is often perceived to be agnostic to asset return forecasts (i.e. expected returns do not play a part in the construction of the portfolio). This feature of the strategy is especially appealing to investors who do not believe or do not want to engage in active investing. However, if we dig into the details, the minimum variance portfolio is not the "no information" portfolio. As we will see in later sections, the strategy depends on the information in the risk model to construct the portfolio and its performance will depend on whether investors favor more diversified and lower volatility portfolios. Indeed, as is true for every strategy that doesn't allocate fully to the market portfolio, there is an implied expected return forecast.

Where is the performance coming from?

On the surface, the MVP strategy is often perceived as a low-risk or low-volatility strategy. All low-risk strategies share the characteristic that they aim to outperform from the belief that stocks with lower risk will outperform stocks with higher risk. From an academic perspective this is abhorrent. One of the central tenets in finance is that investors should expect more reward by taking on more risk. However, the data in many equity markets tell the opposite story; investors that take on less risk tend to outperform those that take on more risk. This contradiction between the data and theory has been dubbed the low-volatility anomaly. There are many papers that study this anomaly and offer hypotheses usually centered on human behavioral biases such as the irrational preference for high volatility stocks with lottery-like payoffs[9], to recent arguments related to benchmarks as a limit to arbitrage (Baker, Bradley, Wurgler [2011]). Our research has also found this to be true with other more exotic ways to measure the risk of a stock. For example, in Cahan *et al* [2010a], we looked at using option risk sentiment measures and found the same result – stocks with lower sentiment risk outperformed those with higher options sentiment risk. Similarly, in Cahan *et al* [2010b] we found that stocks with higher informational risk from high frequency data underperformed stocks with less informational risk. We also found that the low-volatility anomaly applied to stock in European (Alvarez *et al*. [2010a], "Factor Neutralization and Beyond") as well as in Canada (Luo *et al* [2010] & Jussa *et al* [2010]).

Given all the evidence, we do not argue to the validity of the low-risk anomaly and the fact that low risk stocks tend to outperform high risk stocks. However, the important question for this report is whether this is the case for the minimum variance strategy? But the first question should be whether or not the minimum variance portfolio strategy is indeed a low-risk strategy in the conventional sense. Does its performance rely on lower-risk stocks outperforming higher-risk stocks or is there something else hiding behind its complex formulation[10]? In the second part of this report, we break down the strategy and show that there are two parts to the MVP. One part is akin to a low-volatility strategy, but surprisingly the other part can actually turn into the opposite – a strategy that will actually underperform when lower volatility stocks outperform.

The difference is subtle and lies in that the strategy does not actually position itself so that low-volatility *stocks* outperform; instead it aligns itself so that lower risk *portfolios* outperform. Therefore, outperformance depends on whether investors are willing to move prices in order to minimize portfolio risk.

---

[9] See for example, Tversky and Kahneman, 1992; Mitton and Vorkink, 2007; Kumar, 2009, Barberis and Huang 2008.

[10] As we will show later, the construction process is dependent on the inverse of the asset covariance matrix.

To see this suppose that two stocks in the universe have the same risk measure. Then a conventional low-volatility strategy would not be able to discriminate one from the other as far as which one will outperform over the next period. However, the MVP strategy uses an additional risk attribute that goes beyond the individual asset level. This risk attribute measures how well the asset diversifies or lowers the risk of an objective portfolio. In essence, this extra piece of risk depends on the interdependency of an asset with the rest of the assets in the universe. In the case of using the standard deviation of a portfolio to quantify risk, this interdependency part of risk is linked to the covariance (or correlation) of the asset with the rest of the assets in the universe.

What if we had look ahead bias?

Were results in Figure 2 and Figure 3 a fluke? Can just minimizing portfolio variance produce better returns than a value-type strategy? Another way to frame the question is to ask whether investors really demand portfolios or assets that reduce risk. If so then incorporating next month's risk information should produce superior returns relative than using today's information. In other words, if risk-reduction is really important then future risk information about assets and their contribution-to-risk to portfolios should be useful in determining the future price of the assets. This is easy to see in the two asset case. Suppose that the universe consists of two stocks that are uncorrelated and have the same risk and expected return. Then using the Markowitz (1952) mean-variance utility function, an investor who wishes to gain the maximum return for minimum risk should hold an equal amount of each asset. Now suppose that a lucky investor knows next month's risk estimate for each of the assets. She can then adjust her portfolio before any other investor and switch her allocation to underweight (sell) the riskier stock and overweight (buy) the less risky stock. If other investors care about risk minimization then all else equal, they will sell the riskier asset and buy the less risky asset; in the process moving prices in the direction of the investor who had the look-ahead information. However, if risk minimization is not a concern, then the investor with the look-ahead information should not see prices move her way. Can this be tested? Yes, we can run the same backtests as those in Figure 2 and Figure 3 except this time we use the risk model corresponding to the next month. If indeed, investors care about risk minimization then we should see the look-ahead MVP outperform the non look-ahead MVP.

Figure 4: MVP and MVO portfolio IR, using the look-ahead covariance matrix (Jan 1995 – Dec 2010)

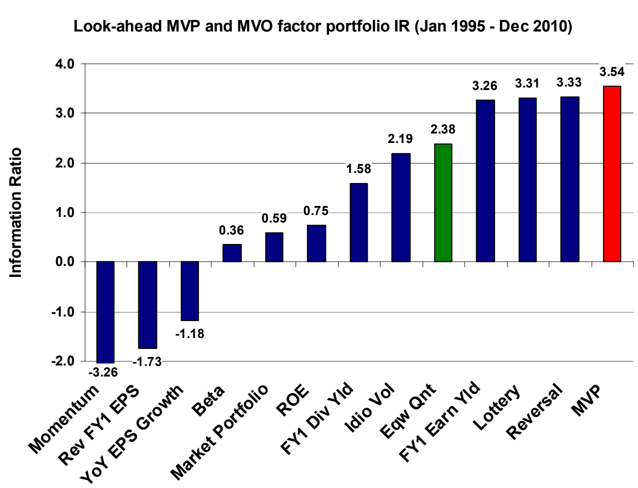

Source: Axioma, Comustat, IBES, Russell, S&P, Deutsche Bank

Figure 5: MVP and MVO portfolio IR, using the look-ahead covariance matrix (Jan 2007 – Dec 2010)

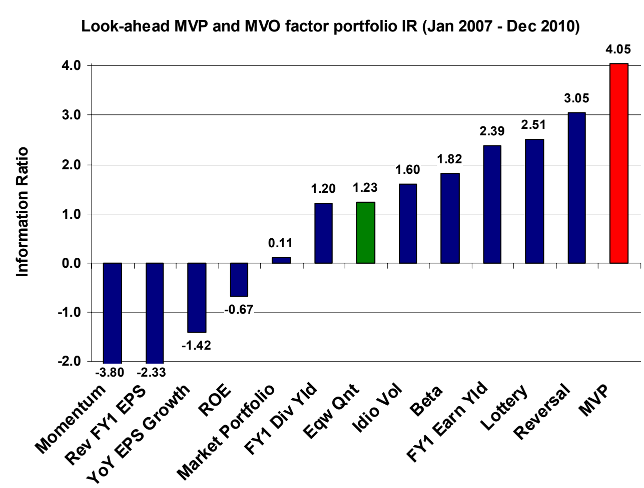

Source: Axioma, Comustat, IBES, Russell, S&P, Deutsche Bank

The results in Figure 4 and Figure 5 are quite compelling. First, they scream out for active managers to focus on forecasting risk as well as expected returns. Second and probably of greater importance to this research, the results suggest that all else equal, investors may indeed have a preference and are willing to pay a premium for portfolio risk-minimization.

One question that pops to mind is whether these results are a result of lower risk, higher return or both. If it's lower risk, then the reason for the better risk-adjusted performance may arise from a better estimate of risk (less noise), instead of investors actually demanding risk diversification. On the other hand, if this performance is due in main part to an increase in return it suggests that investors are moving towards our newly risk-minimized portfolios (or risk-adjusted factor portfolios).

Figure 6 and Figure 7 show that the risk-adjusted improvement is due mainly to an increase in return. Risk actually increases, which suggests that the look-ahead MVP does not gain any significant advantage from risk control in knowing next month's covariance matrix; rather it is profiting from the return on the forward portfolios as investors buy assets that lower portfolio-wide risk and sell-off assets that increase portfolio-wide risk.

Figure 6: MVP and look-ahead MVP mean returns

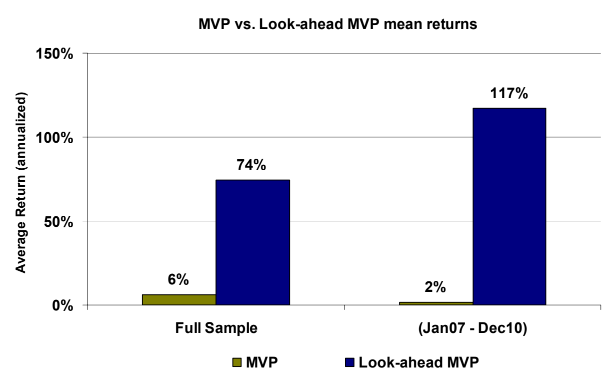

Source: Axioma, Comustat, IBES, Russell, S&P, Deutsche Bank

Figure 7: MVP and look-ahead MVP realized volatility

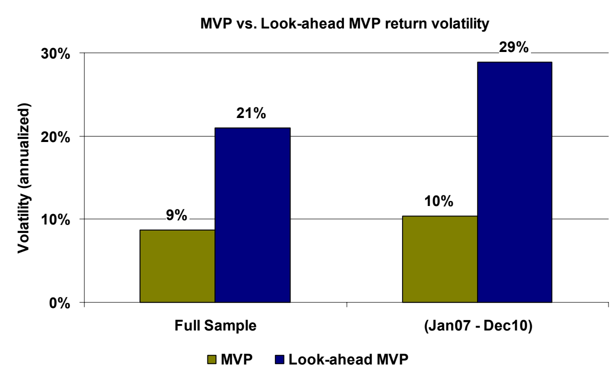

Source: Axioma, Comustat, IBES, Russell, S&P, Deutsche Bank

This finding does not lead to any new or exciting theory on asset pricing. It fact, it simply lends credence to the fact that investors seem to be accounting for risk in their portfolio decision process as Markowitz (1952) suggested in his seminal work on portfolio decision.

Alas, as logic dictates, these results only suggest that we cannot reject the hypothesis that investors move asset prices in order to reduce portfolio risk. However, the results are powerful in that they strongly suggest that future risk information is valuable even if we know nothing about expected returns.

One last comment on these results is that given generic quant factor performance since 2007, it is pertinent to point out the difference between the risk-adjusted performance of the factor portfolios in Figure 5 to those in Figure 3.

Traditional formulation of minimum variance portfolios

There are many ways to formulate and derive the minimum variance portfolio (MVP). The aim is to choose a portfolio that minimizes portfolio risk without taking any particular view on expected returns of the stocks in the universe. Below we present two ways to formulate the problem. The traditional approach presented here is a special case of the mean-variance formulation, which we will present in a later section. We feel that the mean-variance approach is more intuitive from the perspective of a quantitative investor. In addition, as we will see below, the mean-variance approach allows us to more cogently use the minimum variance portfolio strategy to enhance a total alpha model.

Traditional approach

The most common approach that is used to construct the MVP simply involves finding the portfolio with minimum variance such that the weights of the portfolios sum to one (fully invested). Mathematically, this formulation leads to solving the following optimization problem:

$$\min_{\boldsymbol{h}} \boldsymbol{h}'\boldsymbol{V}\boldsymbol{h} \quad \text{such that} \quad \boldsymbol{h}'\boldsymbol{1} = \boldsymbol{1}. \tag{1}$$

where $\boldsymbol{h}$ is the vector of portfolio weights, $\boldsymbol{V}$ is the asset-by-asset covariance matrix (risk model), and $\boldsymbol{1}$ is a vector of ones with the same dimension as $\boldsymbol{h}$. The solution to this problem is well known to be:

$$\tilde{\boldsymbol{h}}_{MVP} = \left(\frac{1}{\boldsymbol{1}'\boldsymbol{V}^{-1}\boldsymbol{1}}\right)\boldsymbol{V}^{-1}\boldsymbol{1}. \tag{2}$$

Three obvious but noteworthy observations with respect to the MVP in equation (2) are:

1) For every covariance matrix $\boldsymbol{V}$, there is a unique MVP, i.e. the MVP is sensitive to the risk model.

2) The MVP will differ with respect to the universe of stocks.

3) The portfolio is long/short. By construction the weights add to 1, but they can be positive and negative.

One way to think about the MVP is that given a covariance matrix $\boldsymbol{V}$ and *any* set of expected stock returns, the MVP corresponds to the portfolio on the left most point of the efficient frontier. For a given $\boldsymbol{V}$, this point will coincide for every efficient frontier corresponding to each expected return vector.

One not so obvious observation is that the formulation in (1) requires a constraint; otherwise the minimum variance portfolio is the zero-holdings portfolio. At the time of this writing, every publication we have seen on the topic uses the fully invested portfolio constraint (sum of weights = 1). In reality, this constraint is completely arbitrary. Our guess is that this is due to the strong academic influence on this topic and their preference to work with fully invested portfolios. However, practitioners who build portfolios on a daily basis know that a fully invested portfolio can be 350% long and 250% short, which is far from a sensible portfolio unless it matches the intended leverage. As we will see in the next section where we explore the empirical portfolio properties, we find that the MVP can indeed take these levels of exposure at different points in time. In our implementation section, we also discuss how this implies that the strategy should be run with less ex-ante risk achievable by changing the weight constraint or setting a risk target in the optimizer.

In a later section, when we formulate the problem using a mean-variance framework the constraint becomes unnecessary. In fact, it is arbitrary in the above formulation as well because the minimum variance portfolio has the property that it does not lose its optimality when constructed using a different positive net weight constraint. In other words, if we set a different net weight constraint, we end up with an equally optimal, but scaled version of the portfolio in (2). The only difference is that it will run at a higher risk (scale up) or lower risk (scale down).

Properties – let the data speak

In this section, we let the data do the talking. We report on the usual statistics: return, risk, beta, style and sector exposure. We also analyze the portfolio from a long and short perspective and look at the portfolio decile properties from which we obtain an additional layer of insight into the strategy. Last we look at the performance across different universes (large cap versus small cap) and find surprisingly that the MVP strategy, as opposed to most quantitative strategies, works well within large-cap stocks. Throughout this descriptive analysis of the strategy, we make sense of the results and report on a few anomalies that we find puzzling from an intuitive sense. We keep these anomalies in mind and later consolidate the results with the properties we find in the analytics section of the report.

Risk and return

The first thing to look at is the performance and the risk of the plain vanilla strategy. In the previous sections we looked at the risk-adjusted returns of the strategy and compared them to those of the market portfolio and a set of generic alpha-factor portfolios. Figure 8 shows the return at each point in time together with the 12-month rolling average. We could analyze the return series to death, but we note that performance was phenomenal from 1995 to 1998, then flat until 2003 and then takes off again on a long and rewarding trend until the 2008 credit crises; after-which it has not performed much better than the market or other generic quant factors. One thing to note in Figure 8 are the two drawdowns (highlighted with the red circle) that correspond to Sep 2008 and March 2009. From the perspective of a low volatility strategy the underperformance in Sep 2008 (and subsequent months) is puzzling given the fact that any low-risk strategy should have outperformed during that time. Readers that have looked at the MVP strategy may blame this on its implied beta exposure (forthcoming), but if this is the case the strategy should have outperformed during the market rally of March 2009, which the chart shows actually resulted in another drawdown. We will keep these two data points in mind as we find out more about the strategy when we look at it analytically. Figure 9 shows us the realized risk of the strategy using a 24-month rolling window. We see nothing odd or special except that the volatility seems to follow that of the general market, especially the large ramp-up in the more recent period.

Figure 8: MVP return and 12m rolling average

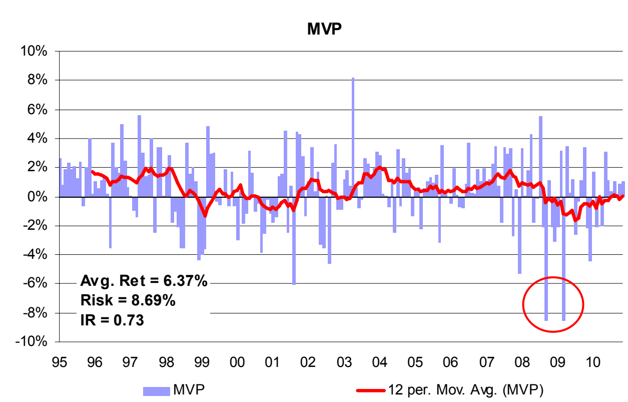

Source: Axioma, Comustat, IBES, Russell, S&P, Deutsche Bank

Figure 9: MVP std. dev. of return 24m rolling window

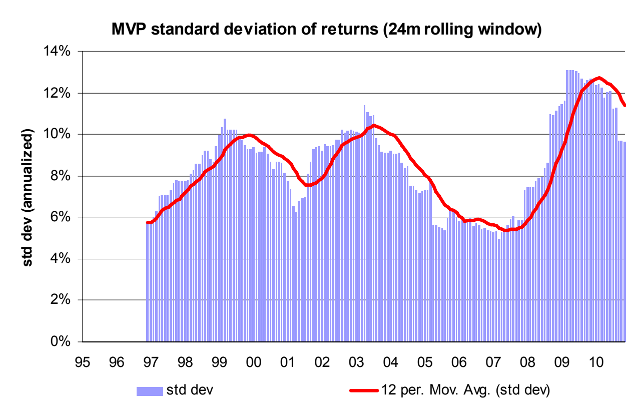

Source: Axioma, Comustat, IBES, Russell, S&P, Deutsche Bank

Another aspect of the strategy that is worth looking at is its performance in up versus down markets. Given its perceived superiority in risk control we would expect limited downside exposure relative to the market. The results in Figure 10 show the up versus down market average returns over the entire sample. From this picture, the strategy looks like a real winner in that it does well in both markets, albeit doing better in up markets. However, over the more recent period (Jan 2007 – Dec 2010) the strategy performance in up versus down market has reversed (see Figure 11). Recent performance shows the strategy to be doing better in down markets. In fact, during the recent period the portfolio has actually produced negative performance in up markets. If we look back to the performance in Figure 8 this is attributable to the significant increase in risk appetite that occurred during the spring of 2009.

Figure 10: MVP average return in up versus down markets (Jan 1995 – Dec 2010)

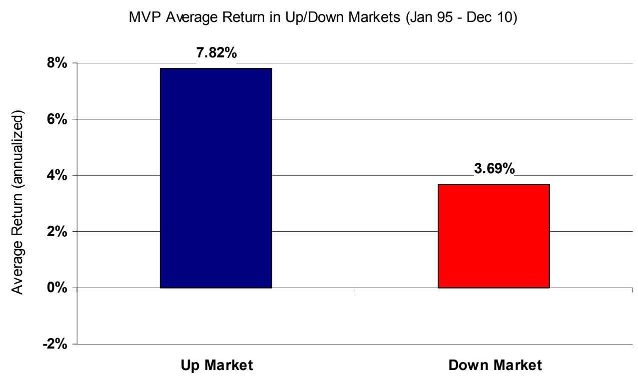

Source: Axioma, Comustat, IBES, Russell, S&P, Deutsche Bank

Figure 11: MVP average return in up versus down markets (Jan 2007 – Dec 2010)

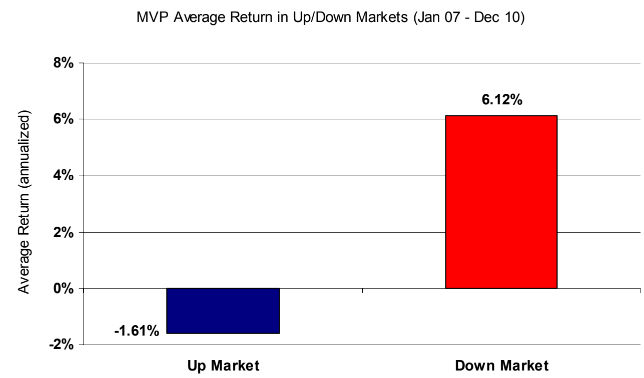

Source: Axioma, Comustat, IBES, Russell, S&P, Deutsche Bank

Long versus short

The MVP is a long/short portfolio. As we stated before, the formulation in (2) forces the weights to sum to one (fully invested portfolio), but it does not restrict the portfolio from taking short positions. This is another practical drawback of the MVP strategy. While shorting on paper many times provides better risk control and performance, taking short positions in the real world can be expensive or downright unachievable. Therefore, it is important to look at the make-up of the portfolio as far its long and short exposure. More importantly we look at the performance on the short side of the portfolio. Indeed, if most of the performance arises from the short-side then it may be likely that the performance reported earlier may be unattainable.

To investigate the long/short makeup of the portfolio, we start by looking at the total long weight (Figure 12) and total short weight (Figure 13) over time. Note from equation (2) that at each point in time the long weight minus the short weight will be equal to one. Therefore, we expect the portfolio exposure to tilt towards the long side as we see is the case.

On major observation about both Figure 12 and Figure 13 is that the fully invested constraint does not actually ensure that the MVP is a sensible portfolio in the context of leverage. Note that there are times (more recent) when the long holdings can reach 300% and the short 200%. This implies that an investor would have roughly $5 of exposure for every $1 invested[11], which is by no means a sensible portfolio. As we explained in the last section the weight constraint is arbitrary and it is better to target a desired risk level instead of running the portfolio with a net weight constraint of one.

Figure 14 and Figure 15 show the MVP long-side and short-side returns, respectively. We note that the short-side of the portfolio performed worse over the entire period and in particular took a big hit during the market rally in March 2009.

[11] Depending on how you define leverage it could imply that the portfolio runs on 5-times the leverage.

Figure 12: MVP long weight exposure

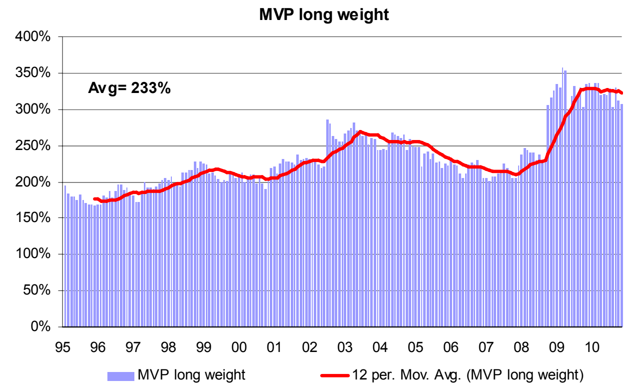

Source: Axioma, Comustat, IBES, Russell, S&P, Deutsche Bank

Figure 13: MVP short weight exposure

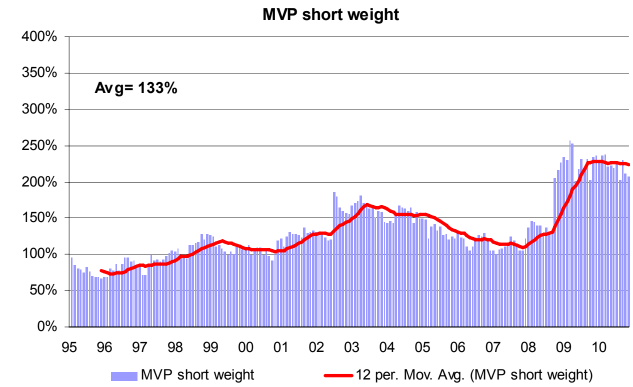

Source: Axioma, Comustat, IBES, Russell, S&P, Deutsche Bank

Figure 14: MVP long side return

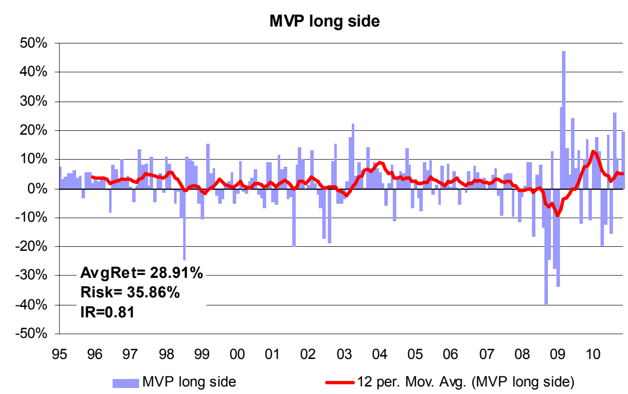

Source: Axioma, Comustat, IBES, Russell, S&P, Deutsche Bank

Figure 15: MVP short side return

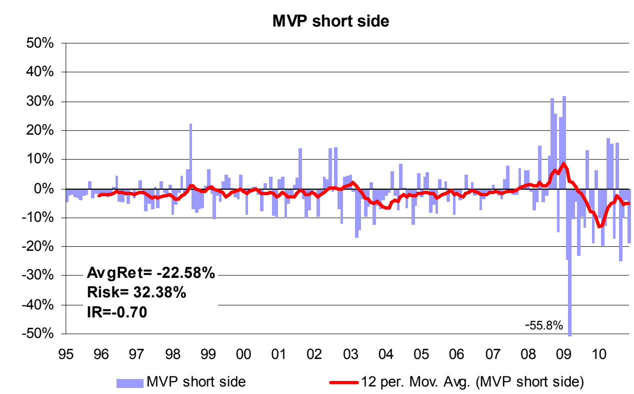

Source: Axioma, Comustat, IBES, Russell, S&P, Deutsche Bank

Decile analysis

We can shed more light on the results above by refining the performance results into the portfolio deciles. Figure 16 and Figure 17 show the average return per decile[12] across the entire sample and over the most recent period, respectively. The only particular statistics that jump out of both graphs is the difference in the magnitude of performance between the top and bottom deciles. From the pure returns perspective this is not a problem since the imbalance is in favor of positive returns.

One interesting result we found was the strategy's performance across up and down markets. We saw that over the entire period, the strategy seems to suffer less on the downside. This may be due purely to the fact that the portfolio is net long when the market was up during the sample. However, we investigate how the deciles performed during up and down markets to ensure that the storyline is correct or if there is something hiding in the background that we can't see in the long/short side analysis.

[12] In this and every other decile analysis, decile 10 corresponds to top 10% holdings of the portfolio and decile 1 corresponds to the lowest 10% of the holdings (always a short position).

Figure 16: MVP decile performance (Jan 95 – Dec 10)

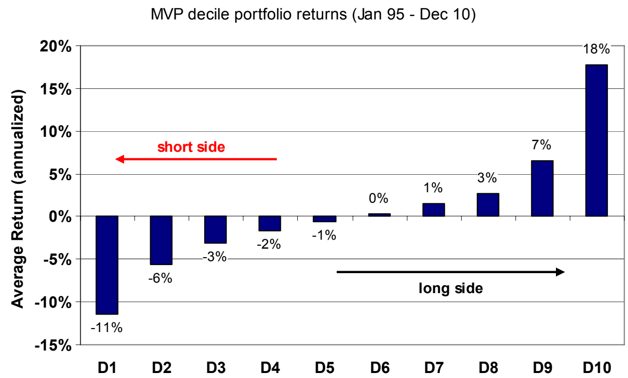

Source: Axioma, Comustat, IBES, Russell, S&P, Deutsche Bank

Figure 17: MVP decile performance (Jan 07 – Dec 10)

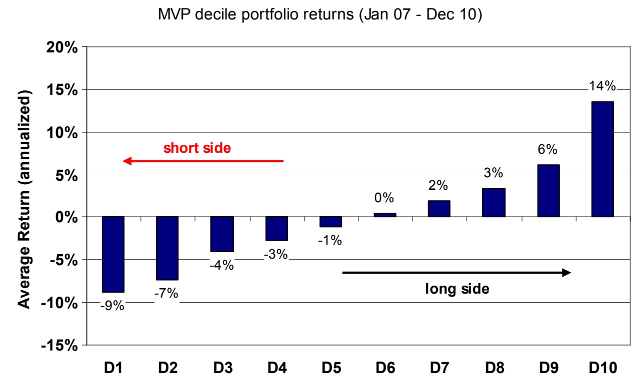

Source: Axioma, Comustat, IBES, Russell, S&P, Deutsche Bank

Figure 18 shows the decile performance of the portfolio in up markets. We see the performance consistent with a long/short strategy that is net long. However, the picture changes in down markets (Figure 19). In down markets we would expect the decile performance to reverse. That is, we would expect the short side of the portfolio (deciles 1-4) to outperform, while the long side (deciles 6-10) should underperform. However the graph shows a surprising result. In down markets the portfolio actually experienced net positive performance coming from the long side of the portfolio and particularly from the decile 10. As we will see analytically below, this is because on the long side of the portfolio, the strategy allocates higher positive weights to assets with less overall risk. These are the least likely to drop during periods of crises. In contrast, we will find that this is not the case for the short-side of the portfolio.

Figure 18: Up market decile returns (Jan 95 – Dec 10)

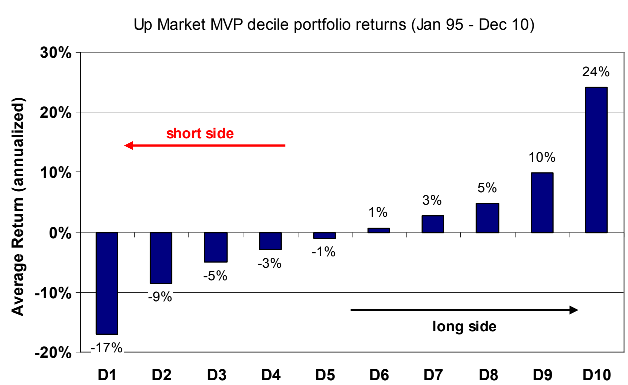

Source: Axioma, Comustat, IBES, Russell, S&P, Deutsche Bank

Figure 19: Down market decile returns (Jan 95 – Dec 10)

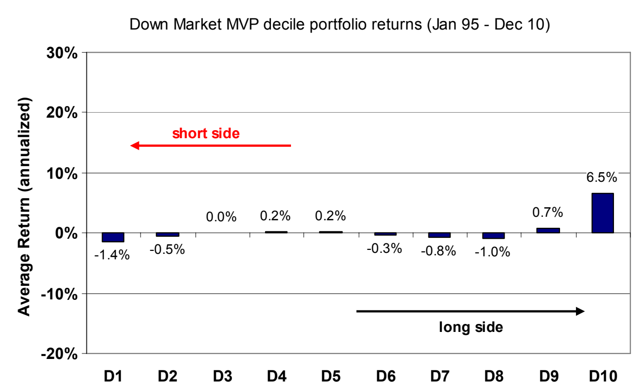

Source: Axioma, Comustat, IBES, Russell, S&P, Deutsche Bank

Note that the decile average returns are balanced except for the two extreme deciles. The top decile, which corresponds to the largest weights, outperformed the lower decile, which corresponds to the largest positive weights and the bottom decile, which corresponds to the largest (in magnitude) negative weights.

Given that this is a minimum variance strategy, the question is whether the negative return on the short-side is warranted from a risk perspective; i.e. is the short side providing risk control? Indeed, Figure 20 and Figure 21 tell us that there is less risk allocated to the shorter side of the portfolio, which explains why the magnitude of the return on the shorter end of the portfolio was of less magnitude that than on the longer side of the portfolio. However, this begs the question: why is the shorter side of the portfolio taking on less risk exposure than the longer side? A low-risk strategy usually puts on more short-exposure towards riskier assets and thus we would expect the MVP to have higher risk exposure on the short side. We will explain this phenomenon in our analytical section in the second part of the report.

Figure 20: MVP decile risk (Jan 95 – Dec 10)

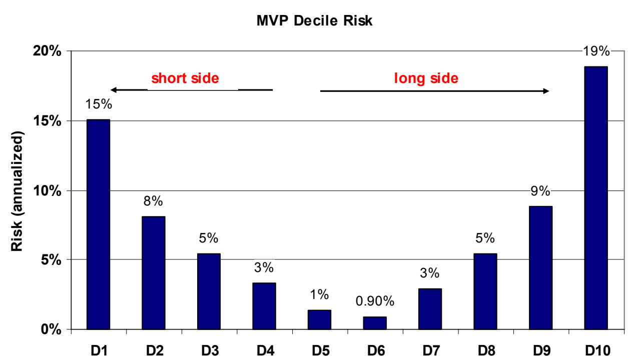

Source: Axioma, Comustat, IBES, Russell, S&P, Deutsche Bank

Figure 21: MVP decile marginal contribution to risk

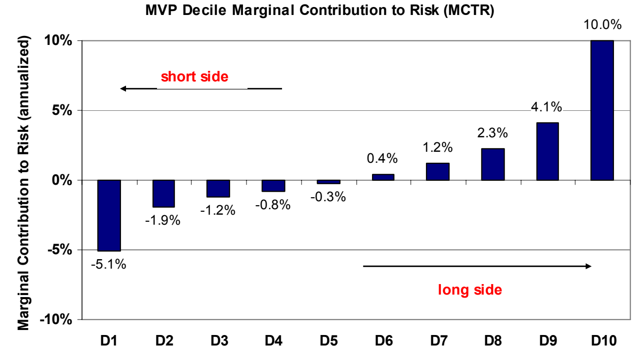

Source: Axioma, Comustat, IBES, Russell, S&P, Deutsche Bank

Beta, style and sector exposures

Perhaps the best way to empirically analyze the MVP strategy is to analyze its exposures across traditional risk factors. Given that the aim of the strategy is to minimize risk, we have certain expectations going into the analysis. First, we expect it to be low-beta. Second, we expect its exposure profile to style risk (e.g. leverage, liquidity, size, etc) to be more conservative than the average stock portfolio of the universe and the market. Last, but not least, we expect the sector allocation to tilt towards those sectors that are more conservative. We start by analyzing the exposure to Predicated Beta[13].

Figure 22: MVP Predicted Beta exposure

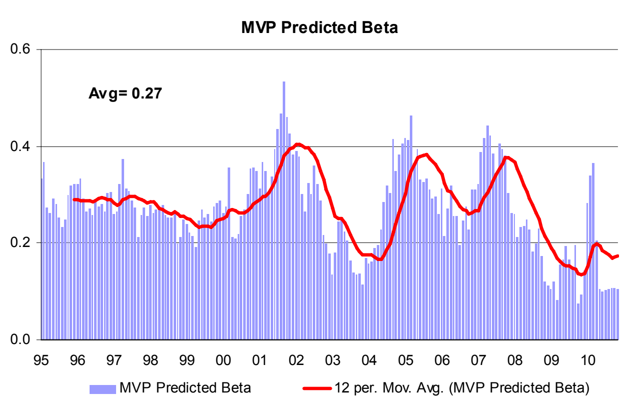

Source: Axioma, Comustat, IBES, Russell, S&P, Deutsche Bank

Figure 23: MVP long versus short side Predicted Beta

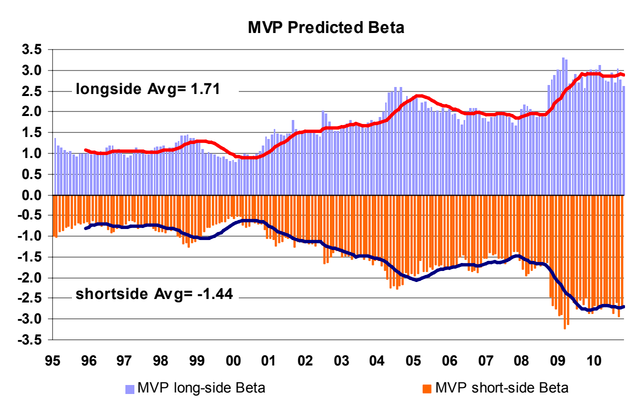

Source: Axioma, Comustat, IBES, Russell, S&P, Deutsche Bank

Figure 22 shows the total exposure to beta. We note that it is always positive, but tends to fluctuate over time. In our analytical section, we will see that the exposure to beta will vary over time with levels of asset specific risk. Figure 23 shows us the beta per side of the

[13] Predicted Beta is the beta of the portfolio computed using the US mid-horizon Axioma fundamental model.

portfolio. The only note is that the two seem to be moving in opposite direction of each other.

To uncover more insight we can break up the beta and specific risk exposures by portfolio deciles. Figure 24 and Figure 25 again point to an imbalance between the top and bottom deciles of the MVP strategy. However, note that the imbalance in beta (Figure 24) is very subtle compared to the imbalance in specific risk (Figure 25).

Figure 24: MVP Predicted Beta by Decile

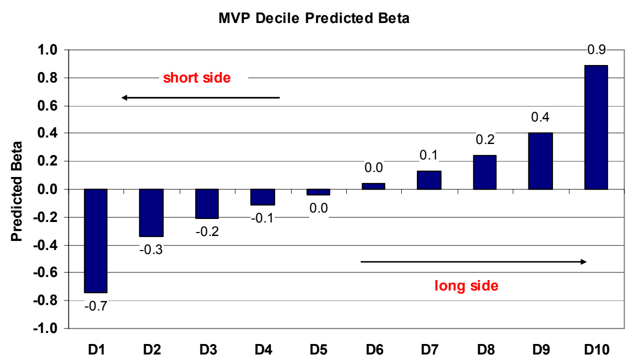

Source: Axioma, Comustat, IBES, Russell, S&P, Deutsche Bank

Figure 25: MVP Specific Risk by Decile

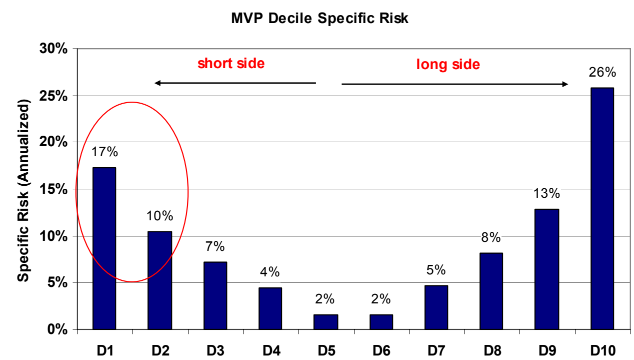

Source: Axioma, Comustat, IBES, Russell, S&P, Deutsche Bank

We also note the peculiar fact that throughout most of the sample period, specific risk has been a much stronger component of asset risk (this has subsided in the recent high correlation environment, see Alvarez *et al*, 2010 "Correlation and Opportunity"). Therefore, the large disparity between the specific risk between the top and bottom decile is suspicious given that a low-volatility strategy would conventionally allocate more short exposure to stocks with higher specific risk. As we will see in our analytical section below, this ambiguity is indeed consistent with the MVP strategy and it relates to a subtle difference between low-volatility strategies at the asset level versus those at the portfolio level.

Sector exposures

Figure 26 and Figure 27 show the MVP sector exposures versus those of the market portfolio over the entire sample period and the more recent period since January 2007.

Figure 26: MVP and Market Portfolio average sector weights (Jan 1995 – Dec 2010)

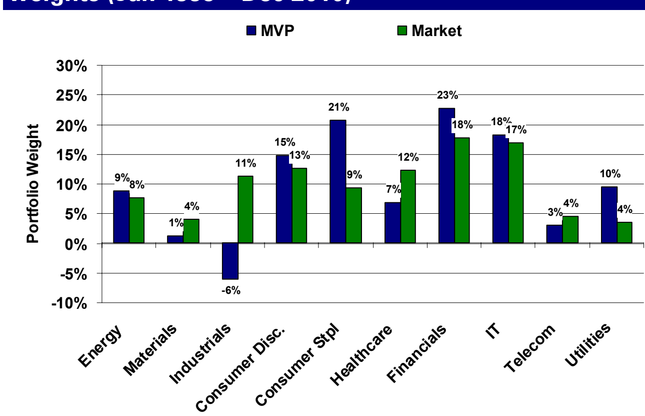

Source: Axioma, Comustat, IBES, Russell, S&P, Deutsche Bank

Figure 27: MVP and Market Portfolio average sector weights (Jan 2007 – Dec 2010)

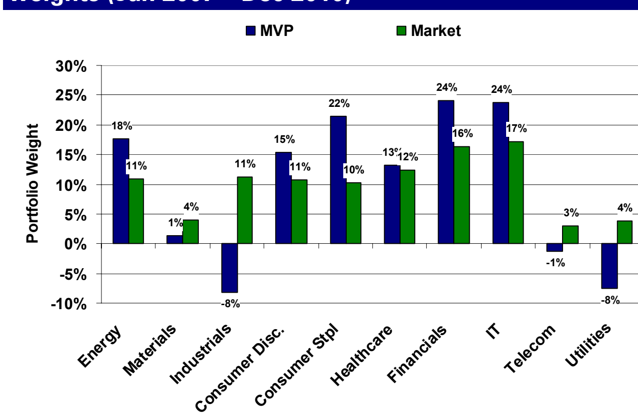

Source: Axioma, Comustat, IBES, Russell, S&P, Deutsche Bank

Figure 28 shows the industry make-up of the portfolio by deciles. There is nothing too surprising in the results; financials seem to make up much of the long and short exposure (more long than short), while the portfolio prefers to be long consumer staples and does not apply much short exposure to this sector.

Figure 28: MVP average sector weights by deciles (Jan 95 – Dec 10)

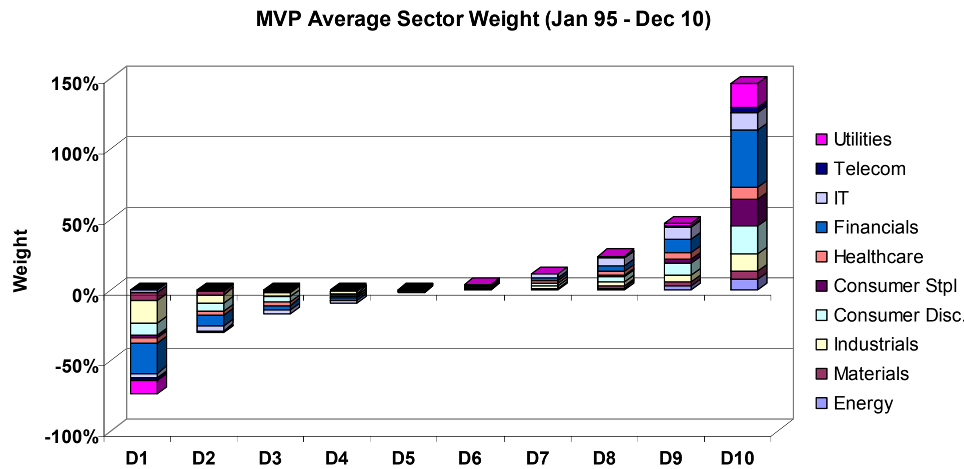

Source: Axioma, Comustat, IBES, Russell, S&P, Deutsche Bank

Style exposures

A descriptive analysis of a strategy is not complete without looking at risk factor style exposures. Figure 29, shows us the MVP style exposures to Axioma's fundamental US model and compares them to the market.

Figure 29: Axioma risk model style[14] factor exposures: MVP vs. Market Portfolio (average over Jan 95 – Dec 10)

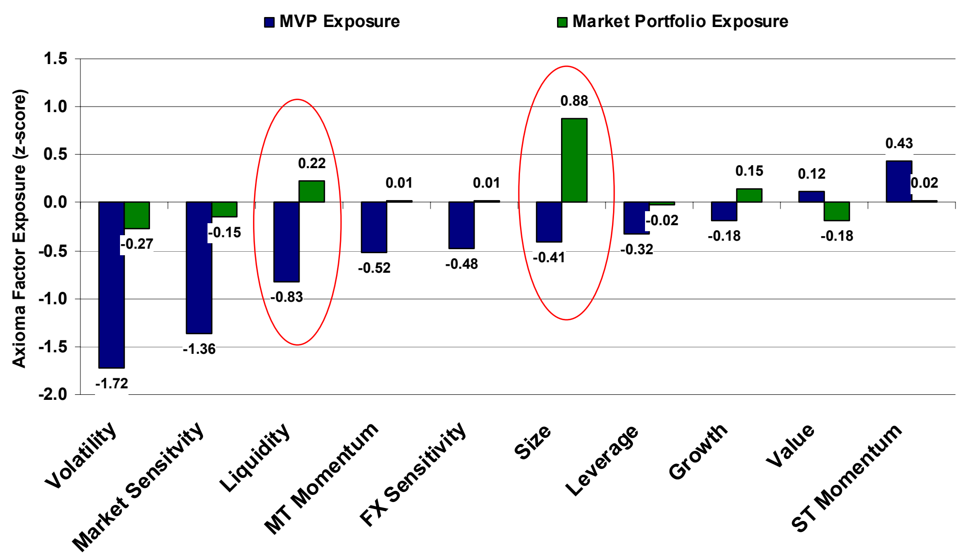

Source: Axioma, Comustat, IBES, Russell, S&P, Deutsche Bank

Note that these exposures are roughly z-scores so that 0 corresponds to the universe average (cap-weighted). Two surprising results that jump out at us in Figure 29 are the size[15] and liquidity[16] factor tilts. This is especially true when comparing to them to the factor exposures of the market portfolio. In both cases the MVP tilts towards what we would consider riskier stocks. This is contrary to what we would expect from a low-volatility strategy or one that was attempting to minimize risk above anything else. Alas, as we will see when we analyze the strategy through an analytical lens, this is consistent with strategy and arises from the way vendor risk models are usually put together. For completeness we look at the exposures to both of these factors throughout time for the MVP and compare them with those for the market portfolio (Figure 30 and Figure 31). One trend to note from both graphs is that the size and liquidity exposures of the MVP have been moving towards less risky levels since 2003.

---

[14] Style risk factors are from the Axioma US fundamental medium horizon risk model (see Axioma Style Handbook, 2010 at Axioma.com).

[15] The Axioma Size risk factor is computed by scoring each stock according to log(average market cap over last 20 days) and then normalizing the scores (z-scores) across the estimation universe.

[16] The Axioma Liquidity factor is computed by scoring each stock according to the 20-day average daily volume (in currency not shares) divided by the stock's 20-day average market cap. The scores are then normalized (z-scored).

Figure 30: Size risk factor exposure: MVP & Market

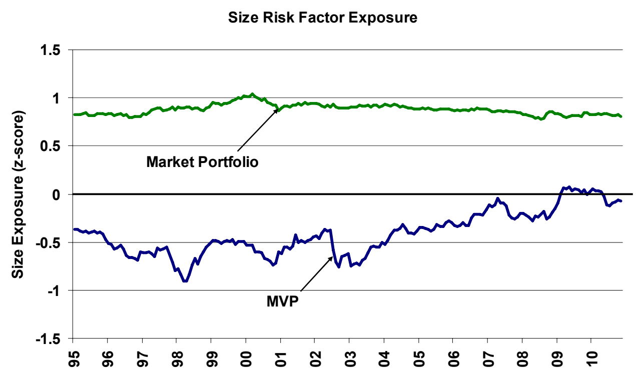

Source: Axioma, Comustat, IBES, Russell, S&P, Deutsche Bank

Figure 31: Liquidity risk factor exposure: MVP & Market

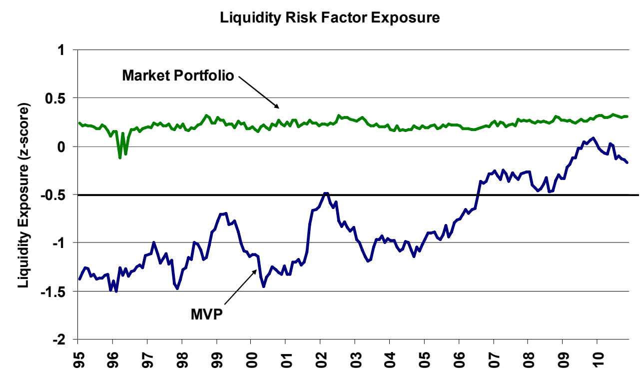

Source: Axioma, Comustat, IBES, Russell, S&P, Deutsche Bank

Last, we look at the decile exposures to these factors in Figure 32 and Figure 33. Again the results are consistent with the recurring theme that the short-side of the portfolio is not in balance with the long-side. As was the case when looking at other decile properties, the tilt is a result of the imbalance between the exposures in the top and bottom deciles. Another statistic that jumps out is that the strategy is that the bottom decile seems to be short assets which are only slightly smaller than the universe average, especially in the bottom decile[17] (Figure 32). This is odd given that a low-volatility strategy would have a stronger short position on size given that smaller assets tend to be more volatile. We find a similar result for the exposure of the bottom decile to liquidity (Figure 33). Note that the bottom decile is short assets that are only slightly more illiquid than the universe average. Again, we would expect a low-volatility strategy to be heavily underweight (or short) assets with low liquidity.

Figure 32: MVP Size risk factor exposure by decile

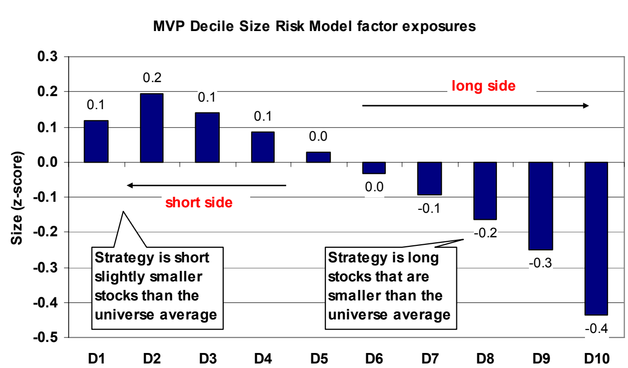

Source: Axioma, Comustat, IBES, Russell, S&P, Deutsche Bank

Figure 33: MVP liquidity risk factor exposure by decile

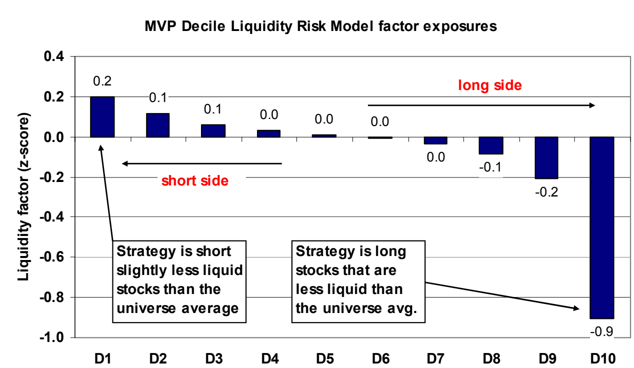

Source: Axioma, Comustat, IBES, Russell, S&P, Deutsche Bank

Correlation with other factors and portfolios

In its pursuit to minimize risk, the MVP strategy must certainly minimize common factor risk (the risk that is common with other assets and portfolios). If this is so then we should expect that the strategy should possess minimal correlation with other portfolios.

[17] Note that the exposures have been multiplied by the portfolio weights. So a positive exposure to size indicates a negative exposure to smaller assets and vice-versa. Similarly, a negative exposure

Figure 34 shows the correlation of the MVP strategy with the market as well as several quant strategy portfolios. The quant portfolios are MVO factor-portfolios (characteristic portfolios see Grinold [1999]). We note the negligible low levels of correlation with the quant factor portfolios and the significant correlation with the market. The market correlation is intuitive given the beta exposure of the factor. Figure 35 shows the returns to the z-score factor portfolios, which are computed using the z-scores directly, except that we scale them so that the resulting portfolio has the same expected risk as the MVP (this way we can put them on the same scale)[18].

Figure 34: MVP return correlation with the Market and MVO factor portfolio returns (Jan 95 – Dec 10)

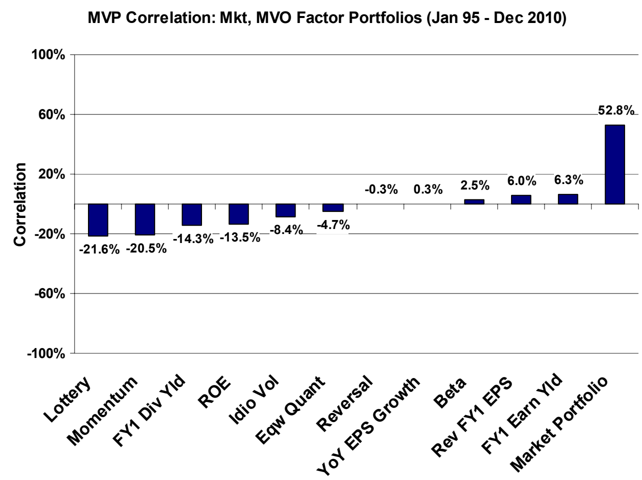

Source: Axioma, Comustat, IBES, Russell, S&P, Deutsche Bank

Figure 35: MVP return correlation z-score factor portfolio returns (Jan 95 – Dec 10)

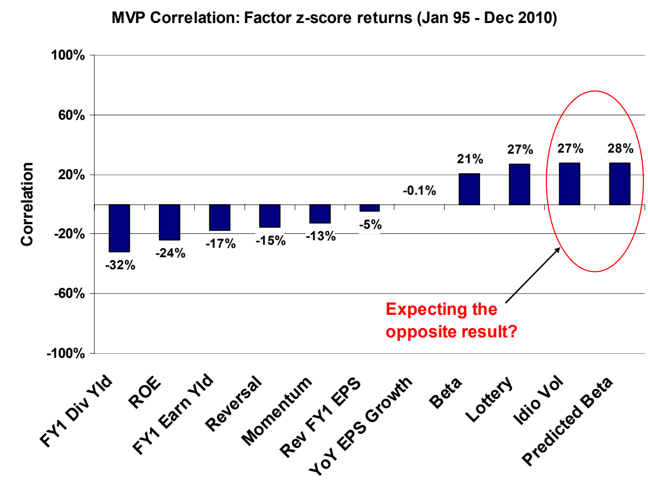

Source: Axioma, Comustat, IBES, Russell, S&P, Deutsche Bank

As we will see in the analytical section below, minimizing common factor risk is indeed the one of the strategy's main goals and this explains the low levels of correlation with the quant MVO and z-score portfolios.

Figure 36: Return correlation (24m rolling window) of the MVP and the Market portfolio

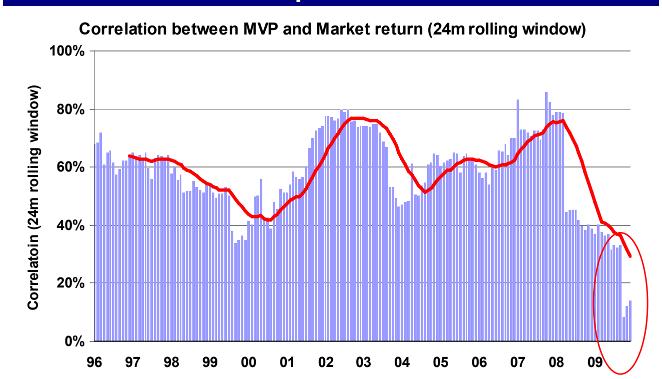

Source: Axioma, Comustat, IBES, Russell, S&P, Deutsche Bank

Figure 37: Return correlation (24m rolling window) of the MVP and a high minus low Beta portfolio

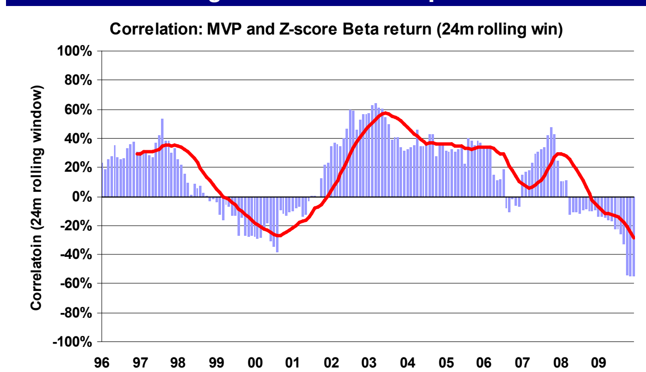

Source: Axioma, Comustat, IBES, Russell, S&P, Deutsche Bank

---

[18] However, when looking at the correlation the scaling does not matter.

Turnover

To implement the strategy in practice, we must examine both the portfolio turnover and its performance decay. Obviously, more portfolio turnover induces more trading costs and if turnover is too large, then the paper performance we saw earlier is likely to be unattainable. The turnover across a period is defined by the sum of the absolute changes in stock weights from period to the next. The more changes in weights induced by the strategy from month to month the more turnover will be reported. However, the turnover depends heavily on the actual magnitude of the weights (the same percentage change will cause portfolios with larger absolute weights to show more turnover). In addition, all else equal, turnover is strongly driven by the volatility of the strategy as well. This is for two reasons. First, all else equal, a portfolio with higher volatility will require larger asset weights. Second, higher volatility will produce larger shifts in the portfolio weights over time. Therefore, to compare the turnover of the two strategies fairly, we have to compare the two strategies at the same volatility levels. By doing this, we are comparing the strategies according to their turnover per unit of risk. Is this fair? Yes because we can always scale down a strategy to lower its risk – reducing turnover at the same time. As we will discuss in the implementation phase this involves formulating and implementing the strategy using a different net weight constraint (a smaller one).

Figure 38, shows the one-way portfolio turnover for the MVP strategy. Needless to say that the numbers over time are eye-popping! This certainly doesn't look like a low-risk strategy (risk is a sticky attribute) and it certainly doesn't resemble a passive strategy. The problem as we hinted to earlier lies in the true size of the portfolio; not the net weight, but the sum of the absolute value of both the long and short weights.

To get a better idea of our true turnover, we can compare the MVP to well-known generic factor portfolios. However, in order to fairly compare the two we need to put them on the same risk scale. As we said before, this is equivalent to comparing their turnover per unit risk. Figure 39 shows the average turnover of the MVP strategy against our generic set of MVO alpha-factor portfolios. To put them on the same risk scale, we actually scale each of the alpha-factor portfolios to have the same ex-ante risk of that of the MVP. The numbers show what we had first suspected: that the MVP strategy indeed has less turnover (per unit risk) than the alpha-factor portfolios.

Figure 38: MVP turnover (one-way)

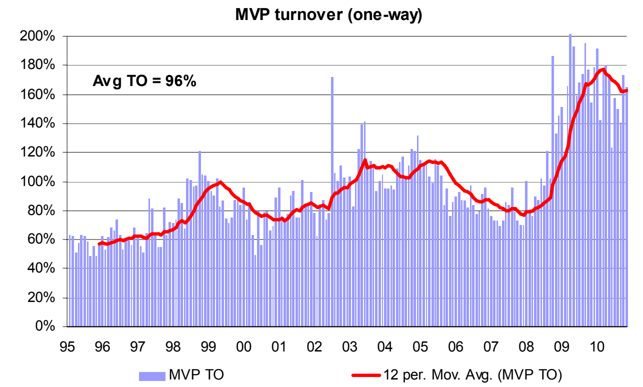

Source: Axioma, Comustat, IBES, Russell, S&P, Deutsche Bank

Figure 39: MVP vs. MVO factor portfolio[19] turnover (one-way). For fair comparison, the factor-portfolios are levered to have the same ex-ante risk as that of the MVP

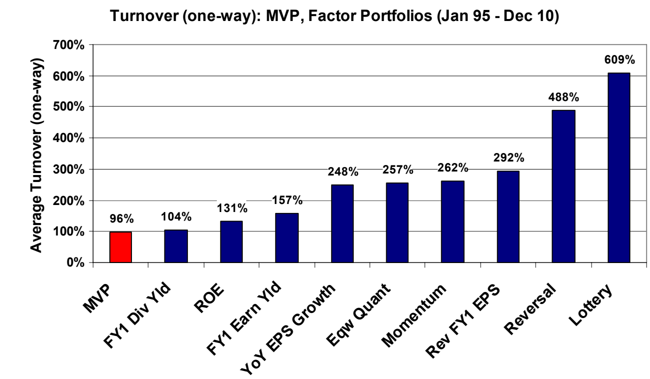

Source: Axioma, Comustat, IBES, Russell, S&P, Deutsche Bank

Performance decay

In the real world, the turnover reported in the previous section would limit the ability of an investor to capture the actual holdings of the MVP at every point in time. Most often a quantitative investor will use an optimizer with a turnover constraint (or transaction cost penalty in the objective function of equation 2), to trade their portfolio into the latest MVP implied by the newest forecast of the variance covariance matrix. The analysis over the entire period (Figure 40) is very promising indeed. The results are representative of a very robust strategy with little performance decay over time. This suggests that the strategy may be of even more value than noted in Figure 2 in that it does not require excessive trading costs to capture the performance. In addition, these results suggest that the excessive portfolio turnover reported in Figure 38 may be inconsequential to capturing performance.

Figure 40: MVP return decay profile (Jan 95 – Dec 10)

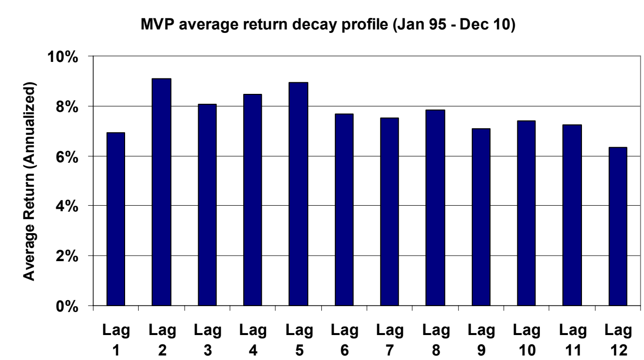

Source: Axioma, Comustat, IBES, Russell, S&P, Deutsche Bank

Figure 41: MVP return decay profile (Jan 07 – Dec 10)

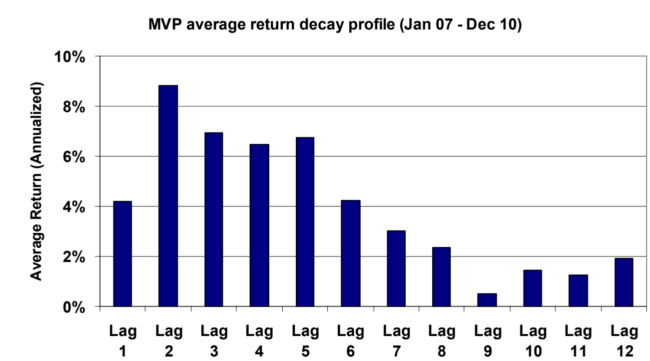

Source: Axioma, Comustat, IBES, Russell, S&P, Deutsche Bank

While not as impressive as what we saw over the entire sample, the decay results over the more recent period (Figure 41) are still favorable; especially relative to most quantitative

[19] As noted in the text, the factor portfolios are the mean-variance optimized portfolios, but scaled to have the same ex-ante risk as that of the MVP. This is done to put the portfolios on the same footing given that, all else equal, turnover will increase with the risk of the strategy. Note that we could have equivalently run the MVP portfolio at lower risk by divesting cash from the strategy (i.e. leveraging down the portfolio). An equivalent way to compare the portfolio turnover of each strategy would have been to look at the turnover per unit risk or an even more comparable method is to look at the autocorrelation of the portfolios. In all these cases, we get the same results on a relative basis.

strategies. The recent downward trend in the portfolio decay may be a result of two phenomena. First, it may be attributed to the strong increase in turnover in the MVP portfolio since 2008 (see Figure 38). This increase is most likely to be a consequence of more variability in the risk model forecast (volatility has been increasingly choppy since 2008). A second possibility for the faster performance decay may be that investors are paying more attention to minimizing portfolio risk or the low-volatility anomaly; consequently arbitraging away the performance of this portfolio minimization strategy.

One last and important observation to make about the decay is that the performance seems to peak for the portfolio two months after it is constructed. This is odd indeed. It points to the argument that investors take their time to react to changes in risk and therefore a portfolio constructed in the past works very well over the following months as well. This observation also points to the puzzling characteristic that under the case of no trading costs limiting the strategy turnover will actually improve performance!

Large-cap versus small-cap

It is well known that most generic quant strategies have historically worked better for small cap relative to large cap universes. It should be evident by now that the MVP strategy is more a quant strategy than a passive strategy so it is necessary to check its efficacy in both the large cap and small universes. We find two very surprising results. First, over the entire sample, the MVP strategy risk-adjusted performance is actually better for large-cap stocks than for small-cap (Figure 42 and Figure 43). In addition for our large-cap universe, with exception to the one-month Reversal factor, the MVP strategy outperformed every other generic quant factor in our study including the equally weighted quant strategy. This is quite an impressive feat given that many of these factors have had very good performance pre-2007.

The fact that the strategy has worked well for large-cap stocks over the entire sample is very interesting. It tells us that portfolio risk minimization may be more important in large-cap universes where there may be more institutional funds, which are mandated to provide risk control. This may also be due to the fact that most vendor risk models are calibrated towards larger-cap stocks[20]

[20] Indeed, many risk model vendors using the square root of market cap (or some close variant) when cross-sectionally estimating the return to the risk model factors, which drive the covariance matrix forecast.

Figure 42: IR: MVP, Market and Factor portfolios Russell 1000 Universe, (Jan 1995 – Dec 2010)

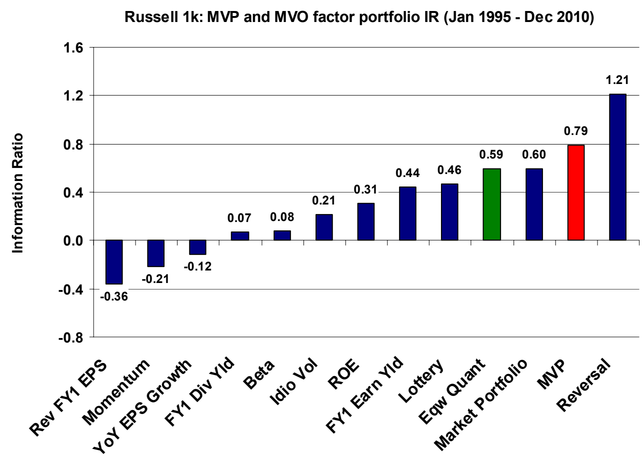

Source: Deutsche Bank

Figure 43: IR: MVP, Market and Factor portfolios Russell 2000 Universe, (Jan 1995 – Dec 2010)

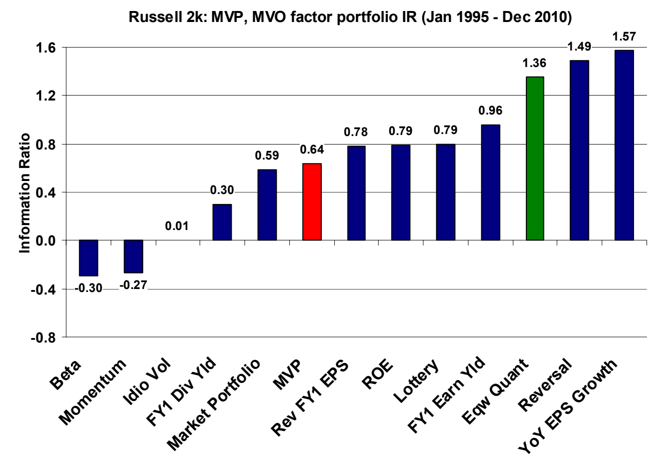

Source: Deutsche Bank

Figure 44: IR: MVP, Market and Factor portfolios Russell 1000 Universe, (Jan 2007 – Dec 2010)

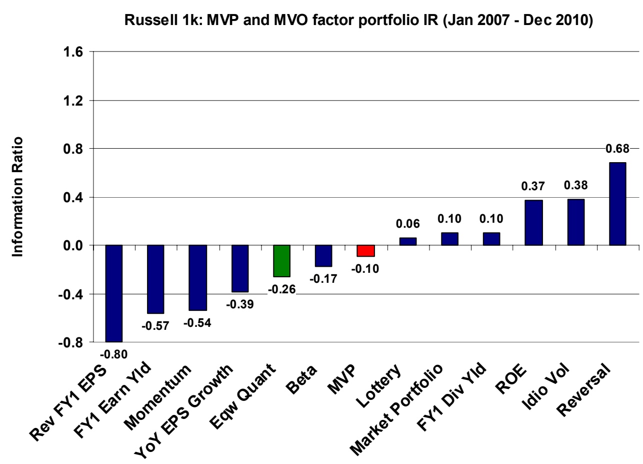

Source: Axioma, Comustat, IBES, Russell, S&P, Deutsche Bank

Figure 45: IR: MVP, Market and Factor portfolios Russell 2000 Universe, (Jan 2007 – Dec 2010)

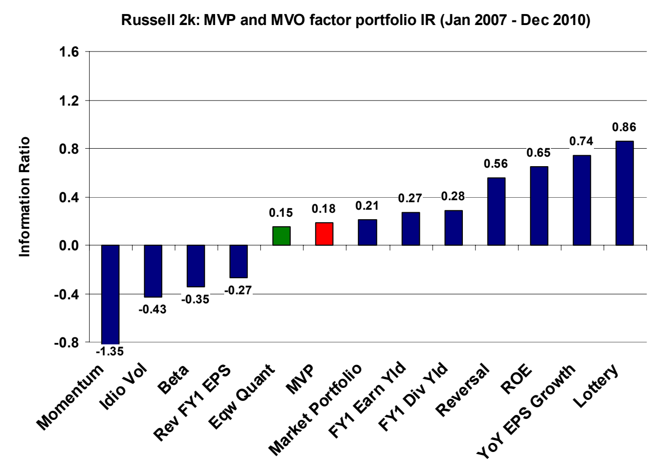

Source: Axioma, Comustat, IBES, Russell, S&P, Deutsche Bank

Last, it is worthy to note that over the recent history however, the MVP has not outperformed in large-cap relative to small-cap (Figure 44 and Figure 45) and does not seem to be a dominant strategy in the large-cap universe relative to other quant strategies.

Borrow costs for the short side

Recent research on borrow costs (see Cahan, *et* al, 2011 "The long and the short") showed that the short side of many quant factors was strongly aligned with hard to borrow stocks. In the real world, this complicates the realization of positive performance due to the high paper-performance realized on the short side of many quant strategies. We found that some of this performance could be attributed to the fact that many quant factors are aligned with shorting high volatile names (see Luo, *et al, 2010,* "Volatility=1/N"). Cahan *et al* (2011) showed that these high volatile stocks were the hardest to borrow. To this end, we want to investigate the borrow cost exposure of the MVP strategy and compare it to a conventional quant factor strategy. Given the perception that the MVP strategy is a low-volatility strategy we expect the short-side of the portfolio to have relatively high exposure to higher volatility stocks and therefore to have higher borrow costs. However, Figure 46 tells us a completely different story. When we look at the borrow cost exposure (higher exposure indicates higher costs), we find that the MVP strategy borrow cost exposure is significantly lower than that of the equal-weighted quant strategy when run at the same risk levels. Figure 47 shows us that the difference between the two has actually become more pronounced over time. Could it be that the MVP strategy is not actually taking such short exposure on highly volatile assets? We find that indeed this is the case when we dissect the strategy analytically in the second part of the report.

Figure 46: Borrow cost factor scores per decile: MVP vs. equal weight quant factor portfolio running at same risk levels (higher score implies higher borrow cost)

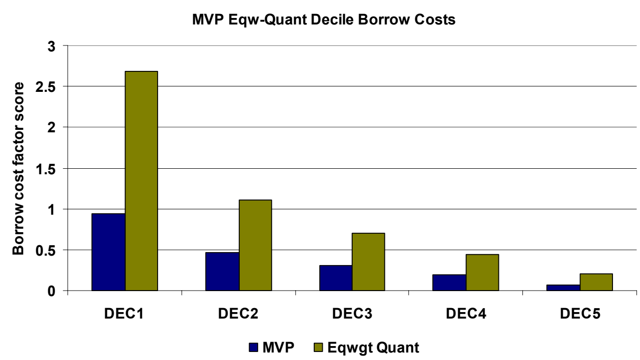

Source: Axioma, Comustat, IBES, Russell, S&P, Deutsche Bank, Data Explorers

Figure 47: Borrow cost factor scores: MVP vs. equal weight quant factor portfolio running at same risk levels (higher score implies higher borrow cost)

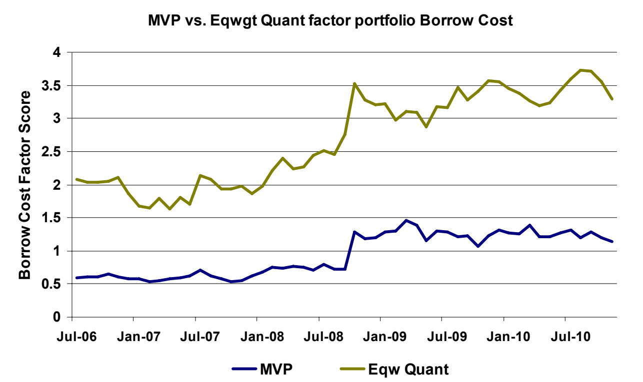

Source: Axioma, Comustat, IBES, Russell, S&P, Deutsche Bank, Data Explorers

MVP under the microscope

In this part of the report, we make use of various analytical results that will open up a spigot of insights and yield some interesting and fascinating properties related to minimum variance portfolios.

We start with a very intuitive characterization of the MVP strategy derived from a simple one-factor risk model (e.g. CAPM) to reveal the strategy's purpose and analyze its composition. We will generalize this characterization to the multi-factor case to find that the MVP strategy is very similar to a strategy that predicts that stocks with high common factor risk will underperform.

The second result is a simple one that provides some surprising properties about the comovement and correlation of the MVP with other portfolios and factors. These results will allow us to understand the MVP more clearly; outlining its strengths and potential pitfalls and more importantly, permitting us to enhance the strategy with a fundamental understanding of what to expect.

Is MVP the anti-Beta?

The first result we investigate was derived by Clarke, Silva and Thorley (2010), and also by Scherer (2010). The derivation provides a simple and intuitive characterization of the minimum variance portfolio under a one factor risk model (e.g. CAPM). Both papers show that under a one-factor risk model, the weight of a stock in the MVP depends solely on the specific risk and relative Beta of the stock. Mathematically, they show that the weight of stock *i*, in the MVP is given by:

$$h_i = \frac{\sigma^2_{MVP}}{\sigma^2_{\varepsilon,i}}\left(1 - \frac{\beta_i}{\beta_{LS}}\right) \tag{3}$$

where h_i is the holding in asset *i*, σ^2_{MVP} is the variance of the MVP portfolio, β_i is the beta of stock *i* to the single risk factor, $\sigma^2_{\varepsilon,i}$ is the specific risk of stock *i*, and β_{LS} is a threshold beta, which, under certain assumptions, can be approximated by the average beta of the universe (see Clark *et al.* 2010 for exact formula and details). We can simplify equation (3) even further by approximating β_{LS} to equal 1 (since the average beta should be close to 1). This approximation leads to an even more intuitive expression for the MVP weights:

$$h_i \approx \frac{\sigma^2_{MVP}}{\sigma^2_{\varepsilon,i}}\left(1 - \beta_i\right). \tag{4}$$

The approximation for the weights in equation (4) lends itself to easy interpretation. The first thing that jumps out at us is that the strategy looks very much like a low beta minus high beta play. That is it will take long positions on stocks whose beta are less than 1, while shorting stocks whose betas are greater than 1. In addition, on the long-side, stocks with smaller beta will garner higher weight, while the short-side will allocate stronger negative weight to stocks with the higher beta.

We can take the approximation one step further to make the latter formula even more intuitive to quant practitioners. If we make the further assumption that the average beta across the universe is equal to 1, then we have that:

$$h_i \approx \sigma^2_{MVP} \frac{-\tilde{\beta}_i}{\sigma^2_{\varepsilon,i}} \tag{5}$$

where $\tilde{\beta}_i$ is the mean-adjusted beta for stock i. Equation (5) makes it even clearer: all else equal, the MVP aligns the portfolio in reverse to stock beta. More importantly, the MVP which is supposed to be agnostic to expected stock returns has revealed itself as a strategy that, all else equal, makes the implicit forecast that stocks with higher beta will underperform stocks with lower beta.

This, of course, flies completely in the face of traditional asset pricing theory. If we believe CAPM and think of beta as a proxy for a stock's expected return then any of the characterizations in (3)-(5) imply that the MVP strategy is forming a portfolio in reverse direction of the expected return forecasts!

Perhaps a better way to think about equations (3)-(5) is to think of beta, not as a proxy for expected return, but instead as a proxy for risk[21]. In this case, the strategy can be justified by aligning it with exploiting the low-volatility anomaly. However, under this interpretation, the strategy is no longer agnostic with regards to expected returns; rather it is making the forecast that lower beta stocks will outperform higher beta stocks. This is far from taking a passive or neutral stance on predicted stock returns.

What about the idiosyncratic volatility?

The other parameter in equations (3)-(5) that plays a role in discriminating the asset weights of the portfolio is the idiosyncratic volatility of each asset, $\sigma^2_{\varepsilon,i}$. However, as we will see in the subsequent section, this parameter is what causes the MVP strategy to deviate from a conventional low-volatility strategy. One obvious fact is that the idiosyncratic risk will play no role in determining the *sign* of the weight of an asset in the MVP. It only plays a role in determining the *magnitude* of the weight.

A subtle, but significant inconsistency with low-volatility strategies

The role that stock level idiosyncratic volatility plays in the MVP strategy reveals something very important about the strategy. In fact, a close look at equation (5) reveals an inconsistency between the MVP strategy and that of other conventional low-volatility strategies, which aim to profit from betting against risky stocks. To see this we have to look at the role of beta and idiosyncratic volatility on the long and short side of the portfolio separately.

The long side of the portfolio is consistent with a low-volatility strategy. This means that if low-risk stocks outperform high-risk stocks, we should expect the long side of the MVP to outperform. This is because the long side puts more weight on stocks with lower beta *and* overweighs stocks with lower idiosyncratic risk.

The inconsistency arises on the short side. To see this note that if all stocks on the short-side had the same beta then the short-side of the portfolio would be overweight (less short) stocks with higher idiosyncratic volatility and underweight (more short) stocks with lower idiosyncratic volatility – the anti-thesis of a low-risk strategy. As we will see below, the correlation between beta and idiosyncratic volatility will complicate the matter even further, but it should be evident that during times when idiosyncratic risk is high relative to beta risk, the short side of the MVP could turn into a risk-seeking strategy.

Is this an anomaly in the strategy? The answer is no because the MVP strategy, in contrast to conventional low-risk strategies, has the objective to minimize portfolio level risk. Unlike, low-risk strategies its objective is not to position itself against stand-alone asset risk. We can actually boil down the difference between the MVP and a low-risk strategy to portfolio mechanics issue: idiosyncratic risk is always additive in a portfolio and whether long or short higher magnitudes of specific risk will induce higher portfolio risk.

[21] In reality, this is two sides of the same coin in that the expected return should be aligned with expected risk, which is of course what CAPM and other asset-pricing theory is all about. The reason we suggest thinking about beta as common factor risk is to treat it in the context of the low-volatility anomaly as we discussed in the first section.

Finally, we can use two real-life examples to see how the MVP weights are related to beta and idiosyncratic risk[22]. In addition, we will see how this relationship will be heavily influenced by the dominant source of risk at any particular point in time.

First, we observe the relationship between the MVP weights with beta and idiosyncratic risk as of Dec 31, 2010 in Figure 48 and Figure 49, respectively. We note this is a period in which asset specific risk is at historical lows. Note that beta has a negative linear-like behavior with the MVP weights (Figure 48). In contrast, idiosyncratic volatility does not share the same linear behavior, especially on the short-side of the portfolio. Note that the short-side of the portfolio allocates less short exposure (higher weight) to many stocks with very high specific risk. This short-side specific risk exposure is inconsistent with a low-volatility strategy.

Now we look at a period in which specific risk was a very strong driver of asset total risk. Figure 50 and Figure 51 show the MVP weights relative to beta and specific risk, respectively as of Jan 31, 2001. This was a period when specific risk levels were at historic high levels.

Figure 48: MVP weights vs. Predicted Beta (Dec 2010)

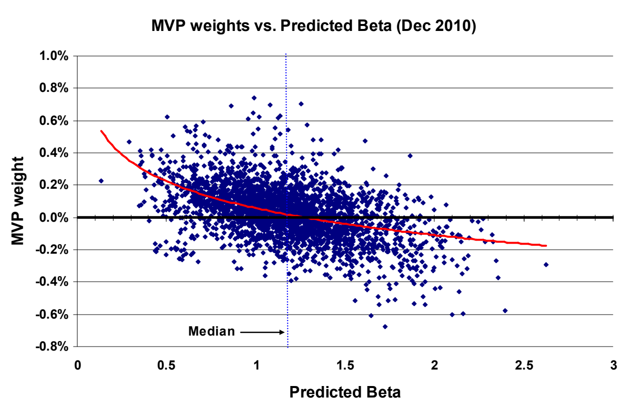

Source: Axioma, Comustat, IBES, Russell, S&P, Deutsche Bank, Data Explorers

Figure 49: MVP weights vs. Specific Risk (Dec 2010)

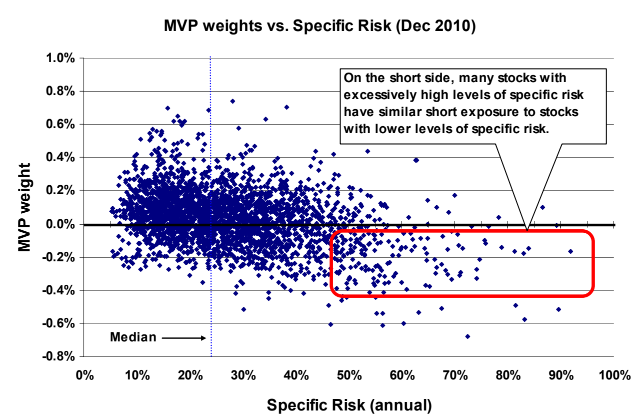

Source: Axioma, Comustat, IBES, Russell, S&P, Deutsche Bank, Data Explorers

Figure 50: MVP weights vs. Predicted Beta (Jan 2001)

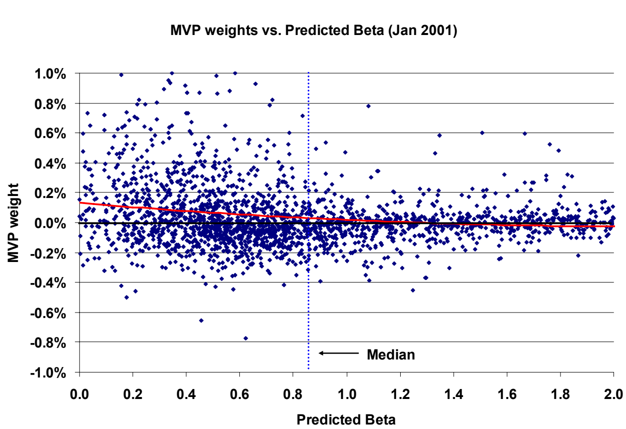

Source: Axioma, Comustat, IBES, Russell, S&P, Deutsche Bank, Data Explorers

Figure 51: MVP weights vs. Specific Risk (Jan 2001)

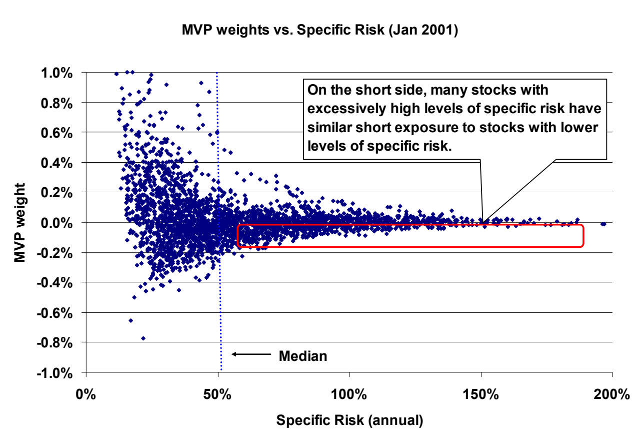

Source: Axioma, Comustat, IBES, Russell, S&P, Deutsche Bank, Data Explorers

---

[22] We proxy the idiosyncratic risk with specific risk because this MVP is actually constructed using the risk model and not the simple one-factor beta model from which we obtained the characterization.

First, note how the MVP weights are much less linear with respect to Predicted Beta in Figure 50 relative to Figure 48. This is because specific risk dominates the strategy and is throwing off the natural linear relationship between beta and the MVP weights. Second, note how the short-side of the portfolio tends too allocate significantly more short exposure (lower weight) to assets with lower specific risk; and less short exposure (higher weight) to assets with higher specific risk. It should be obvious that during this period the short-side of the MVP strategy is not a low-risk strategy; rather it will outperform when higher risk stocks outperform lower risk stocks.

Can this short-side idiosyncratic volatility preference hurt the portfolio? The answer seems to be yes, but it depends on whether asset specific risk is dominant relative to market or common factor risk. Perhaps the apex of this analysis can be boiled down to Figure 52. It shows the 12-month average return of the MVP strategy versus the levels of specific risk. The figure clearly shows that high levels of specific risk coincide with underperformance of the MVP strategy. This suggests that this higher specific risk preference on the short-side of the portfolio may hurt performance, especially when specific risk is the dominant source or asset risk.

Figure 52: MVP 12-month average return vs. Specific Risk levels

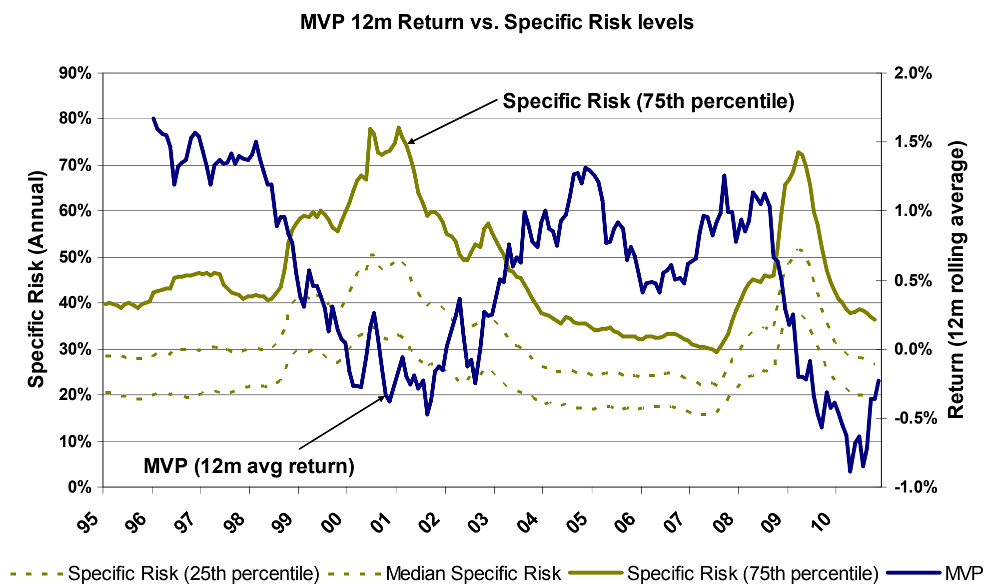

Source: Axioma, Comustat, IBES, Russell, S&P, Deutsche Bank, Data Explorers

But Beta and idiosyncratic volatility are highly correlated

A hidden element in the characterization above is the correlation between beta and idiosyncratic volatility. This is a well known relationship, and it is very intuitive: a risky stock is expected to have both higher beta and higher idiosyncratic risk; while a less risky stock is expected to have lower beta and lower idiosyncratic risk. To see this, Figure 53 shows the relationship between Predicted Beta and Idiosyncratic Volatility values for stocks in our universe on Dec 31, 2010. To see that this is not a one-time relationship, Figure 53 shows the rank correlation between the two factor scores over our sample. Note that the average rank correlation hovers between 40% and 50% over the sample period.

If we look back to equation (5), we see that the relationship between these two attributes affect the weights in the MVP portfolio in a significant way. We saw a preview to this in Figure 48. There we saw how the relationship between the MVP weights and beta was nonlinear when comparing the long-side of the portfolio with the short-side. If we look back at the chart, we see that indeed, as beta gets large the short-side of the portfolio allocates less short exposure per unit beta. This is another cause for concern when putting the strategy in the context of a low-risk strategy. It says that two betas that are equidistant but on opposite sides of the average beta (or the long/short threshold in equation 3) will be treated differently on the long-side versus short-side of the portfolio. The long-side will tend to allocate a higher magnitude long-exposure relative to the short exposure allocated by the short-side. This not only creates an imbalance in the holdings, which explains some of the results in the first section, but it will cause the MVP to profit less robustly when low-risk stocks outperform high-risk stocks.

Figure 53: Predicted Beta vs. Idiosyncratic Volatility factor scores as of Dec 31, 2010

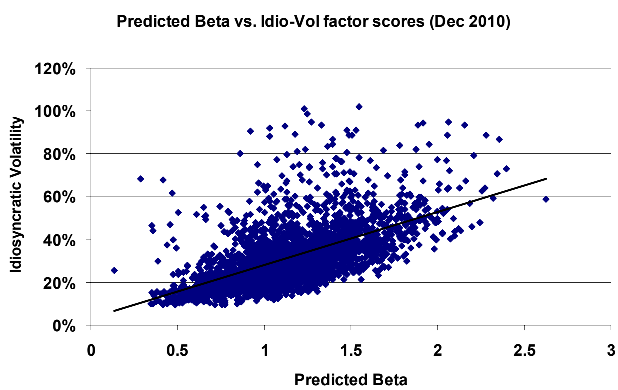

Source: Axioma, Comustat, IBES, Russell, S&P, Deutsche Bank

Figure 54: Rank correlation between cross-sectional factor scores of predicted beta and idiosyncratic vol

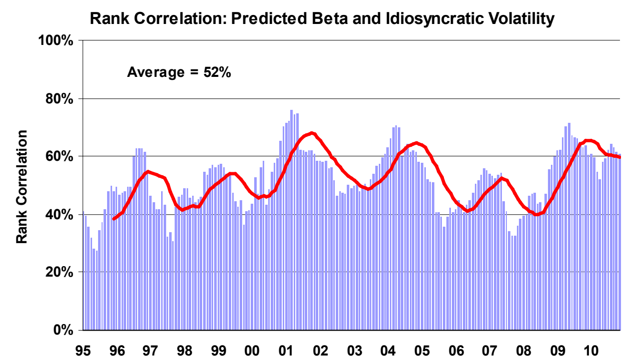

Source: Axioma, Comustat, IBES, Russell, S&P, Deutsche Bank

Risk is a multi-factor affair

It is well accepted that systematic risk is best captured using multiple factors. Most quantitative practitioners use multi-factor risk models either from well-known vendors such as Axioma, Barra or Northfield, or have custom versions of their own. Therefore, we must ask ourselves whether everything we did in the past section is valid or extrapolates to the multi-factor space or is it only particular to the simple one-factor CAPM model.

The characterization in equation (3) used a clever shortcut to express the inverse of a covariance matrix when the covariance structure is driven by one factor. However, extrapolating this to the multi-factor case is intractable. Luckily, there is a paper – often overlooked – that has an intuitive expression for the product of the inverse of the covariance matrix with another vector[23]. Using the result from Stevens [1998] (see equation 9), we can write the holdings of the MVP as:

$$h_i = \frac{\sigma^2_{MVP}}{\sigma^2_{\varepsilon,i}}\left(1 - \sum_{j \neq i} \beta_j\right) \tag{6}$$

where the β_j ;'s are the implied multi-variable regression betas we would obtain from regressing stock *i* on all the other stocks in the universe. These multivariate betas are implied from the risk model (i.e. there is not an actual regression taking place).

Equation (6) tells us that stocks with higher exposure (or covariance) to other stocks will be underweight in the MVP, while stocks with lower exposure to other stocks will be overweight. In a factor risk model the exposure or covariance between two assets is driven only by the common factors in the model. Therefore, we can argue that equation (6) is a

[23] In their case they did it to characterize the mean-variance optimal portfolio, but it can be generalized as we noted above. In any case, the result can be directly applied if we use the mean-variance formulation of the MVP in the last part of the report.

similar characterization of equation (3) except for the multi-factor case. It is not hard to show that the sum of the betas in equation (6) is driven by the actual exposures of the stocks to the common factors in the model. Therefore, instead of beta to the market factor as shown in equation (3), the formula would involve exposures to the common factors.

Comovement with other portfolios, factors and attributes

As we saw in a previous section, the MVP generates returns that have negligible correlations with MVO and z-score alpha-factor portfolios (see Figure 34 and Figure 35). Was this empirical result a fluke or is there something special about the MVP strategy that makes it inherently uncorrelated with quant strategies? The answer can be found analytically using two lines of matrix algebra. It can be shown (see appendix) that the expected (ex-ante) correlation of the MVP portfolio with any other portfolio, P is given by:

$$corr(r_{MVP}, r_P) = \frac{\sigma_{MVP}}{\sigma_P} \sum_i h_{P,i} \tag{7}$$

This simple formula provides many valuable properties and insights. The first thing to note about the formula is that there is no longer a covariance matrix! Correlations no longer depend on the covariance structure of the assets[24].

The second important relationship we can appreciate from the formula is that the ex-ante correlation between the MVP and any other portfolio *P*, is inversely proportional to the volatility of portfolio *P*. This explains why the correlation with the market (see Figure 34) rises in times when market volatility is low and drops when market volatility is high.

Third, we can see that for any dollar-neutral portfolio, or more generally any portfolio with a net zero weight will have zero ex-ante correlation with the MVP.

Fourth, we see that the correlation will be directly proportional to the sum of the weights in the portfolio. An immediate corollary to this is that the correlation of the MVP with a portfolio that is net positive will be positive and vice-versa.

We also note that we could easily replace the portfolio P in equation (7) with factor z-scores. The equation then states that the correlation of the MVP with a factor should be 0 or close to 0 depending on the algorithm used to z-score the factor. This points to another powerful property: mainly that the MVP will be uncorrelated (ex-ante) to most classes of high-minus-low factor portfolios.

In addition, equation (7) will help us decipher other questions that arise from certain empirical data and it will also lend itself useful when mixing the strategy with other alpha strategies.

Last, it is pertinent to say that these results are all in the context of a forecast. This is because they are all ex-ante and they will only be realized if the variance covariance matrix is exactly right. We know that this is never the case and the recent volatility in risk estimates has made forecasting even more difficult. However, we have to work with the "best" estimate we have and achieving a better risk forecast is something we will leave for another day.

Below is a short list of important results from this section.

1) The correlation and exposure between the MVP and any dollar neutral portfolio is 0.

---

[24] A nice way to explain this come from Lee (2010), who noted that the marginal contribution to risk (MCTR) of every asset in the portfolio was the same

2) The correlation between the MVP and a portfolio with negative weights is negative.

3) The correlation between the MVP and a portfolio with positive weights is positive.

4) The magnitude of the correlation between the MVP and any portfolio P, will be directly proportional to the magnitude of the net weight of the portfolio and inversely proportional to the volatility of the portfolio.

A quick recap of what we know

Up to this point, we have seen a mountain of statistics, results, characterizations and analytics about the MVP to drive anyone mad. We now summarize many of our more important findings and their links to the empirical results found in the first part of the report.

1) The strategy is not passive or a conventional low-volatility strategy.

2) The MVP is a long/short portfolio.

3) The MVP is aligned with a low-beta minus high-beta strategy.

4) The MVP as a whole is <u>not</u> aligned with a low idiosyncratic volatility strategy.

   a. The long-side of the MVP is aligned with lower idiosyncratic risk; i.e. all else equal, it will outperform when low idiosyncratic risk stocks outperform stocks with higher idiosyncratic risk.

   b. In contrast, the short-side of the MVP is aligned with higher idiosyncratic risk; i.e. all else equal, it will outperform when higher idiosyncratic risk stocks outperform stock with lower idiosyncratic risk.

5) All else equal, the strategy will overweight stocks that are less correlated across the universe and will underweight stocks that have higher correlation across other stocks in the universe.

6) The fact that most factor risk models are calibrated towards higher market capitalization stocks will cause a bias in the portfolio towards smaller and less liquid stocks. In past data, we also find that small and highly illiquid stocks tended to have lower beta due to stronger market dislocations, stale pricing and other data issues.

7) Analytically, the ex-ante (expected) correlation between the MVP and any dollar-neutral or holdings-neutral portfolio is <u>zero</u>. We find this to be consistent when looking out-of-sample (any error is the covariance matrix's fault).

8) Analytically, the ex-ante correlation between the MVP and a net-short portfolio is negative. Similarly, the correlation with a net-long portfolio is positive. The larger the net long (short) position the more positive (negative) the correlation.

9) All else equal, the *magnitude* of the ex-ante correlation between the MVP and any portfolio *P* will be inversely proportional to the risk of the portfolio *P*.

(EMVP) Enhanced minimum variance portfolio strategy

Improving on traditional minimum variance

How can we improve upon MVP performance without veering too far from the intended purpose of the strategy? In this section we look at one way to enhance the minimum variance portfolio construction in order to improve both performance and risk control. In addition, we will look at two methods that can be used to combine minimum variance portfolios with alpha models to generate portfolios with improved risk-adjusted return.

In a break with most practioner minimum-variance portfolio research, we do not data-mine the optimization to find the right settings for industry, style or beta exposures. Neither, do we overlay factor upon factor on the strategy and choose the one that best works over our sample; then justifying after the fact that this or that style tilt works because it "pairs" well with a minimum variance portfolio.

In the following sections we will show why most practioner publications show worse risk-adjusted performance after neutralizing the MVP for industry, style or beta tilts. In addition, we show a very interesting property – already known to many – that shows that the minimum variance portfolio will arise naturally in an optimal portfolio that has a total holdings constraint (e.g. the dollar neutral constraint).

Letting the analytics guide us…

It is always necessary to verify theory with empirical testing and many times we allow the data to be our guide and judge the efficacy, properties and other attributes of a strategy. However, relying solely on the data and simulations can lead to data mining, biases, and may obfuscate important properties or attributes of the strategy. In the following, we use some of the analytical results obtained in previous sections to shed light on possible extensions of the minimum variance portfolio strategy. More importantly, we will use these results to guide us towards enhancing the strategy without data mining or torturing the optimizer. In our quest, we will see why many portfolio construction techniques that work for traditional quant factors are useless for minimum variance strategies. And we will see that and we will know what to avoid so that we do not waste time, data mine for performance or deviate too far from the intended purpose of the strategy.

By now it should be clear that one side of the MVP relies on the mispricing of risk so it will tend to outperform when investors shift their risk appetite downward (become more risk averse) and underperform when investors increase their risk appetite (become less risk-averse). The other side of the MVP is diversification or portfolio risk minimization. As we saw in the previous section this objective is not entirely consistent with a low-volatility strategy in which outperformance depends on lower risk assets outperforming those with higher risk.

We will look to see if we can improve the MVP by shifting the strategy slightly so that it still attempts to minimize risk, but at the same time stays consistent with a low-volatility strategy.

Neutralizing minimum variance portfolios

Can we improve minimum variance performance by neutralizing it for certain factors? The first exposure we would consider neutralizing in our MVP is the nagging beta exposure we saw back in Figure 22. In our analytical section, we saw how the nature of the MVP strategy could be biased towards beta. The bias was mainly due to the effect of the specific (idiosyncratic) risk on the weights and the high correlation between specific risk and beta. So would neutralization help the portfolio? To investigate we apply the conventional quant constraint of neutralizing the beta exposure; the sector exposures; and then proceeding to the style neutral MVP. The results in Figure 55, Figure 56, Figure 57 and Figure 58 show these results. Neutralization seems to deteriorate risk-adjusted performance, lowering return while in many cases increasing risk and the correlation to the market.

Figure 55: IR: MVP and neutralized-MVP (Jan 95 – Dec 10)

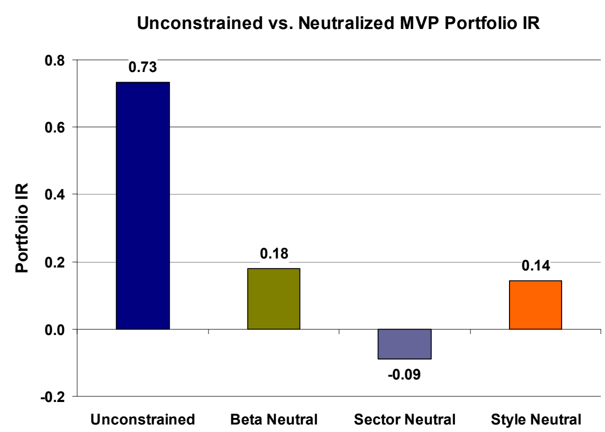

Source: Axioma, Comustat, IBES, Russell, S&P, Deutsche Bank

Figure 56: Average Return: MVP and neutralized-MVP (Jan 95 – Dec 10)

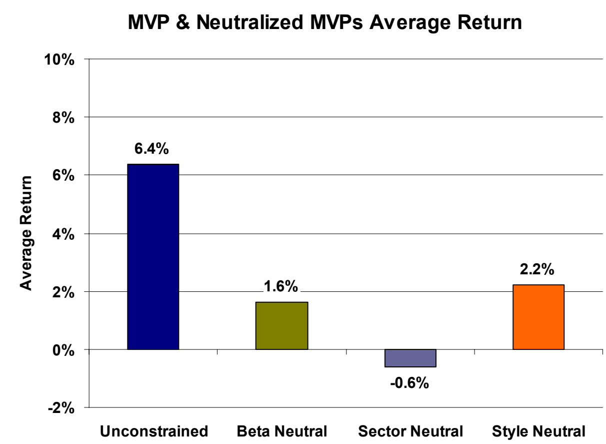

Source: Axioma, Comustat, IBES, Russell, S&P, Deutsche Bank

Figure 57: Realized Volatility: MVP and neutralized-MVP (Jan 95 – Dec 10)

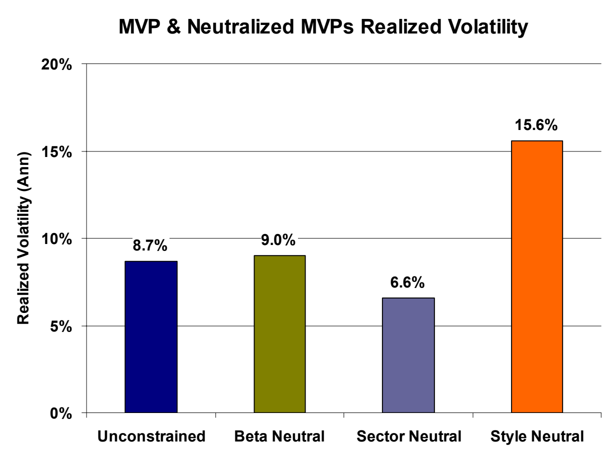

Source: Axioma, Comustat, IBES, Russell, S&P, Deutsche Bank

Figure 58: Correlation to the Market: MVP and neutralized-MVP (Jan 95 – Dec 10)

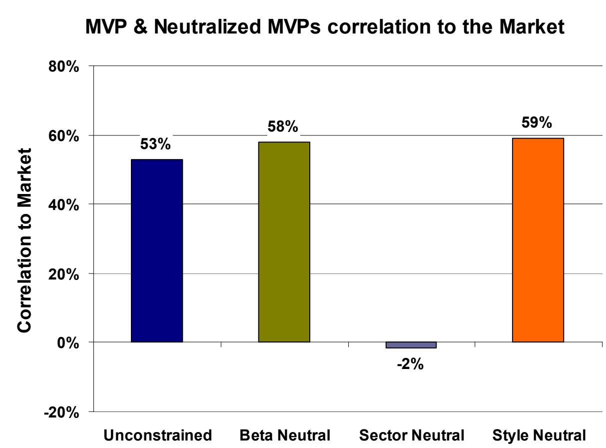

Source: Axioma, Comustat, IBES, Russell, S&P, Deutsche Bank

What's going on? Isn't neutralization supposed to correct for correlation to a factor? The answer is: it depends what type of neutralization technique is being used. For example, optimizers use what we call exposure-neutralization. This type of neutralization is achieved through a constraint and guarantees that the final portfolio will have an exposure of 0 (or close to 0) to the attribute being neutralized. However, it does not guarantee that the final portfolio will be uncorrelated to the underlying neutralizing attribute. While this may seem unintuitive, take the example of generating a portfolio with an exposure constraint of zero to the Software industry. While the resulting portfolio may be guaranteed to have a zero exposure to the Software industry, it is not guaranteed to have an exposure of zero to the Telecom industry. However, the Telecom industry in turn is highly correlated to the Software industry so the portfolio exposure to the Telecom industry – which is not 0 – will induce a correlation of the constrained portfolio with the Software industry. For those that are familiar with risk models it is easier to see. A risk model factor can be thought to be a portfolio with an exposure of 1 to the factor and an exposure of 0 to every other factor in the risk model. Therefore, the Software industry factor can be thought of a factor with an exposure of 1 to the software industry and an exposure of 0 to every other factor. However, the Software industry factor returns are highly correlated to every other industry factor in the risk model because all industry factors are correlated through the market. So having an exposure of zero to these other factors doesn't exactly make the portfolio representing the Software industry factor uncorrelated to the other industry factors.

In fact, it is not hard to show that neutralizing for beta (z-scores) the final covariance of the neutralized MVP with a high minus low beta portfolio (z-score portfolio) will result to be a function of the original beta exposure.

Another type of neutralization is correlation neutralization, which we introduced in Luo *et al* [2010c] "Volatility=1/N" and discussed the differences in another report Alvarez *et al*, 2010, "Factor Neutralization and Beyond". This does not work for the case of the MVP strategy because we found that the correlation with z-score attributes will be zero or close to zero.

Enhancing MVP to be consistent with a low-volatility strategy

Given that neutralization does not seem to work is there a way to enhance the strategy without deviating too far from its intended purpose and stripping it of its most useful properties? The answer is yes, but care must be taken not to change the nature of the strategy. The enhancement we propose is a slight one that aims to make the strategy more consistent with a conventional low-risk strategy. However, we find that it improves risk-adjusted performance, lowers the correlation with the market and conserves the most desirable properties such as its low correlation with other quantitative strategies.

The enhancement is simple and straightforward. We assume that the specific risk of each asset is the same and re-compute the asset-by-asset covariance matrix. We then use this covariance matrix to construct our enhanced minimum variance portfolio (EMPV). The aim of the enhancement is to the remove the role that specific risk plays on the short-side of the strategy, which tends to overweight (apply less short exposure) to stocks with higher specific risk. Looking back to equation (5), this enhancement has the effect of converting the strategy into a pure low-beta strategy. Similarly, in the case of multiple risk factors, we can use equation (6) to argue that the strategy will become a low common-factor risk strategy.

Figure 59 shows the effect of the enhancement on the exposure to specific risk. Note the large difference in the specific risk exposure between the first deciles of the MVP and EMVP. The difference indicates that the EMVP has much more short exposure to assets with higher specific risk, which will make it more aligned with a low-volatility strategy. In other words, we would expect the EMVP to outperform the MVP when stocks with high levels of specific risk underperform.

Figure 60 shows us the difference in beta exposure. Again we see that the EMVP strategy has more short exposure to stocks with higher beta. This is actually because the enhancement has mitigated the nonlinear effect on beta exposure resulting from the correlation between beta and specific risk (see Figure 48 and Figure 53).

Figure 59: MVP and EMVP specific risk exposure by decile (Jan 95 – Dec 10)

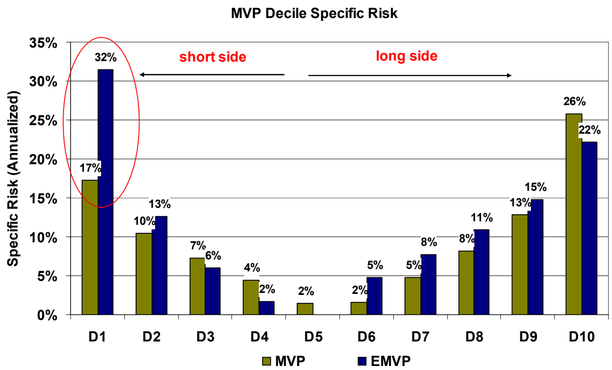

Source: Axioma, Comustat, IBES, Russell, S&P, Deutsche Bank

Figure 60: MVP and EMVP predicted beta exposure by decile (Jan 95 – Dec 10)

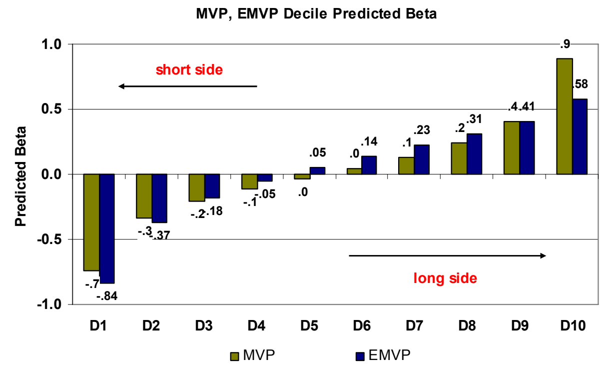

Source: Axioma, Comustat, IBES, Russell, S&P, Deutsche Bank

The EMVP strategy is now more consistent with a low-volatility strategy. However, the important question is whether it actually increases the performance of the strategy and at what costs?

In Figure 61 we show that the EMVP has a higher average return than the MVP over both the full sample and more recent period. The next question is whether we pay for the extra performance with added risk. Figure 62 shows that the increase in risk of the EMVP is quite marginal for both the entire and more recent periods. This result is quite significant in that it suggests that the levels of asset specific risk are not very relevant in the portfolio minimization problem. So we have better return, for negligible increase in risk which leads us to Figure 63 which shows that the EMVP outperforms the MVP on a risk-adjusted basis. As a final check, we look to see if the results are consistent when using the look-ahead risk model. The reason for this test is to see if having a better forecast of specific risk would cancel out the outperformance of the EMVP. Figure 64 indicates that even with look-ahead bias, the EMVP still outperforms the MVP. The results are surprising in that having a better forecast of asset level specific risk does not improve performance.

Last we want to point out that we can think of this enhancement as shrinking the correlations of the covariance matrix. To see this, note that a stock with higher specific risk will less correlated across other stocks in the universe, while stocks with lower specific risk will tend to be more correlated across the other stocks in the universe. Setting all the stock specific risk levels to the universe average will have the effect of increasing the overall correlation of high specific risk stocks and decreasing the correlation of stocks with originally low levels of specific risk.

Figure 61: EMVP vs. MVP mean returns

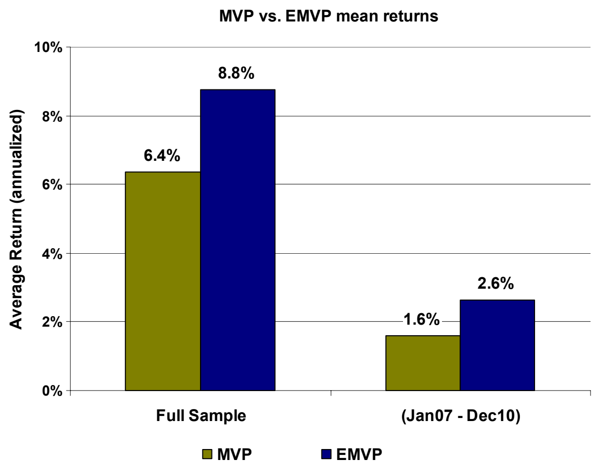

Source: Axioma, Comustat, IBES, Russell, S&P, Deutsche Bank

Figure 62: EMVP vs. MVP volatility (realized)

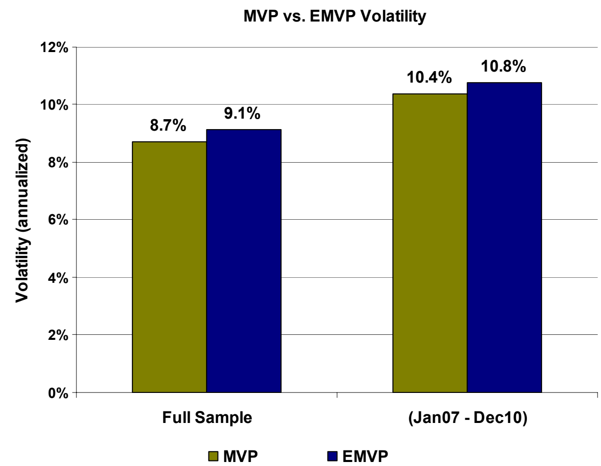

Source: Axioma, Comustat, IBES, Russell, S&P, Deutsche Bank

Figure 63: EMVP vs. MVP portfolio IR

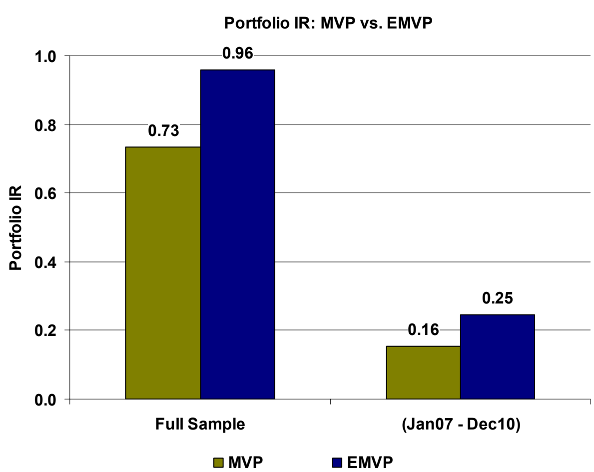

Source: Axioma, Comustat, IBES, Russell, S&P, Deutsche Bank

Figure 64: Look-ahead EMVP vs. MVP portfolio IR

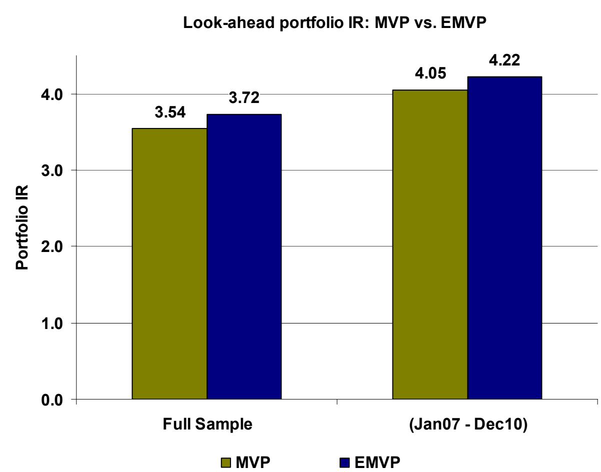

Source: Axioma, Comustat, IBES, Russell, S&P, Deutsche Bank

Last, we want to make sure that the EMVP has some of the same desirable properties we found for the MVP. Figure 65 shows that the low correlations with traditional quant strategy portfolios remain negligible. In addition and probably of more value is that the correlation with the market has dropped off. Last we show that the EMVP is not much different that the MVP when looking at their correlation across time (Figure 66). We also note looking back at Figure 52 that the correlation is lower during periods when idiosyncratic risk levels are higher. This is consistent with the methodology used to construct the EMVP.

Figure 65: EMVP, MVP Correlation with market and MVO factor portfolios

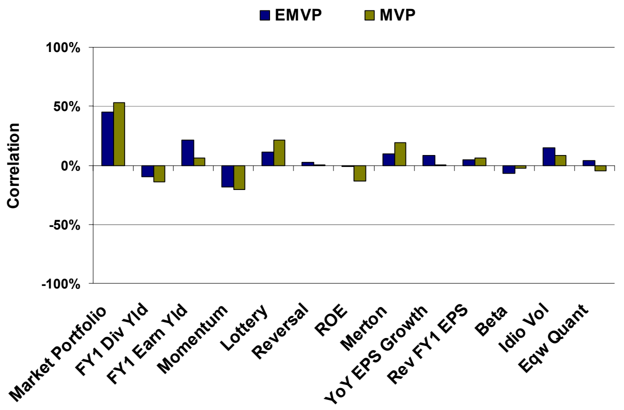

Source: Axioma, Comustat, IBES, Russell, S&P, Deutsche Bank

Figure 66: EMVP, MVP realized correlation (24-month rolling window)

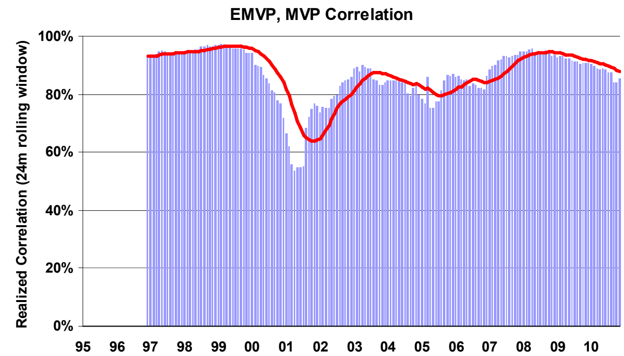

Source: Axioma, Comustat, IBES, Russell, S&P, Deutsche Bank

What happens to borrow costs?

Our prior research analyzing borrow costs (see Cahan *et al* 2011) showed that a strategy which was short stocks with high idiosyncratic volatility incurred relatively high amounts of borrow costs. This is not great news for our EMVP strategy, given that it will tend to have more short exposure to stocks with high specific risk. Therefore, we have to be careful that borrow costs do not increase substantially.

To see the effect of our enhancement on borrow cost exposure note that Figure 67 shows that borrow costs actually increased rather substantially for the bottom decile (most short exposure), but decreased for deciles 3,4 and 5. This is consistent with the change in the strategy produced by the enhancement. The enhancement (EMVP) has pushed stocks with higher specific risks towards the bottom decile and so the borrow costs have also shifted towards that decile. However, if we add up the total borrow cost on the short side of the EMVP we find that it is roughly the same – in fact lower – as that for the MVP.

Figure 67: MVP vs. EMVP Borrow Cost factor scores per decile Jul 06 – Dec 10 (higher score implies higher costs)

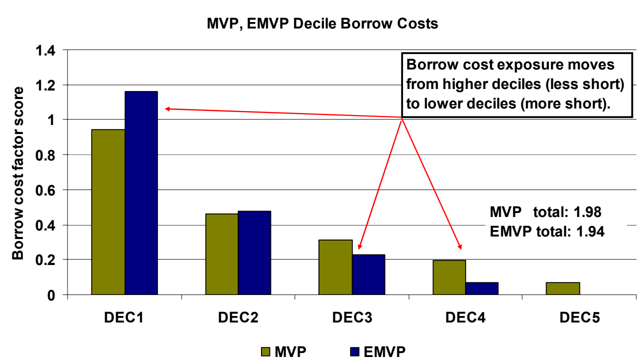

Source: Axioma, Comustat, IBES, Russell, S&P, Deutsche Bank, Data Explorers

Figure 68: MVP vs. EMVP Borrow Cost factor scores for bottom decile (higher score implies higher costs)

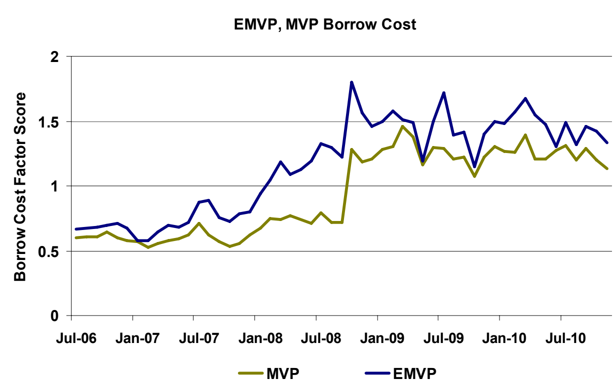

Source: Axioma, Comustat, IBES, Russell, S&P, Deutsche Bank, Data Explorers

Figure 68 compares the evolution of borrow costs for the bottom deciles (largest short exposures) both the MVP and EMVP. Note that EMVP has consistently higher borrow costs, but the difference over the more recent period has leveled off. In addition, if we look back to Figure 47, we find that the borrow costs are still quite low when compared to a traditional quant strategy running at a similar risk level.

How can we use the EMVP in practice?
To use the EMVP we can simply back out the implied alphas and use in the optimizer using a mean-variance framework. In fact, the next section shows how we can use the mean variance framework to construct the MVP as well. For completeness, the EMVP implied alphas have the following form:

$$\boldsymbol{\alpha}_{EMVP} = \boldsymbol{V}\boldsymbol{h}_{EMVP} = \boldsymbol{V}\left(\frac{\tilde{\boldsymbol{V}}^{-1}\boldsymbol{1}}{\boldsymbol{1}'\tilde{\boldsymbol{V}}^{-1}\boldsymbol{1}}\right), \tag{8}$$

where $\tilde{\boldsymbol{V}}$ is the variance covariance matrix constructed assuming equal asset specific risks.

A different formulation for minimum variance portfolios

The formulation in equation (1) involves minimizing a portfolio without maximizing any particular return forecast. Next to the covariance matrix, the lack of a return forecast is of fundamental importance to the minimum variance portfolio. Given that there is not a return forecast, is there a way to pose the problem in the usual mean-variance (i.e. Markowitz 1952 optimization) framework? The answer is yes. In this section we show an equivalent formulation to the minimum-variance problem framing it in a mean-variance optimization context. Why do this? There are several reasons, but in most part we do it to make the problem more accessible to techniques and interpretations already developed for mean-variance optimal portfolios. However, as we will show, framing the problem in a mean-variance context adds one analytical benefit: we do not need the constraint that the weights sum to 1 or any particular number. While this is a subtle difference we find that working with the true ideal portfolio actually improves performance, but we do not pursue this avenue here because this requires finding the right scaling procedure, which could be somewhat tricky and is something we will discuss and investigate in a future report.

A mean-variance formulation
Mathematically, the mean-variance formulation of the MVP is given by solving the traditional mean-variance optimization problem:

$$\max_{\boldsymbol{h}} \ \boldsymbol{h}'\boldsymbol{1} - 2\lambda \boldsymbol{h}'\boldsymbol{V}\boldsymbol{h} \tag{9}$$

where $\boldsymbol{h}$ is the vector of optimal holdings (weights), $\boldsymbol{V}$ is the variance-covariance matrix of asset returns, and $\boldsymbol{1}$ is a vector of ones that represent the expected returns. The solution has the same form as that obtained in the traditional mean-variance optimization problem, except that alpha is replaced with a vector of ones:

$$\boldsymbol{h}_{MVP} = \frac{1}{\lambda}\boldsymbol{V}^{-1}\boldsymbol{1} \tag{10}$$

Note that this version of the MVP is similar to that in equation (2) except that the term in the parenthesis is missing. To get the same result as that obtained in equation (2) we would just add the fully invested constraint (weights sum to 1).

The fact that the vector of alphas is replaced with a vector of ones implies that we are not making relative return forecasts or that any one stock will outperform another. The idea is the same as before: let the risk model drive the portfolio decision by choosing the set of weights that generate the portfolio with the least amount of variance without any relative expected return information to differentiate the stocks.

Adding it to our QCD model

The QCD is our flagship alpha model (see Luo *et al* 2010, "DB Quant Handbook"). In backtests and during its live performance the model generates good risk-adjusted performance. Running the model at the same ex-ante risk of the MVP strategy, the QCD optimized portfolio produces 34% average annual returns and a realized annualized risk of 13%. The risk-adjusted return (IR) turns out to be 2.7, which looking back at Figure 2 is an enormous feat. The question is now whether the MVP or its close cousin the EMVP can add value to this already effective model. Given the return performance of the MVP reported throughout this report, we cannot expect the MVP to add value by increasing performance over time. Instead, we have to rely on the MVP's two most distinguishable properties; it is a portfolio with minimum risk and its diversification power with most alpha-factor portfolios.

There are many ways to combine strategies, but many of the techniques we are familiar with combine strategies using expected returns. Therefore, we leverage the formulation developed in the last section, which allows us to work with the MVP in a mean-variance framework[25]. This will allow us maximum flexibility to choose any number of factor modeling techniques for combining alpha factors. In this report, we keep things simple and use the Grinold & Kahn factor optimization approach, which consists in finding the weights that maximize the expected return of the final alpha while minimizing its risk. In line with keeping things simple we use naïve estimates for both the factor covariance matrix and the forecast of returns. The idea is if naïve and simple works, then sophistication can only help improve performance for the better. For more specialized and sophisticated techniques see Luo, Y. *et al* [2011] "Robust Factor Modeling". Specifically, we use the past 24-month returns of each strategy as a forecast for the mean and the sample covariance matrix[26]. Because, we have seen that the MVP is uncorrelated with most alpha strategies, we fix the correlation in the risk model to be 0. Last, we avoid taking negative weights by setting a weight to zero if the rolling 24-month average is less than zero[27].

Figure 69: QCD vs. QCD+MVP portfolio IR over (Nov 1999 – Dec 2010)

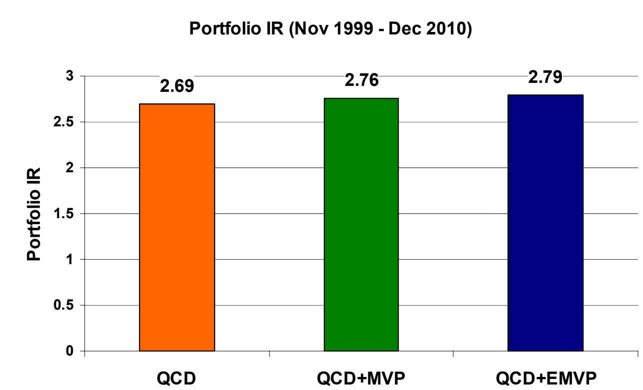

ource: Axioma, Comustat, IBES, Russell, S&P, Deutsche Bank

Figure 70: QCD vs. QCD+MVP portfolio IR over (Jan 2007 – Dec 2010)

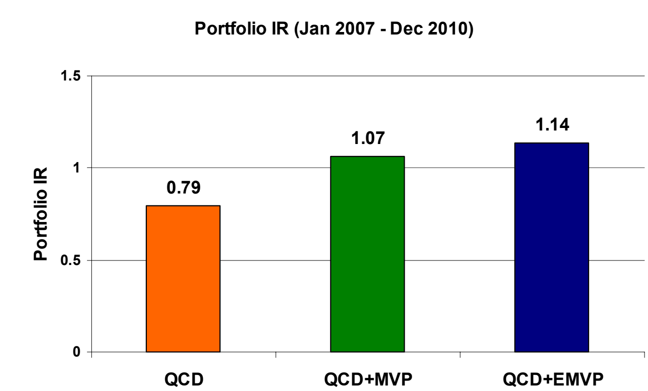

Source: Axioma, Comustat, IBES, Russell, S&P, Deutsche Bank

[25] This is not actually necessary. In theory, we can combine stand-alone portfolios in the same way we combine alpha factors. However, it is much simpler and more intuitive to work in the conventional framework.

[26] An expanding window mean and covariance matrix provides very similar results.

[27] The QCD model 24-month moving average is always positive so we are guaranteed to have exposure to the model. The MVP and EMVP strategies had 35 out of 135 periods with a negative 24-month moving average. In these cases the weight to either strategy was set to 0.

How about factor tilts?

We have seen various ad-hoc attempts to make the MVP strategy look better in backtests. One method is to apply a style-tilt towards generic alpha factors such as value, quality, momentum or growth. When a factor tilt works well in the backtest, the resulting portfolio is then touted to have better performance, less risk and less turnover than the original MVP. There is then a fundamental justification given for the better performance…

However, these results are a sell-fulfilling prophecy. The lower risk comes from the fact that a “hard” factor tilt results in a portfolio that is very similar to linear combination of the MVP and the corresponding factor portfolio[28]. For example, a “hard” factor tilt on value will result in a portfolio that is a linear combination of the MVP and the MVO Value factor portfolio we have looked at above. The harder the tilt the more weight is given to the factor portfolio. The resulting portfolio will take on the properties of both portfolios and a “hard” tilt will give more weight to the factor portfolio. Ultimately the resulting portfolio properties will be a mix of both portfolios.

A better way to tilt towards a factor is simply to combine the MVP strategy to an alpha strategy as we did in the last section with our QCD model. In the manner, we can use a bona fide alpha weighting technique instead of setting an arbitrary tilt.

Implementing the EMVP

In this last section of the report, we focus on the actual implementation of the EMVP. This will indicate whether the strategy can produce similar performance under risk, turnover constrains and after accounting for transaction costs.

We implement two strategies. The only difference between the two will be the target risk and the leverage at which they will run. Below, are the following steps and optimization settings for implementing the strategy:

1) We back out the implied alphas as suggested in our prior section.

2) Target risk at 6% for the high-risk strategy and 3% for the low-risk strategy.

3) Leverage constraint of 3x for the high-risk strategy and 2x for the low-risk strategy.

4) Turnover constraint of 40% two-way for each of the strategies.

5) Use 20bps linear t-cost per buy/sell in the objective function (this will slow down turnover as well).

6) No constraints or neutralization of betas, industry or styles.

Figure 71 and Figure 72 show the after-cost results for the high-risk and low-risk strategies, respectively. The high-risk strategy produced an after-cost Sharpe ratio of 1.13, while the low-risk strategy produced an after-cost Sharpe ratio of 0.99. These results show that indeed this strategy is a promising one on its own.

[28] It is true that risk model factors can be very different than the alpha-factors we have studied above with the same name. However, given the variance covariance matrix we could easily construct a factor-mimicking portfolio with very similar properties to the risk model factor and then mix it with the MVP.

Figure 71: High risk EMVP portfolio

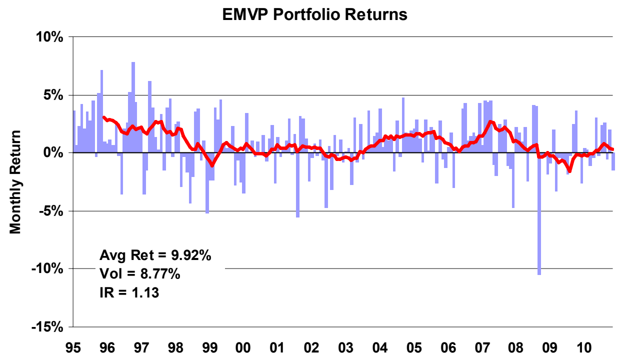

Source: Axioma, Comustat, IBES, Russell, S&P, Deutsche Bank

Figure 72: Low risk EMVP portfolio

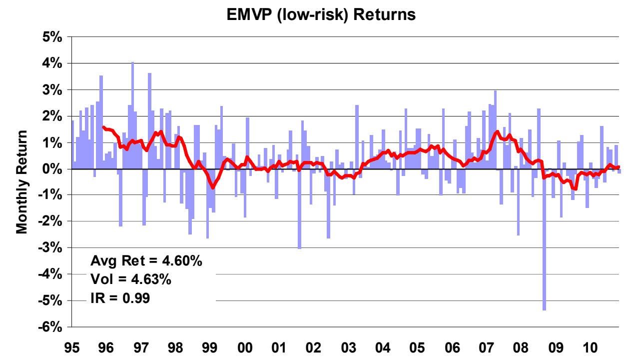

Source: Axioma, Comustat, IBES, Russell, S&P, Deutsche Bank

Bibliography

Alvarez, M., Kassam, A., Mesomeris, S. [2010]. "Factor Neutralization and Beyond", Deutsche Bank Quantitative Strategy, September 21, 2010.

Alvarez, M., Chen, Z., Luo, Y., Cahan, R., Jussa, J., "Portfolios Under Construction: Correlation and Opportunity", Deutsche Bank Quantitative Strategy, Dec 2, 2010.

Barberis, N., Huang, M., 2008, "Stocks as Lotteries: The Implications of Probability Weighting for Security Prices.", *American Economic Review*, vol. 98, no. 5

Cahan, R., Luo, Y., Alvarez, M., Jussa, J., Weiner, S., 2010 "Signal Processing: The Options Issue",Deutsche Bank Quantitative Strategy, May 12, 2010.

Cahan R., Luo, Y., Alvarez, M., Jussa, J., Chen, Z., 2011 "Signal Processing: The long and the short of it", Deutsche Bank Quantitative Strategy, Jan 18, 2011.

Clarke R., de Silva H., Thorley S. [2011]. "Minimum-Variance Portfolio Composition", *Journal of Portfolio Management*, Winter 2011

Baker, M., Brendan B., Wurgler, J., [2011]. "Benchmarks as limits to arbitrage: Understanding the low-volatility anomaly", Financial Analyst Journal, Volume 67, Number 1, 2011

Grinold, R., Kahn. R. [1999]. *Active Portfolio Management: A Quantitative Approach for Producing Superior Returns and Controlling Risk*, McGraw-Hill, 2nd Edition.

Haugen, R.A., and Baker, N.L. [1991], "The efficient market inefficiency of capitalization-weighted stock portfolios", *Journal of Portfolio Management*, Sprint 1991.

Jussa, J., Cahan, R., Alvarez, M., Luo, Y., Chen, Z., "Canada Quant: New Options in Canada", Deutsche Bank Quantitative Strategy, November 23, 2010.

Lee, J., Stefek, D., [2006], "Do Risk Models Eat Alphas?", *Journal of Portfolio Management*, Vol. 34, No. 4.

Lee, W., [2010], "Risk-Based Asset Allocation: A New Answer to an Old Question?", forthcoming in *Journal of Portfolio Management*.

Luo, Y., Cahan, R., Jussa, J., Alvarez, M. [2010]. "Portfolios Under Construction: Volatility=1/N", Deutsche Bank Quantitative Strategy, June 16, 2010.

Luo, Y., Cahan, R., Jussa, J., Alvarez, M. [2010]. "QCD Model: DB Quant Handbook", Deutsche Bank Quantitative Strategy, July 22, 2010.

Luo, Y., Cahan, R., Jussa, J. Alvarez, M., [2011], "Portfolios Under Construction: Robust Factor Modeling", Deutsche Bank Quantitative Strategy, January 21, 2011.

Luo, Y., Jussa, J., Cahan, R., Alvarez, M., [2010], "Canada Quant: Introducing Canada Quantitative Strategy", Deutsche Bank Quantitative Strategy, October 24, 2010.

Markowitz, H., [1952], "Portfolio selection", *Journal of Finance*, 7:77-91

Mitton, Todd, and Keith Vorkink, 2007. "Equilibrium Underdiversification and the Preferenece for Skewness." *Review of Financial Studies*, vol. 20, no. 4

Scherer, B., [2010], "A New Look At Minimum Variance Investing", Working paper, SSRN: http://ssrn.com/abstract=1681306

Stevens, G., [1998], "On the Inverse of the Covariance Matrix in Portfolio Analysis", *The Journal of Finance*, Vol LIII, No. 5.

Tversky, Amos, And Daniel K., 1992, "Advances in Prospect Theory: Cumulative Representation of Uncertainty." *Journal of Risk and Uncertainty*, vol. 5, no. 4

Appendix A

MVP comovement with portfolios

Let $\boldsymbol{h_p}$ and σ_P be the holdings and volatility of a portfolio *P*, respectively. Let $\boldsymbol{h}_{MVP}$ and σ_{MVP} be the holdings and the volatility of the MVP portfolio. Let **V** be the variance-covariance matrix of asset returns then:

$$corr(r_{MVP}, r_P) = \frac{\boldsymbol{h}'_{MVP}\boldsymbol{V}\boldsymbol{h_p}}{\sigma_{MVP}\sigma_P} = \sigma^2_{MVP}\frac{\boldsymbol{1}'\boldsymbol{V}^{-1}\boldsymbol{V}\boldsymbol{h_p}}{\sigma_{MVP}\sigma_P}$$

$$corr(r_{MVP}, r_P) = \frac{\sigma_{MVP}}{\sigma_P}\sum_i h_{P,i}\ .$$

Appendix 1

Important Disclosures

Additional information available upon request

For disclosures pertaining to recommendations or estimates made on a security mentioned in this report, please see the most recently published company report or visit our global disclosure look-up page on our website at http://gm.db.com/ger/disclosure/DisclosureDirectory.eqsr.

Analyst Certification

The views expressed in this report accurately reflect the personal views of the undersigned lead analyst(s). In addition, the undersigned lead analyst(s) has not and will not receive any compensation for providing a specific recommendation or view in this report. Miguel-A Alvarez/Rochester Cahan/Zongye Chen/Javed Jussa/Yin Luo

Hypothetical Disclaimer

Backtested, hypothetical or simulated performance results discussed on page 10 herein and after have inherent limitations. Unlike an actual performance record based on trading actual client portfolios, simulated results are achieved by means of the retroactive application of a backtested model itself designed with the benefit of hindsight. Taking into account historical events the backtesting of performance also differs from actual account performance because an actual investment strategy may be adjusted any time, for any reason, including a response to material, economic or market factors. The backtested performance includes hypothetical results that do not reflect the reinvestment of dividends and other earnings or the deduction of advisory fees, brokerage or other commissions, and any other expenses that a client would have paid or actually paid. No representation is made that any trading strategy or account will or is likely to achieve profits or losses similar to those shown. Alternative modeling techniques or assumptions might produce significantly different results and prove to be more appropriate. Past hypothetical backtest results are neither an indicator nor guarantee of future returns. Actual results will vary, perhaps materially, from the analysis.

Regulatory Disclosures

1. Important Additional Conflict Disclosures

Aside from within this report, important conflict disclosures can also be found at https://gm.db.com/equities under the "Disclosures Lookup" and "Legal" tabs. Investors are strongly encouraged to review this information before investing.

2. Short-Term Trade Ideas

Deutsche Bank equity research analysts sometimes have shorter-term trade ideas (known as SOLAR ideas) that are consistent or inconsistent with Deutsche Bank's existing longer term ratings. These trade ideas can be found at the SOLAR link at http://gm.db.com.

3. Country-Specific Disclosures

Australia: This research, and any access to it, is intended only for "wholesale clients" within the meaning of the Australian Corporations Act.

EU countries: Disclosures relating to our obligations under MiFiD can be found at http://globalmarkets.db.com/riskdisclosures.

Japan: Disclosures under the Financial Instruments and Exchange Law: Company name - Deutsche Securities Inc. Registration number - Registered as a financial instruments dealer by the Head of the Kanto Local Finance Bureau (Kinsho) No. 117. Member of associations: JSDA, The Financial Futures Association of Japan. Commissions and risks involved in stock transactions - for stock transactions, we charge stock commissions and consumption tax by multiplying the transaction amount by the commission rate agreed with each customer. Stock transactions can lead to losses as a result of share price fluctuations and other factors. Transactions in foreign stocks can lead to additional losses stemming from foreign exchange fluctuations. "Moody's", "Standard & Poor's", and "Fitch" mentioned in this report are not registered as rating agency in Japan unless specifically indicated as Japan entities of such rating agencies.

New Zealand: This research is not intended for, and should not be given to, "members of the public" within the meaning of the New Zealand Securities Market Act 1988.

Russia: This information, interpretation and opinions submitted herein are not in the context of, and do not constitute, any appraisal or evaluation activity requiring a license in the Russian Federation.

Deutsche Bank Securities Inc.

North American location

Deutsche Bank Securities Inc.
60 Wall Street
New York, NY 10005
Tel: (212) 250 2500

Deutsche Bank Securities Inc.
One International Place
12th Floor
Boston, MA 02110
United States of America
Tel: (1) 617 217 6100

Deutsche Bank Securities Inc.
222 South Riverside Plaza
30th Floor
Chicago, IL 60606
Tel: (312) 537-3758

Deutsche Bank Securities Inc.
3033 East First Avenue
Suite 303, Third Floor
Denver, CO 80206
Tel: (303) 394 6800

Deutsche Bank Securities Inc.
1735 Market Street
24th Floor
Philadelphia, PA 19103
Tel: (215) 854 1546

Deutsche Bank Securities Inc.
101 California Street
46th Floor
San Francisco, CA 94111
Tel: (415) 617 2800

Deutsche Bank Securities Inc.
700 Louisiana Street
Houston, TX 77002
Tel: (832) 239-4600

International Locations

Deutsche Bank Securities Inc.
60 Wall Street
New York, NY 10005
United States of America
Tel: (1) 212 250 2500

Deutsche Bank AG London
1 Great Winchester Street
London EC2N 2EQ
United Kingdom
Tel: (44) 20 7545 8000

Deutsche Bank AG
Große Gallusstraße 10-14
60272 Frankfurt am Main
Germany
Tel: (49) 69 910 00

Deutsche Bank AG
Deutsche Bank Place
Level 16
Corner of Hunter & Phillip Streets
Sydney, NSW 2000
Australia
Tel: (61) 2 8258 1234

Deutsche Bank AG
Filiale Hongkong
International Commerce Centre,
1 Austin Road West,Kowloon,
Hong Kong
Tel: (852) 2203 8888

Deutsche Securities Inc.
2-11-1 Nagatacho
Sanno Park Tower
Chiyoda-ku, Tokyo 100-6171
Japan
Tel: (81) 3 5156 6770

Disclaimer

The information and opinions in this report were prepared by Deutsche Bank AG or one of its affiliates (collectively "Deutsche Bank"). The information herein is believed to be reliable and has been obtained from public sources believed to be reliable. Deutsche Bank makes no representation as to the accuracy or completeness of such information.

Deutsche Bank may engage in securities transactions, on a proprietary basis or otherwise, in a manner inconsistent with the view taken in this research report. In addition, others within Deutsche Bank, including strategists and sales staff, may take a view that is inconsistent with that taken in this research report.

Opinions, estimates and projections in this report constitute the current judgement of the author as of the date of this report. They do not necessarily reflect the opinions of Deutsche Bank and are subject to change without notice. Deutsche Bank has no obligation to update, modify or amend this report or to otherwise notify a recipient thereof in the event that any opinion, forecast or estimate set forth herein, changes or subsequently becomes inaccurate. Prices and availability of financial instruments are subject to change without notice. This report is provided for informational purposes only. It is not an offer or a solicitation of an offer to buy or sell any financial instruments or to participate in any particular trading strategy. Target prices are inherently imprecise and a product of the analyst judgement.

As a result of Deutsche Bank’s recent acquisition of BHF-Bank AG, a security may be covered by more than one analyst within the Deutsche Bank group. Each of these analysts may use differing methodologies to value the security; as a result, the recommendations may differ and the price targets and estimates of each may vary widely.

Deutsche Bank has instituted a new policy whereby analysts may choose not to set or maintain a target price of certain issuers under coverage with a Hold rating. In particular, this will typically occur for "Hold" rated stocks having a market cap smaller than most other companies in its sector or region. We believe that such policy will allow us to make best use of our resources. Please visit our website at http://gm.db.com to determine the target price of any stock.

The financial instruments discussed in this report may not be suitable for all investors and investors must make their own informed investment decisions. Stock transactions can lead to losses as a result of price fluctuations and other factors. If a financial instrument is denominated in a currency other than an investor's currency, a change in exchange rates may adversely affect the investment.

All prices are those current at the end of the previous trading session unless otherwise indicated. Prices are sourced from local exchanges via Reuters, Bloomberg and other vendors. Data is sourced from Deutsche Bank and subject companies.

Past performance is not necessarily indicative of future results. Deutsche Bank may with respect to securities covered by this report, sell to or buy from customers on a principal basis, and consider this report in deciding to trade on a proprietary basis.

Derivative transactions involve numerous risks including, among others, market, counterparty default and illiquidity risk. The appropriateness or otherwise of these products for use by investors is dependent on the investors' own circumstances including their tax position, their regulatory environment and the nature of their other assets and liabilities and as such investors should take expert legal and financial advice before entering into any transaction similar to or inspired by the contents of this publication. Trading in options involves risk and is not suitable for all investors. Prior to buying or selling an option investors must review the "Characteristics and Risks of Standardized Options," at http://www.theocc.com/components/docs/riskstoc.pdf If you are unable to access the website please contact Deutsche Bank AG at +1 (212) 250-7994, for a copy of this important document.

The risk of loss in futures trading, foreign or domestic, can be substantial. As a result of the high degree of leverage obtainable in futures trading, losses may be incurred that are greater than the amount of funds initially deposited.

Unless governing law provides otherwise, all transactions should be executed through the Deutsche Bank entity in the investor's home jurisdiction. In the U.S. this report is approved and/or distributed by Deutsche Bank Securities Inc., a member of the NYSE, the NASD, NFA and SIPC. In Germany this report is approved and/or communicated by Deutsche Bank AG Frankfurt authorized by the BaFin. In the United Kingdom this report is approved and/or communicated by Deutsche Bank AG London, a member of the London Stock Exchange and regulated by the Financial Services Authority for the conduct of investment business in the UK and authorized by the BaFin. This report is distributed in Hong Kong by Deutsche Bank AG, Hong Kong Branch, in Korea by Deutsche Securities Korea Co. This report is distributed in Singapore by Deutsche Bank AG, Singapore Branch, and recipients in Singapore of this report are to contact Deutsche Bank AG, Singapore Branch in respect of any matters arising from, or in connection with, this report. Where this report is issued or promulgated in Singapore to a person who is not an accredited investor, expert investor or institutional investor (as defined in the applicable Singapore laws and regulations), Deutsche Bank AG, Singapore Branch accepts legal responsibility to such person for the contents of this report. In Japan this report is approved and/or distributed by Deutsche Securities Inc. The information contained in this report does not constitute the provision of investment advice. In Australia, retail clients should obtain a copy of a Product Disclosure Statement (PDS) relating to any financial product referred to in this report and consider the PDS before making any decision about whether to acquire the product. Deutsche Bank AG Johannesburg is incorporated in the Federal Republic of Germany (Branch Register Number in South Africa: 1998/003298/10). Additional information relative to securities, other financial products or issuers discussed in this report is available upon request. This report may not be reproduced, distributed or published by any person for any purpose without Deutsche Bank's prior written consent. Please cite source when quoting.