Minimum Variance and Tracking Error: Combining Absolute and Relative Risk in a Single Strategy

Matthew Van Der Weide

FACTSET

Minimum Variance and Tracking Error: Combining Absolute and Relative Risk in a Single Strategy

Contents

Introduction

During the ongoing financial crisis, it has become clear that investing in equities with an eye toward the long-term return premium can lead to substantial periods of drawdown. It makes sense, therefore, to see a renewed interest in products and strategies aimed at trying to capture this premium at a potentially lower risk level. This strategy is not something new; the low volatility phenomenon was described in the early 1990s by Haugen and Baker. More recently we have seen a related but slightly different approach, namely, minimum variance, as described by Clarke, de Silva and Thorley (2006). Not only have managers deployed these strategies, but such is their popularity that we have recently seen these techniques being deployed by benchmark providers and provided as commercially available ETFs.

Although the two concepts are related and share a common goal (a better risk return profile), there are slight nuances in either approach. Low volatility strategies describe the outperformance of stocks with a lower price fluctuation. Minimum variance takes into account correlations between assets and looks at the volatility of the portfolio rather than the individual securities. Rather than selecting stocks with low and potentially shorting securities with a high (idiosyncratic) volatility, minimum variance focuses on constructing generally long-only portfolios in a way that minimizes the absolute risk.

Here, we will consider minimum variance investing, not as an advocate, but to examine some of the characteristics of these strategies and how they can potentially be blended with traditional passive market cap weighted investing to improve returns.

Absolute vs. Relative Risk

In the investment process, we can often define two levels of decision making. The first decision is at the asset class level, the second on the sub asset class or portfolio level. On these two levels we generally also find two notions of risk: absolute and relative. At the (tactical) asset allocation level, managers decide between asset classes, in general on the basis of a risk and return in an absolute sense. A broad benchmark can provide the long-term risk and return characteristics. An equity index, for example, traditionally offers a higher return than a fixed income benchmark; however, this often comes at a higher risk. Therefore a combination of X% in equities and 100-X% in fixed income can be a compelling mix to match a firm's risk-return profile.

Subsequent to asset class level mix, firms must decide how or where to invest: managing internally or externally, employing a single manager or multiple managers? These decisions are typically made on a benchmark relative basis. Does a particular manager have the skill to beat the benchmark, and, if so, at what risk level? In other words, rather than

looking at value at risk or a Sharpe ratio as for an asset class, this decision would be based on tracking error and information ratio.

Within this context, the minimum variance strategy can offer another interesting alternative. Traditionally when looking at asset allocation and more specifically equities, we would look at a market cap weighted index to compare it to other asset classes. In that light, we could see a minimum variance equity strategy either as a separate asset class or as a different style within the asset class. Traditionally we would also consider the manner in which minimum variance portfolios are constructed, focusing on absolute risk, which is typical of an asset class decision rather than a decision within the equities asset class.

What we will consider is a united approach, making a decision on both fronts at the same time, minimizing risk absolute risk while constraining relative risk.

Approach

In our research, we define our universe as a European index consisting of roughly 600 mid- and large-cap companies over the period of 2001-2011. To construct our portfolio, we used an optimizer, and we purchased securities from the index only. To forecast risk, we used a statistical factor-based risk model. We programmed the optimizer to minimize the absolute risk of the portfolio. The only constraint in the optimization was on cash, ensuring the portfolio was fully invested at all times. We ran our simulation from December 2000 to December 2011, rebalancing the portfolio every six months. This gave us our base case, a portfolio invested in equities with the lowest predicted standard deviation. As expected, our portfolio had a significant tracking error against the European index. We then ran six more strategies in which we added constraints on the predicted tracking error against the index. The tracking error limits were set at 25, 20, 15, 10, 5, and 0.5. The outer extremes of these limits functioned as controls to our approach. With a tracking error limit of 0.5, we expected to closely track the index, while with a tracking error limit of 25, we expected results very similar to our base case. Even though we expected a significant tracking error against the index, we considered it unlikely that it would surpass this level. In hindsight looking at the predicted tracking errors of the base scenario, we indeed observed that at only in December 2008 our predicted tracking error significantly exceeded the value of 25.

We did not set a limit on the number of companies. The most concentrated of the strategies had on average 158 and as a minimum 85 names, and we considered this to be diversified enough to not set a constraint.

Results

Chart 1: Cumulative Return Chart

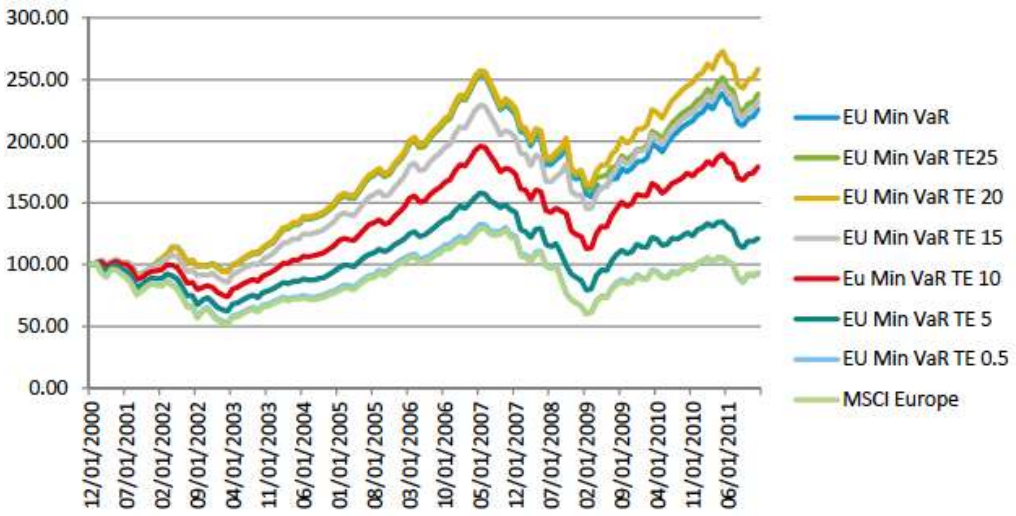

Looking at the cumulative return chart, we observe that over the complete time frame the index is the least performing strategy. To put this in context, we have two severe bear periods (the technology bubble and the financial crisis) that would favor a strategy focused on downside protection. Furthermore, we see that our control strategies behave as expected. The strategy with a 0.5% tracking error constraint tracks the index very closely, and the tracking error constraint of 25 gives very similar results as the base strategy with no constraint on tracking error.

Table 1: Annualized Return

Strategy	2001	2002	2003	2004	2005	2006	2007	2008	2009	2010
EU Min VaR	-2.53	3.36	17.18	25.16	22.00	26.68	-2.83	-23.84	8.58	20.63
EU Min VaR TE 25	-2.77	3.32	17.37	25.40	21.76	26.93	-1.99	-23.00	12.07	20.72
EU Min VaR TE 20	-2.77	3.32	18.75	25.40	21.76	26.93	-1.99	-23.00	20.17	20.72
EU Min VaR TE 15	-2.81	-4.29	17.22	23.44	20.85	26.69	-1.66	-22.88	23.10	18.44
EU Min VaR TE 10	-5.51	-13.84	16.41	21.36	20.85	26.69	-1.66	-28.26	26.48	12.15
EU Min VaR TE 5	-10.82	-22.11	16.03	16.92	23.18	23.66	-0.73	-37.23	30.35	10.06
EU Min VaR TE 0.5	-15.06	-29.18	16.06	12.94	26.75	20.33	2.27	-43.93	33.30	9.89
MSCI Europe	-15.26	-30.50	15.76	12.65	26.68	20.18	3.17	-43.29	32.55	11.75

Moreover we see that a tighter tracking error constraint brings the results closer to that of the benchmark. The exception seems to be with the strategies with a 15 and 20% tracking error constraint that from financial crises onwards seem to catch up, or even overtake, the base strategy. Looking at the results more closely, as depicted in Table One, we see that in 2009 the looser strategies seemed to benefit from having a larger exposure to financials; given their volatility these were largely ignored by minimum variance strategies.

Table 2: Return Statistics

Portfolio	Annualized Return	Annualized Std Dev	Beta	Alpha	Sharpe Ratio	Information Ratio	Max Drawdown
EU Min VaR	7.71	11.19	0.45	0.64	0.44	0.66	- 38.98
EU Min VaR TE 25	8.22	10.98	0.45	0.68	0.49	0.72	- 36.84
EU Min VaR TE 20	9.03	11.10	0.46	0.74	0.56	0.80	- 35.85
EU Min VaR TE 15	7.98	11.33	0.50	0.66	0.46	0.75	- 36.70
EU Min VaR TE 10	5.46	11.94	0.62	0.46	0.22	0.70	- 42.42
EU Min VaR TE 5	1.77	14.02	0.81	0.17	- 0.08	0.55	- 49.89
EU Min VaR TE 0.5	- 0.72	16.81	1.00	- 0.01	- 0.21	- 0.15	- 54.98
MSCI Europe	- 0.59	16.86	1.00	0.00	- 0.20	#N/A	- 53.72

Looking at Table Two, all strategies except our control case with a tracking error constraint of 0.5% seem to outperform over the full horizon and show information ratios that by themselves could be a reason to use such strategies for longer term investment. However, our aim was to produce strategies that can offer an equity premium at lower risk, so we should consider not information ratio but rather a Sharpe ratio. We clearly see that minimum variance portfolios have a much better Sharpe ratio and also have a lower maximum drawdown.

Table 3: Up and Down Statistics

Description	% Positive Returns	% Negative Returns	Annualized Upmarket Return	Annualized Downmarket Return	Annualized Positive Return	Annualized Negative Return
EU Min VaR	67.42	32.58	32.33	-14.42	21.56	-11.39
EU Min VaR TE 25	68.18	31.82	33.23	-14.27	21.56	-10.97
EU Min VaR TE 20	67.42	32.58	35.29	-14.44	22.44	-10.95
EU Min VaR TE 15	65.91	34.09	36.21	-16.88	21.99	-11.48
EU Min VaR TE 10	60.61	39.39	39.38	-23.20	21.11	-12.93
EU Min VaR TE 5	56.82	43.18	44.31	-31.61	22.22	-16.74
EU Min VaR TE 0.5	53.79	46.21	51.88	-38.70	24.98	-20.56
MSCI Europe	53.79	46.21	52.35	-38.74	25.18	-20.58

So where do these strategies outperform? Earlier research describes that these strategies “win by not losing,” or “not losing as much.” To confirm, in Table Three we show the average up and down market returns, which tend to favor minimum variance strategies. We first see that there are more positive returns. The higher number of positive returns, however, comes at a cost; the average positive return is somewhat dampened, further illustrated in Chart Two. We also see this further pronounced on the down side where we have both fewer negative returns and lesser magnitude. Those stocks favored by the minimum variance bias delivered more often while being less volatile, a confirmation of the “win by losing less” adage.

Chart 2: Annualized Up and Down Market Return

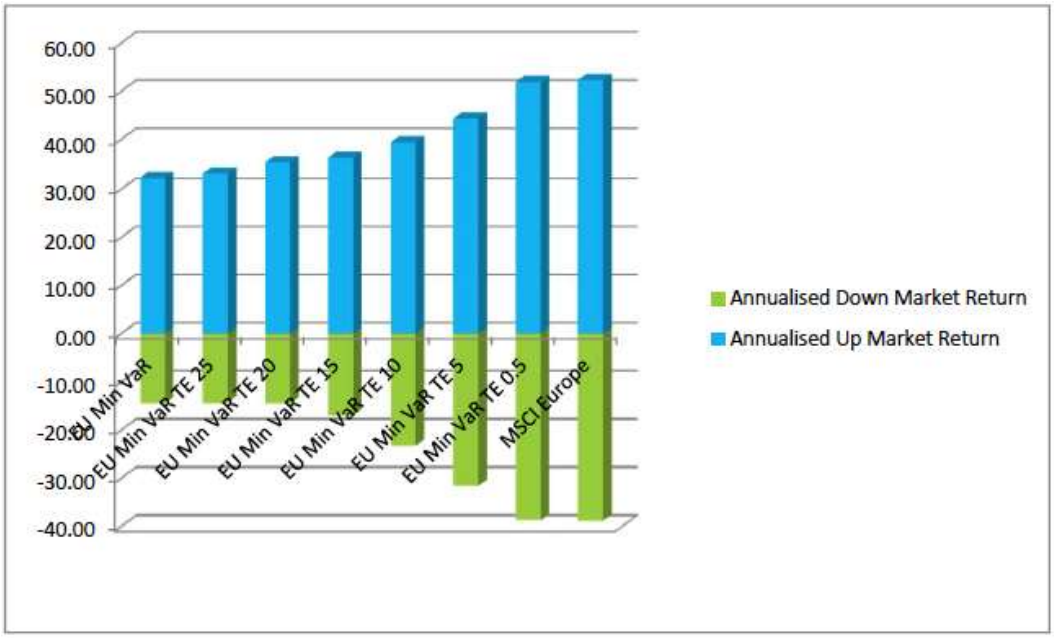

In the up and down market returns, we see a steady pattern in decreasing up- and downside potential with the looser strategies, showing our findings are robust. If there is any discontinuity, it lies potentially in the strategy with a tracking error constraint of 20, where we see the lowest down market return, even lower than the base strategy. This also seems to have helped in the period after 2008 looking at the total cumulative return. The differences are minimal, though, which is also explained by the fact that over the whole period we saw only two periods where the constraint of a tracking error of 20 was actually hit; for other periods we expect to end up with a very similar portfolio as our base strategy.

Conclusion and Application

Horizon and single region analysis notwithstanding, our research seems to be in line with earlier findings: minimum variance strategies offer the possibility of capturing an equity premium at a lower risk or at a better Sharpe ratio. Our primary conclusion, however, is that we believe it is possible to combine both absolute and relative risk in a single strategy. By setting a tracking error constraint we can bring the minimum variance portfolio closer to the market cap weighted index. In effect, we can use tracking error as an appetite measure for risk away from the minimum variance portfolio. In this case, there was an optimum return when setting a tracking error constraint of 20; however, the results are not significantly different enough to make a strong conclusion. Although this might be an additional benefit worth researching further, our main conclusion remains the fact that it is possible to scale between absolute and relative risk.

By itself, this conclusion can be an interesting alternative in an asset class allocation or market timing. Rather than increasing equity exposure by investing more in nominal terms, investors can gain more “traditional” equity exposure by tilting more towards relative risk. Rather than this being a discrete on/off switch, we see that this can be a continuous decision scale, with a connection to conviction level. With low expectations of a market rally, staying on the higher tracking error limit can be a more conservative approach, whereas stronger convictions may lead to a tighter tracking error constraint.