9-2019

Volatility timing under low-volatility strategy

Poh Ling NEO

Chyng Wen TEE
Singapore Management University, cwtee@smu.edu.sg

Volatility Timing under Low-Volatility Strategy

POH LING NEO CHYNG WEN TEE

September 20, 2019

Abstract

The authors show that the slope of the volatility decile portfolio's return profile contains valuable information that can be harvested to time volatility and market condition. During good market condition, high-volatility portfolio produces the highest return, and vice versa. The authors proceed to devise a volatility timing strategy based on statistical tests on the slope of the volatility decile portfolio return profile. Volatility timing is achieved by being aggressive during strong growth periods, while being conservative during market downturns. Superior performance is obtained, with a 30% increase in Sharpe ratio and an order of magnitude improvement on cumulated wealth. The authors also demonstrate that stocks in the high-volatility portfolio are more strongly correlated compared to stocks in the low-volatility portfolio. The profitability of the volatility timing strategy can be attributed to holding a diversified portfolio during bear markets, while holding a concentrated growth portfolio during bull markets.

Poh Ling Neo is a senior lecturer at the School of Business, Singapore University of Social Sciences, 463 Clementi Road, Singapore 599494. Email: plneo@suss.edu.sg

Chyng Wen Tee is an associate professor at the Lee Kong Chian School of Business, Singapore Management University, 50 Stamford Road, 05-01, Singapore 178899. Email: cwtee@smu.edu.sg

Low-volatility has been established for some time now as one of the winning factors in smart beta investing. Holding low-volatility stocks can provide superior risk-adjusted returns compared to holding high-volatility stocks. In fact, high-volatility portfolio generates very poor returns over time. The benefit of holding a low-volatility portfolio can be attributed to the compounding effect—a more stable return coupled with a lower volatility enhances the long-term risk-adjusted performance in investment.

The fact that low risk stocks have higher expected returns is an interesting anomaly in the field of finance. It is also a persistent anomaly—research has shown that US stocks with high volatility earned abnormally low returns over the 1963-2000 period. Blitz and van Vliet (2007) show that this anomaly also extends to all equity markets in the world.

Most asset pricing models postulate a positive relationship between a stock portfolio's expected returns and volatility. If high-volatility stocks do indeed lead over time to higher returns, then one would expect stocks that produce returns with low-volatility should generate lower returns over time. However, Black, Jensen and Scholes (1972) and Haugen and Heins (1975) were among the first to show that high-volatility (or high-beta) stocks had lower returns than equilibrium suggests. On the contrary, low-volatility (or low-beta) stocks earned higher returns than conventionally expected. In short, they report a strong negative relationship between return and volatility in both the stock and the bond market.

Since then, there has been a long string of financial research that models stock return volatilities as negatively correlated with stock returns (see, for instance, Cox and Ross (1976)). A number of empirical analyses have also shown significant negative relationship between expected returns and volatilities, lending support to claims that stock returns are negatively correlated with stock volatilities (see Li, Yang et al. (2005)).

Our objective in this work is to explore the time-varying component of the low-volatility strategy, and study the feasibility of augmenting its performance by volatility timing. We show that although a low-volatility portfolio does generate higher risk-adjusted returns relative to high-volatility portfolio when aggregated over a multiple-decade period studied extensively in the low-volatility literature, the return profiles of the volatility decile portfolio are in fact time-varying. During good market condition, high-volatility portfolio produces the highest return, while during bad market condition, low-volatility portfolio exhibits the least negative return. We use PCA analysis to show that the variation in the shape of the volatility decile portfolio's return profile comprises of 10% variation in slope, which contains important information as a lead indicator on market condition.

Based on this insight, we proceed to devise a volatility timing strategy which switches between holding low-, mid-, or high-volatility portfolio based on statistical tests on the slope of the return profile of the volatility decile portfolio. Superior performance is obtained, with an increment of 30% in Sharpe ratio and an order of magnitude improvement in cumulated wealth. This is achieved via successful volatility timing by being aggressive in holding a high-volatility portfolio during strong growth periods, while being conservative in holding a low-volatility portfolio during market downturns.

Correlation analysis is also performed to to investigate the role of diversification in our strategy. We found that stocks in the high-volatility portfolio are consistently more strongly correlated compared to stocks in the low-volatility portfolio. In other words, the profitability of the volatility timing strategy can be attributed to holding a diversified portfolio during bear markets, while holding a concentrated growth portfolio during bull markets. This fits consistently within the framework and empirical studies of market correlation structure and its impact on diversification (see Chua, Kritzman and Page (2009)). Thus, our results connect two separate strands of research in the portfolio management space.

Low Volatility Strategies

Contrary to basic finance principles, high-beta and high-volatility stocks have long underperformed low-beta and low-volatility stocks in more than 40 years of US data (see Baker, Bradley and Wurgler (2011)). A number of portfolio strategies have been developed over time to exploit this negative relationship between volatility and return. These include the minimum variance portfolio, the low-volatility portfolio, the maximum diversification portfolio, the bet-against-beta (BAB) portfolio, and the volatility-managed portfolio.

The minimum variance portfolio by Clarke, de Silva and Thorley (2006) uses the full covariance matrix, coupled with an optimization algorithm to demonstrate that minimum-variance portfolios on the 1,000 largest US stocks over the 1968-2005 period can achieve a volatility reduction of about 25%, while delivering comparable or even higher average returns than the market portfolio.

The low-volatility portfolio by Blitz and van Vliet (2007) and Baker and Haugen (2012) use only the diagonal of the historical covariance matrix. They find that portfolios of stocks with the lowest historical volatility are associated with Sharpe ratio improvements and statistically significant improvement in alpha. Robustness tests clearly indicate the existence of a volatility effect—low-risk stocks exhibit significantly higher risk-adjusted returns than the market portfolio, while high-risk stocks significantly underperform on a risk-adjusted basis.

The maximum diversification portfolio by Choueifaty and Coignard (2008) uses an objective function that maximizes the ratio of weighted-average asset volatilities to portfolio volatility. Like minimum variance, maximum diversification portfolios equalize each asset's marginal contributions, given a small change in the asset's weight. They show that the resulting maximum diversification portfolios have higher Sharpe ratios than the market capitalization weighted indices, and exhibit both lower volatilities and higher returns in the long run.

Apart from using historical volatility, it is also common in standard asset pricing models to decompose volatility into beta (β), which measures covariance with the market or other common factors, and an idiosyncratic volatility component. Both methods have also been explored in the low-volatility strategies literature. The bet-against-beta (BAB) portfolio by Frazzini and Pedersen (2014) argues that because constrained investors bid up high-beta assets, high-beta is consequently associated with low-alpha. A bet-against-beta (BAB) factor, which is long leveraged low-beta assets and short high-beta assets, produces significant positive risk-adjusted returns.

Clearly, ranking stocks on their historical volatility is related to ranking stocks on their historical capital asset pricing model betas. In addition to beta, Ang et al. (2006) show that stocks with low idiosyncratic volatility significantly outperform stocks with high idiosyncratic volatility. In a subsequent study, Ang et al. (2009) expand their research where they apply their method to a global market and arrive at the same conclusion. Stocks with high idiosyncratic volatility, calculated as the variance of the residuals in a Fama-French regression, have bad performance whereas stocks with low idiosyncratic volatility performs well.[1]

More recently, Moreira and Muir (2017) develop this strategy further in the form of volatility-managed portfolios. They construct portfolios that scale monthly returns by the inverse of their previous month's realized variance, decreasing risk exposure when variance was high recently and vice versa. They call these volatility-managed portfolios[2]. They provide evidence showing that volatility managed portfolios that take less risk when volatility is high produce large alphas, increase Sharpe ratios, and produce large utility gains for mean-variance investors.

In this work, we begin by extending the low-volatility portfolio analysis to include recent period. We will use this as a base case to illustrate our volatility timing strategy. We focus on the US equity market. Stock return data are obtained from the Center for Research in Security Prices (CRSP). Our US equity data include all available common stocks on CRPS between January 1963 and December 2016. The Fama-French factor data are downloaded from Kenneth R. French's web site. In our robustness tests, we consider alphas with respect to the market factor and factor returns based on size (SMB) and book-to-market (HML). Other US firm-level data are from obtained COMPUSTAT.

Following the standard low-volatility portfolio construction procedure, we rank stocks by their market capitalization. The top 1000 stocks with the greatest market capitalization are selected to form the volatility decile portfolios at each period. Every month, we calculate the realized volatility of each stock over the past month, and use this to group them into decile portfolio based on their volatility. This procedure is carried out across the entire 1963-2016 period included in our study. The mean realized return and volatility profiles, along with Sharpe ratios, are presented in Exhibit I. The upper figure shows the average realized volatility of the volatility decile portfolio on the month following portfolio construction, while the lower figure shows the average realized return. As expected, when aggregated across the 1963-2016 period studied, there is a distinct negative relationship between return and volatility, manifesting in the downward slope in the return decile portfolio. This finding is consistent with existing literature in this area, where high-volatility portfolio significantly underperforms, while low-volatility portfolio outperforms, in terms of both raw return and risk-adjusted returns (Sharpe ratio).

Exhibit I: Volatility and return of decile portfolio

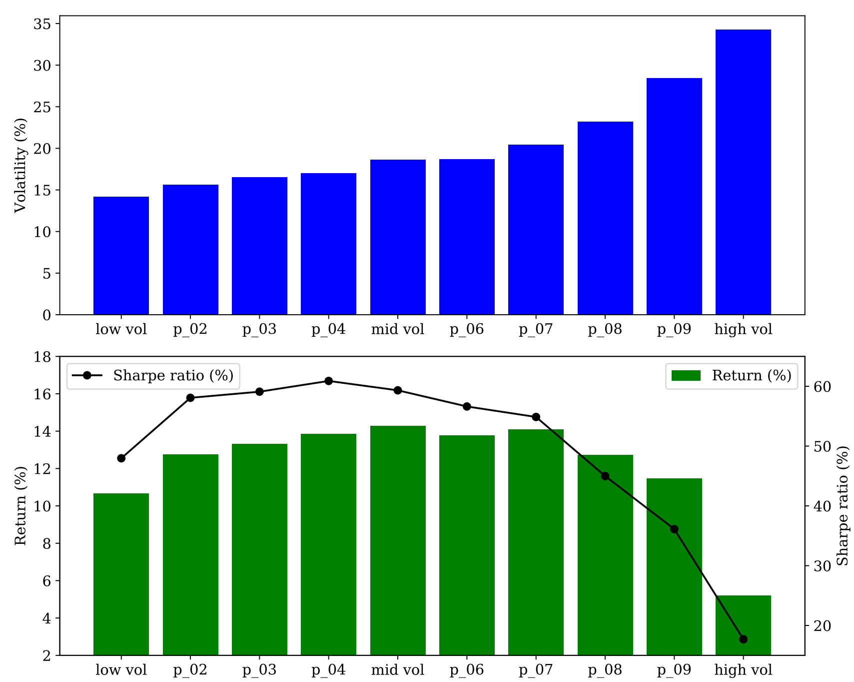

Numerous research efforts have been devoted to explore possible explanations for this low-volatility anomaly. Beveratos et al. (2017) attribute this to high dividends by demonstrating that there is a strong correlation between low-volatility and high-dividend yield. Li, Sullivan and Garcia-Feijóo (2016) argue that the relatively high returns of low-volatility portfolios are likely driven by market mispricing volatility characteristic. Driessen et al. (2019) show that low-volatility stock portfolios have negative exposure to interest rates, and their outperformance can be partially attributed to the premium for interest rate exposure.

In this paper, our main focus is not to furnish possible explanation for the low-volatility effect, but rather to investigate its behavior and persistence. Next, we will explore the time-varying characteristic of the risk-return relationship, and assess the feasibility of using this as a lead indicator to time volatility.

Time-varying Return Profile

First, we investigate the variation in the shape of the volatility decile portfolio's return profile by performing principal component analysis (PCA) on the time series. PCA is a statistical tool widely employed to investigate the number of factors involved in explaining shape movement. Using PCA to investigate the factor structures of movement in the return profile of the volatility decile portfolio will allow us to identify the main drivers of the shape variation and their relative importance. Exhibit II plots the eigenvectors of the three most significant principal components, which collectively account for approximately 98% of the variation in the shape of the volatility decile portfolio's return profile.

Exhibit II: PCA of volatility decide portfolio return profile

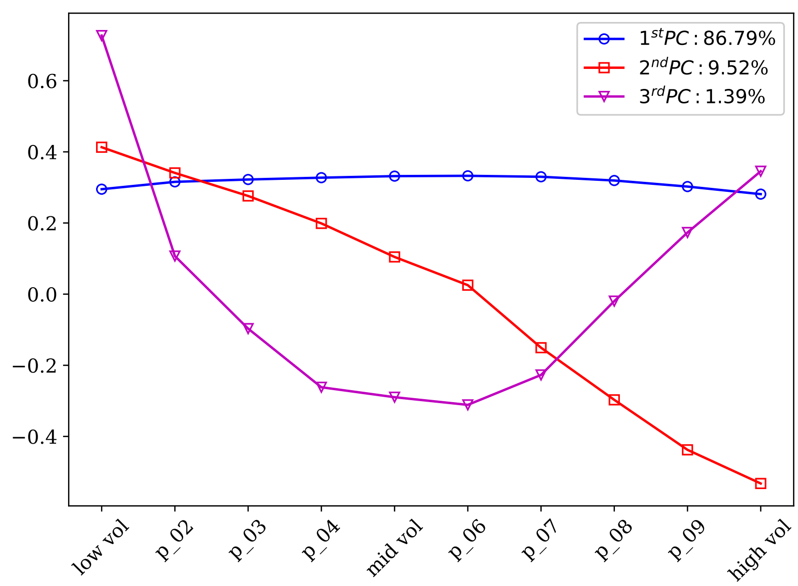

Based on the shape of the three principal components in Exhibit II, it should be clear that the first principal component (PC) captures parallel movement in the return profile common across the volatility decile portfolio. The second PC captures the change in the slope, while the third PC accounts for the variation in the curvature. The explanatory power of each PC is measured by the magnitude of the eigenvalues, and are also labeled in Exhibit II. The first PC accounts for 86.79% of the return profile's shape variation, followed by the second PC, which also plays a prominent role at 9.52%. This is a strong indication that the slope of the volatility decile portfolio's return profile also varies over time, and has important contribution to the shape of the return profile.

Clearly, low-volatility stocks are not expected to consistently outperform high-volatility stocks across all market regimes. For instance, it should be intuitive to expect that, under periods of strong growth, certain high volatility stocks should be expected to outperform the rest of the market. A quick empirical study confirms this. To illustrate the time-varying characteristic of the return profile, we divide the period studied (1963-2016) into "good" and "bad" market conditions by comparing monthly market return against monthly Treasury Bill return. Following Bessembinder (2018), we use Treasury bills as a benchmark to measure stock market performance—if market return outperforms T-bill return, we classify the period as "good" market condition, while the opposite is true for "bad" market condition. We proceed to calculate the aggregated volatility and return profile separately under "good" and "bad" market conditions. There are 647 monthly observations, of which 380 (59%) months were "good", and 267 (41%) months were "bad".

Exhibit III: Volatility and return profile under different market conditions.

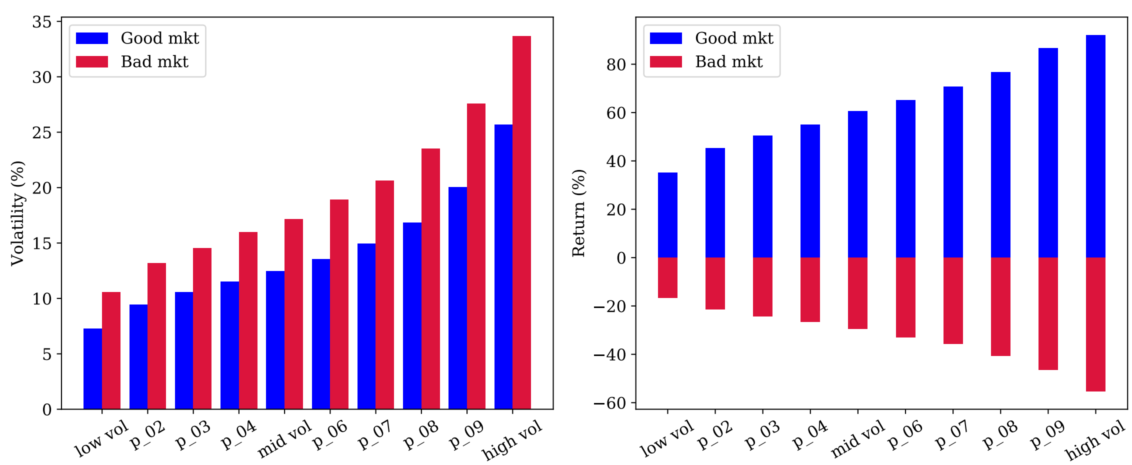

Exhibit III plots the average realized volatility (left figure) and return (right figure) of the volatility decile portfolio under good and bad market conditions. As expected, the volatilities of all portfolios under bad market condition are higher when compared to those during good market condition. Nevertheless, under both good and bad market conditions, the volatility trend stays the same, increasing monotonically across the volatility decile portfolios, and only differ in their magnitude.

Return profile is where the interesting time-varying characteristic of the risk-return relationship manifests. All portfolios have better return performance under good market condition when compared to the bad market conditions. However, under good market condition, the high-volatility portfolio is seen to have the highest return. On the other hand, under bad market condition, the low volatility portfolio becomes the best performer, with the smallest loss registered. This results in an opposite risk-return trends under different market conditions—high volatility stocks are associated with high return during good market condition and bad return during bad market condition, and vice versa. This implies that market condition has important information that can be used to extend mean-variance investment analysis.

An important observation on our results is that the slope of the volatility decile portfolio return profile will change sign depending on the prevailing market condition. When the shape is download sloping, indicating the negative relationship between return and volatility, holding low-volatility portfolio is optimal. On the other hand, when market is experiencing strong growth during boom time, the return slope becomes upward sloping, which in turn indicates a positive relationship between return and volatility. During this time, it is optimal to hold high-volatility portfolio instead. Finally, when aggregated across both market conditions, high-volatility portfolio underperforms, while low-volatility portfolio overperforms, resulting in the well-known risk-return relationship in Exhibit I.

We have shown that the slope of the volatility decile portfolio's return profile can become upward sloping from time to time, indicating that at certain period it is in fact optimal to hold high-volatility stocks instead. Generally, high-volatility (or high-beta) stocks tend to belong to growth industries, and eliminating them would naturally result in a more value-oriented portfolio. For this reason, low-volatility (or low-beta) stocks will naturally underperform in a bull market and outperform in a bear market. Is it possible to time volatility and further improve return by targeting high-volatility stocks during periods of growth, and holding low-volatility during market downturns? We attempt to answer this question in the next section.

Adding Volatility Timing to the Strategy

It is widely accepted that an equity portfolio's volatility is generally dominated by its co-movement with the market. It is therefore natural to expect that shifting from a portfolio that is concentrated in the market factor to a more diversified portfolio by adding other factors with attractive premia, such as low-volatility, would improve the portfolio's risk-return profile. In other words, the main benefit of holding a low-volatility portfolio is not merely limited to reducing volatility exposure, but also to reallocate exposure from the market factor to other sources of return. Lower volatility is simply one of the consequences of this risk diversification. However, it is also commonly acknowledged that the primary cost of low-volatility strategies is underperformance in upward-trending market environments, which investors may find undesirable.

In the context of portfolio optimization, diversification in a bull market is not a desirable characteristic. On the contrary, investors generally prefer unification on the upside. However, Chua, Kritzman and Page (2009) has demonstrated that observed correlations are higher than normal correlations on the downside and lower on the upside. In other words, diversification works well during good times—when it is not needed—yet suffers a significant drop in efficiency during down markets. Conventional approaches to portfolio construction generally ignore this correlation asymmetry, which will inevitably cause static portfolios to be unnecessarily diversified when market is experiencing strong growth, yet being concentrated during market downturns.

If the main benefit of holding low-volatility portfolio stems from diversification, and yet naively following it will subject investors to underperformance during upward-trending market (see Chow, Hsu and Li (2014)), it could potentially be beneficial to incorporate a timing strategy to further improve the performance of low-volatility portfolio. Research by Fleming, Kirby and Ostdiek (2001) on the benefit of market timing have demonstrated that efficient timing strategies can potentially outperform static portfolios. The purpose of adding timing strategy is to allow one to optimally hold a portfolio in which assets diversify each other more on the downside, but move together more on the upside.

Here we formulate a volatility timing strategy which keeps track of the shape of the volatility decile portfolio's return profile. We define slope as the return difference between the return of the high-volatility portfolio and low-volatility portfolio, i.e.

$$\text{Slope} = r_{\text{high-vol}} - r_{\text{low-vol}}$$

where $r_{\text{high-vol}}$ and $r_{\text{low-vol}}$ are the returns of the high-volatility and low-volatility portfolios respectively. When market condition is good (bad), this slope is expected to be positive (negative), as we have demonstrated in the previous section. Recall that our PCA analysis has revealed that changes in the slope of the return profile accounts for approximately 10% of the overall variation—there are useful information to be harvested in the slope parameter.

Exhibit IV shows the distribution of slope parameter under good and bad market conditions. Histograms and fitted kernel density estimates (KDE) for the slope under both (good and bad) market conditions are presented. The good (bad) market condition has a mean slope parameter of 3.2% (-4.8%), with a standard deviation of 6.2% (6%). As one would expect, the average slope is positive (negative) under good (bad) market condition. Although there is a certain amount of overlap in the slope distributions under the two market conditions, it is noteworthy that when the slope parameters are high (right tail), market condition is unequivocably good, while when the slope parameters are low (left tail), market condition is clearly bad. Our calculation also reveals an autocorrelation of 16.1% in the slope parameter in these cases. This observation indicates the feasibility of using a hypothesis testing framework to generate the signals for volatility timing. When the slope parameter is sufficiently high (or low), if it can be established that the magnitude is statistically significant, we can then infer that the market condition is good (or bad), and use this information to hold the right portfolio.

Exhibit IV: Slope distribution of the volatility decile portfolio's return profile.

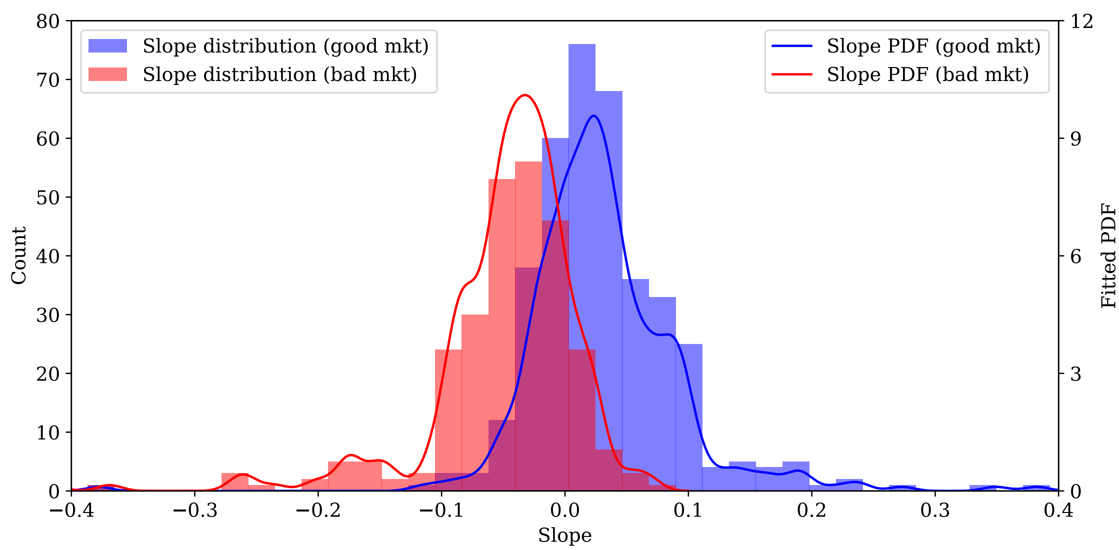

Given this insight, one can readily expect a volatility timing strategy that dynamically allocate portfolio weights based on the slope of the volatility decile portfolio's previous month's return profile to outperform a static strategy that always hold a specific (e.g. low-volatility) portfolio[3]. The main goal of our volatility timing strategy is to selectively hold high-volatility stocks during bull market to benefit from strong stocks growth, while switching to hold low-volatility stocks during market downturns to benefit from diversification.

We begin with the simplest version of the volatility timing strategy, which will be referred to as "vol timing" in subsequent discussions. In this formulation, our null hypothesis is that the volatility decile portfolio's return profile is downward sloping, since the aggregated return profile takes this shape (see

Exhibit I). Hence, the default portfolio to hold is the low-volatility portfolio. If the slope becomes positive in a way that is statistically significant for a given month, we reject the null and accept the alternative hypothesis—that the market condition is good. We will then switch to hold the high-volatility portfolio instead for the following month. As and when the slope parameter is no longer statistically significant, we will fall back to the default low-volatility portfolio. Using the Student's t-distribution with a moving window sample size of $n = 12$ and a significance level of $\alpha = 0.1\%$, our results demonstrate that timing volatility in this way can generate superior return compared to simply holding low volatility stocks. Holding the low-volatility portfolio as default along with setting a high significance level threshold is a sensible choice, since all previous empirical studies have shown that high-volatility stocks underperform in the long run. In other words, it is better to occasionally miss out on periods of strong market growth than to be holding high-volatility stocks during market downturns.

We measure the performance of our volatility timing strategy via wealth plot, realized volatility, and excess returns. These are presented in Exhibit V. For ease of comparison, we also include the performance of the static strategies of holding low-, mid-, and high-volatility portfolios. The performance of these static strategies are well documented in the literature—high-volatility portfolio has the lowest return, cumulated wealth, and the highest realized volatility. Low-volatility portfolio performs noticeably better: it has the smallest realized volatility, higher return and Sharpe ratio, and a much higher cumulated wealth compared to the high-volatility portfolio. The performance of the volatility timing strategy elaborated above is plotted as "Vol timing" in Exhibit V. This strategy holds low-volatility portfolio by default, and only switches to hold high-volatility portfolio when the slope is significantly positive, indicating a good market condition. Note that the strategy's realized volatility and return are both higher than the low-volatility portfolio due to occasionally holding high-volatility stocks when the market condition is good. However, the Sharpe ratio is in fact marginally lower compared to the low-volatility portfolio, implying that this simple version of volatility timing strategy does not generate superior risk-adjusted return compared to the static low-volatility portfolio.

Exhibit V: Performance and wealth plots of the volatility timing strategy.

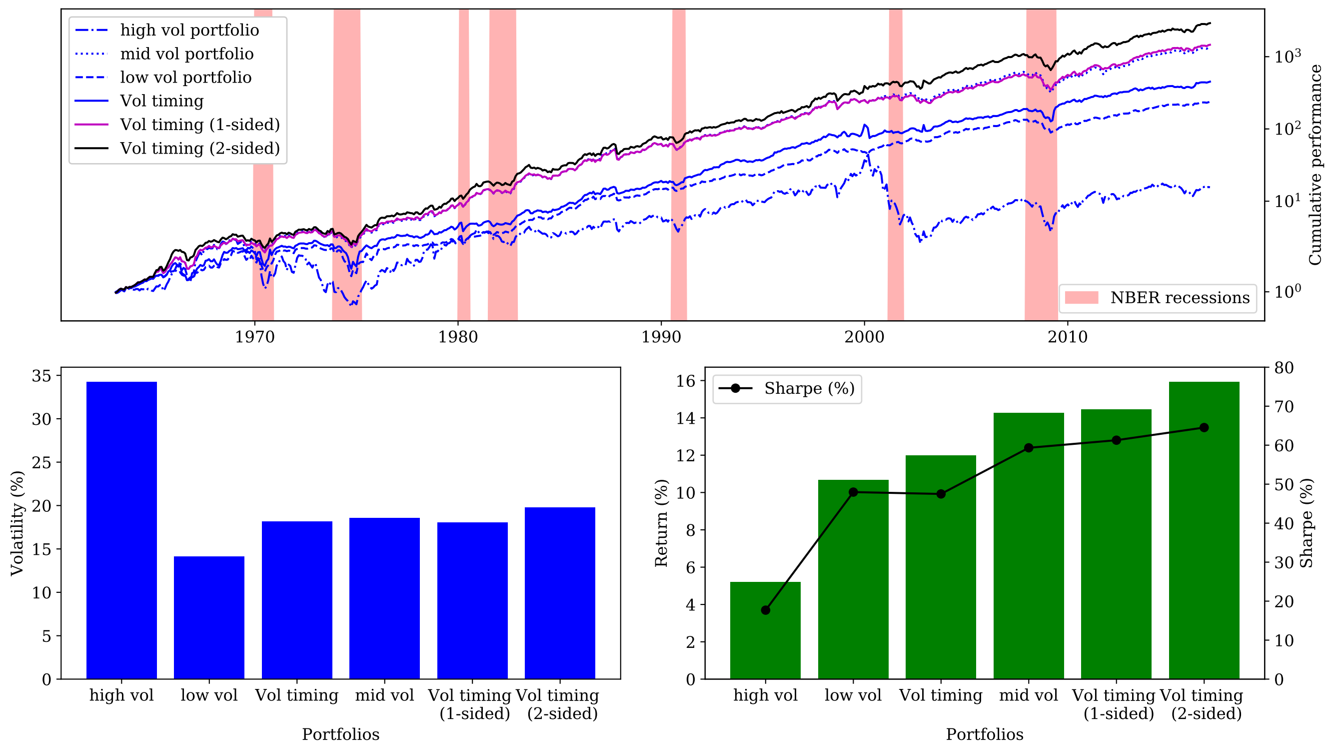

Interestingly, the mid-volatility portfolio beats both low-vol and the basic version of our vol timing strategy by a comfortable margin. In essence, the mid-volatility portfolio compensates for the slightly higher volatility by exhibiting noticeably higher return. This suggests that our volatility timing strategy can be further refined. Instead of holding the low-volatility portfolio as default, which assumes that the market condition is bad, we hold mid-volatility portfolio by default, thereby assuming that market condition is neutral. We consider two refinements on our volatility timing strategy. The first is "Vol timing (1-sided)". In this case, we hold the mid-volatility portfolio as default. If the slope of the return profile is significantly negative, we have sufficient evidence to claim that the market condition is bad, and we switch to hold the low-volatility portfolio. As the name implies, in this 1-sided test we only assess how holding low-volatility stocks can help during bad market condition, while riding on the higher return of the mid-volatility portfolio during normal market condition. We also consider a further refinement which we call "Vol timing (2-sided)". As the name suggests, this includes all the feature of the 1-sided strategy, but also test the slope for good market condition—if the return slope profile is sufficiently positive, we accept the alternative hypothesis that the market condition is good, and we in turn switch to hold the high-volatility portfolio to benefit from strong market growth. As before, the test is run across each month over the full period studied, and we decide whether to hold mid-volatility portfolio (default), or low- or high-volatility portfolio if we reject our null hypothesis depending on whether market condition is bad or good, respectively.

On average, the three vol timing strategies investigated achieved an accuracy of 57%, where we measure accuracy as holding the right portfolio at the right time. The default portfolios are held on average 81% of the periods, while the alternative portfolios are held on 19% of the periods. The results of the two refined volatility timing strategies are labeled as "Vol timing (1-sided)" and "Vol timing (2-sided)" in Exhibit V. We are able to obtain superior performance. Compared to the low-volatility portfolio, the vol timing (2-sided) strategy achieve a Sharpe ratio improvement from 0.48 to 0.65, and a cumulated wealth improvement in excess of 1125%. This strong performance can be intuitively understood by referring to both Exhibit I and III. Both vol timing (1-sided and 2-sided) strategies ride on the higher return of the mid-volatility portfolio under normal market condition. Only when there is statistically significant evidence based on the slope parameter to reject the null and accept the alternative of bad market condition will we switch to hold the low-volatility portfolio to further optimize risk-reward exposure. For vol timing (2-sided), we also test for good market condition and selectively hold high-volatility stocks during strong growth periods. A summary of our main results can be found in Exhibit VI.

Exhibit VI: Portfolio statistics of market-timed portfolios

	Volatility (%)	Sharpe (%)	Wealth ($)	Default Portfolio	Alternative Portfolio
Vol timing (2-sided)	19.8	64.5	2913.7	mid-vol	low-/high-vol
Vol timing (1-sided)	18.1	61.3	1461.7	mid-vol	low-vol
Mid vol	18.6	59.3	1345.1	-	-
Vol timing	18.2	47.5	451.3	low-vol	high-vol
Low vol	14.1	47.9	237.8	-	-
High vol	34.2	17.7	15.6	-	-

For the three volatility timing strategies, the default and alternative portfolios are also labeled in Exhibit VI. All 6 strategies are ranked according to their cumulated wealth. As expected, the high-volatility portfolio has the worst performance. Low-volatility portfolio performs comparatively better with less than half the realized volatility and approximately three times the return. The basic vol timing strategy which holds low-volatility portfolio by default and occasionally holds high-volatility stocks during periods of strong growth exhibits higher cumulated wealth at 451.3, but holding higher volatility stocks inevitably drags down the Sharpe ratio. It is noteworthy that the mid-volatility portfolio performs substantially better than the basic vol timing strategy, with a cumulated wealth of 1345.1. Based on this observation, we refine our vol timing strategies by using the mid-volatility portfolio as default to ride on its higher risk-adjusted return. The vol timing (1-sided) strategy use statistical test to switch to low-volatility portfolio during market downturns, and achieve a marginal reduction in realized volatility from 18.6% to 18.1%, with corresponding improvements in Sharpe ratio and cumulated wealth. The best performing strategy is vol timing (2-sided), with a cumulated wealth of 2913.7, more than double the nearest competitor—by also testing for good market condition and holding high-volatility stocks occasionally, a Sharpe ratio improvement to 0.65 can be achieved.

Robustness Test & Discussions

We compute CAPM and Fama-French alphas of volatility decile portfolio to check the robustness of our results. For CAPM, we run the following regression:

$$r_{t,j} = \alpha_j + \beta_{j,\text{MKT}} r_{t,\text{MKT}} + \epsilon_{t,j}, \qquad j = 1, \cdots, 10$$

where $r_{t,j}$ is the j^{th} volatility decile portfolio return at time t, $r_{t,\text{MKT}}$ is the excess return on the market portfolio, and α_j is the CAPM alpha. For Fama-French, we run the following regression:

$$r_{t,j} = \alpha_j + \beta_{j,\text{MKT}} r_{t,\text{MKT}} + \beta_{j,\text{SMB}} r_{t,\text{SMB}} + \beta_{j,\text{hml}} r_{t,\text{HML}} + \epsilon_{t,j}, \qquad j = 1, \cdots, 10$$

where $r_{t,\text{MKT}}$ is the excess return on the market portfolio, $r_{t,\text{SMB}}$ is the daily return on the small-minus-big (SMB) factor, and $r_{t,\text{HML}}$ is the daily return on the high-minus-low (HML) factor, and α_j is the Fama-French alpha.

Exhibit VII reports the alphas, betas and adjusted-R^2 of the CAPM and Fama-French regressions for the low-, mid- and high-volatility portfolios and the volatility timing strategies (1-sided and 2-sided). The regressions are carried out using the portfolios' realized monthly returns. The Newey and West robust t-statistics are also provided in parentheses below. ** and * indicate significance at the 1% and 5% levels respectively. Note that both CAPM and Fama-French alphas increase progressively over the strategies ranked by performance. This shows that the returns of the volatility timing strategies are not explained by the well-known market, size, and book-to-market factors.

Exhibit VII: CAPM & Fama-French Alphas

Portfolio	Panel A: CAPM			Panel B: Fama-French				
	$\alpha(\%)$	β_{mkt}	$R^2(\%)$	$\alpha(\%)$	β_{mkt}	β_{SMB}	β_{HML}	$R^2(\%)$
High vol	-0.442** (-2.449)	1.664** (26.613)	72.721	-0.374** (-2.708)	1.434** (23.100)	0.843** (9.428)	-0.351** (-3.786)	83.021
Low vol	0.203** (2.541)	0.596** (23.722)	66.349	0.069 (0.912)	0.645** (26.423)	-0.003 (-0.086)	0.290** (6.871)	72.251
Mid vol	0.326** (5.107)	0.995** (46.209)	89.262	0.204** (3.550)	1.012** (55.479)	0.100** (2.564)	0.237** (5.644)	91.313
Vol timing	0.288* (2.402)	0.720* (16.237)	49.702	0.155 (1.309)	0.700* (14.078)	0.273* (3.093)	0.218* (2.893)	53.821
Vol-timing 1-sided	0.350** (5.525)	0.958** (42.107)	87.994	0.241** (3.895)	0.975** (44.140)	0.089** (2.231)	0.213** (5.078)	89.747
Vol-timing 2-sided	0.453** (5.678)	0.998** (38.295)	81.708	0.360** (4.580)	0.993** (40.537)	0.154** (2.819)	0.163** (3.137)	83.109

This aside, we also perform correlation analysis to investigate the reason behind the outperformance, which is presented in Exhibit VIII. Portfolio correlation is calculated using time series data, i.e. the realized returns. Highlighted regions denote NBER-defined recession periods[4] . The general portfolio correlation is defined as

$$\bar{\rho} = \frac{2\sum_{i=1}^{N}\sum_{j>i}^{N} w_i w_j \rho_{ij}}{1 - \sum_{i=1}^{N} w_i^2}$$

where ρ_{ij} is the pair-wise correlation between stock i and stock j, while w_i and w_j are their weights in the portfolio, respectively.

We see that the moving window average correlation is substantially higher in high-volatility portfolio. It is also noticeable that correlation spikes up during market recessions. Although not reported here, we have also calculated portfolio correlation using a factor model based framework, and obtain similar results in Exhibit VIII.

The analysis presented here lends support to our assertion earlier that the outperformance achieved by our volatility timing strategy, in particular the vol timing (2-sided) version, can be attributed to successfully holding a concentrated growth portfolio during bull markets, while switching to hold a diversified low-volatility portfolio during bear markets.

Exhibit VIII: Correlation of low- and high-volatility portfolios.

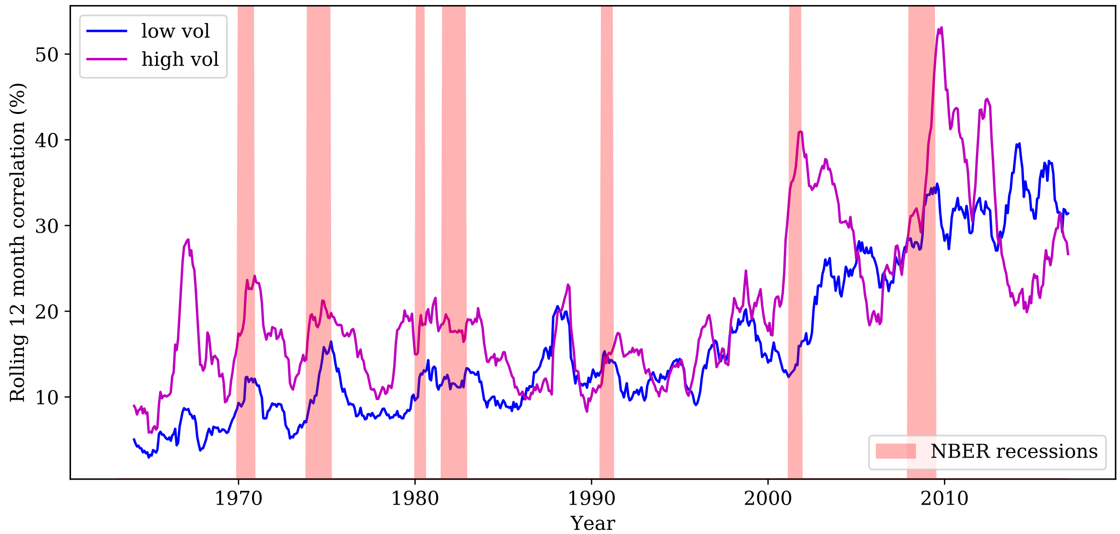

Conclusions

We add to the low-volatility literature by developing a volatility timing framework that can be used to further enhance the performance of low-volatility strategy. We show that the shape of the volatility decile portfolio's return profile is time-varying, and contains important information about market condition that can be harvested to form volatility timing strategy. When market condition is good, high-volatility portfolio outperforms, while during bad market condition, low-volatility portfolio is optimal. Our volatility timing strategy monitors the slope of the return profile for signals to switch between holding low-, mid-, or high-volatility portfolio, with the objective of holding concentrated growth stocks during good market condition, and holding diversified low-volatility stocks during bad market condition. A refined version of our volatility timing strategy based on a two-sided test that holds mid-volatility portfolio by default and selectively switch to low- or high-volatility portfolios during bad or good market condition is able to generate superior performance relative to other comparable strategies.

Portfolio correlation analyses show that high-volatility portfolio are more correlated compared to low-volatility portfolio. Our results thus connect research on low-volatility strategies with studies done on portfolio correlation and diversification, highlighting a feasible way to further improve portfolio risk-return performance.

An important follow-up work is to investigate the transaction costs of incorporating volatility timing to low-volatility strategies. van Vliet (2018) has demonstrated that a turnover level of around 30% is enough to create an effective low-volatility strategy. An in-depth comparison on the profitability of the volatility timing strategy, along with other related strategies like momentum, in the presence of transaction cost will be of interest to investors. Nevertheless, our analyses reported in this paper already reveal important insights on the variation and behavior of the volatility decile portfolio's return profile, and the feasibility of volatility timing using this information.

Notes

[1]Fu (2009) uses an EGARCH model to estimate expected idiosyncratic volatility and sort stocks into decile portfolios. His results show no evidence of a low volatility anomaly. Quite the opposite, the high volatility portfolio clearly outperforms the low volatility portfolio.

[2]This is different from risk parity approach, which ignores information about expected returns and covariances. Volatility-managed approach only increase or decrease overall risk exposure based on total volatility.

[3]Static here refers to holding portfolio corresponding to a specific portfolio volatility characteristic, e.g. low-volatility. It will still involve rebalancing as the low volatility stocks do change over time based on their realized volatility.

[4]NBER recession periods are obtained from `https://www.nber.org/cycles.html`.

References

Ang, A., Hodrick, R. J., Xing, Y. and Zhang, X. (2006), 'The cross-section of volatility and expected returns', *Journal of Finance* **61** (1), pp. 259–299.

Ang, A., Hodrick, R. J., Xing, Y. and Zhang, X. (2009), 'High idiosyncratic volatility and low returns: international and further U.S. evidence', *Journal of Financial Economics* **91** (1), pp. 1–23.

Baker, M., Bradley, B. and Wurgler, J. (2011), 'Benchmarks as limits to arbitrage - understanding the low-volatility anomaly', *Financial Analysts Journal* **67** (1), pp. 40–54.

Baker, N. L. and Haugen, R. A. (2012), 'Low risk stocks outperform within all observable markets of the world', *Working Paper*, [Online] Available at SSRN: `https://ssrn.com/abstract=2055431`.

Bessembinder, H. (2018), 'Do stocks outperform Treasury bills?', *Journal of Financial Economics* **129** (3), pp. 440–457.

Beveratos, A., Bouchaud, J.-P., Ciliberti, S., Laloux, L., Lempérière, Y., Potters, M. and Simon, G. (2017), 'Deconstructing the low-vol anomaly', *Journal of Portfolio Management* **44** (1), pp. 91–103.

Black, F., Jensen, M. C. and Scholes, M. S. (1972), 'The capital asset pricing model - some empirical tests', *Studies in the Theory of Capital Markets*, ed. by M. C. Jensen, New York: Praeger.

Blitz, D. C. and van Vliet, P. (2007), 'The volatility effect', *Journal of Portfolio Management* **34** (1), pp. 102–113.

Choueifaty, Y. and Coignard, Y. (2008), 'Toward maximum diversification', *Journal of Portfolio Management* **35** (1), pp. 40–51.

Chow, T.-M., Hsu, J. C. and Li, F. (2014), 'A study of low-volatility portfolio construction methods', *Journal of Portfolio Management* **40** (4), pp. 89–105.

Chua, D. B., Kritzman, M. and Page, S. (2009), 'The myth of diversification', *Journal of Portfolio Management* **36** (1), pp. 26–35.

Clarke, R. G., de Silva, H. and Thorley, S. (2006), 'Minimum-variance portfolios in the US equity market', *Journal of Portfolio Management* **33** (1), pp. 10–24.

Cox, J. C. and Ross, S. A. (1976), ‘The valuation of options for alternative stochastic processes’, *Journal of Financial Economics* **3** (1-2), pp. 145–166.

Driessen, J., Kuiper, I., Nazliben, K. and Beilo, R. (2019), ‘Does interest rate exposure explain the low-volatility anomaly?’, *Working Paper*, [Online] Available at SSRN: `https://ssrn.com/abstract=2831157`.

Fleming, J., Kirby, C. and Ostdiek, B. (2001), ‘The economic value of volatility timing’, *Journal of Finance* **56** (1), pp. 329–352.

Frazzini, A. and Pedersen, L. H. (2014), ‘Betting against beta’, *Journal of Financial Economics* **111** (1), pp. 1–25.

Fu, F. (2009), ‘Idiosyncratic risk and the cross-section of expected stock returns’, *Journal of Financial Economics* **91** (1), pp. 24–37.

Haugen, R. A. and Heins, A. J. (1975), ‘On the evidence supporting risk premiums in the capital market’, *Wisconsin Working Paper 4-75-20*, [Online] Available at SSRN: `https://papers.ssrn.com/sol3/papers.cfm?abstract_id=1783797`.

Li, Q., Yang, J., Hsiao, C. and Chang, Y.-J. (2005), ‘The relationship between stock returns and volatility in international stock markets’, *Journal of Empirical Finance* **12** (5), pp. 650–665.

Li, X., Sullivan, R. N. and Garcia-Feijóo, L. (2016), ‘The low-volatility anomaly: market evidence on systematic risk vs. mispricing’, *Financial Analysts Journal* **72** (1), pp. 36–47.

Moreira, A. and Muir, T. (2017), ‘Volatility-managed portfolio’, *Journal of Finance* **72** (4), pp. 1611–1644.

Van Vliet, P. (2018), ‘Low volatility needs little trading’, *Journal of Portfolio Management* **44** (3), pp. 33–42.