



Volatility Timing - Its Time Has Come

June 2016

Dr. Vinay Nair, 55ip¹

¹ Dr. Nair is chairman of 55ip and a visiting professor at The Wharton School.

CONTENTS

Introduction	3
Volatility Timing and Volatility Managed Portfolios (VMPs).....	3
Forecasting Volatility is Possible and Practical	4
Predicting Up or Down Markets: Volatility Timing Makes a Form of Market Timing Feasible	7
A Simple Example of Volatility Timing and Performance	8
Transaction Costs.....	9
Volatility Timing and Portfolio Holdings	10
Conclusion.....	11

Introduction

Market timing is generally known as the act of attempting to predict the future direction of the market. Attempts at market timing reflect the general desire to improve portfolio performance over a “buy-and-hold” portfolio. Often investors explicitly indulge in this temptation and at times they get swayed by the experts, pundits and media spurring one-off actions to buy “this” or sell “that”. Even when they are not actively paying attention to their portfolios, the decision of when to cash out or to start an investment often leads them down the path of market timing.

Yet, there is little evidence that market returns are predictable at short time horizons. In fact, there is a strong view among researchers that it is not possible to predict market returns at short time horizons on a sustainable basis. Both theory and empirical findings suggest that short-term market return forecasting is difficult, and perhaps even impossible. Theory suggests that asset returns should not be easily predicted using readily-available information and forecasting techniques, and a broad interpretation of four decades of empirical work suggests that the data support the theory (e.g., Fama, 1970, 1991).²

How then might it be possible to improve portfolio performance over a buy-and-hold portfolio?

This paper highlights an alternative approach by attempting to predict not the magnitude of future market returns, which is difficult if not impossible, but the volatility of future market returns, which are more quickly visible. Focusing on the volatility or risk dimension may enable performance that is superior over a buy-and-hold portfolio.

Volatility Timing and Volatility Managed Portfolios (VMPs)

Is forecasting market volatility a form of timing? In a sense, yes, but it is a very different sort of timing. Let’s call it “Volatility Timing” or a term we prefer, “Volatility Managed Portfolios” (VMPs). VMPs hold promise for investors who are looking for better prescriptions than the traditional buy-and-hold (and pray) approach. Investors’ need for new investment solutions is particularly important now. Expected returns are at multi-generational lows across asset classes. For example, a measure of forward-looking real equity yield and real bond yield are at the bottom 10th percentile today when compared to the level of these measures since 1900.³

Even as these future returns have fallen, liabilities and spending needs of investors have been slow to adjust. For example, even as forward-looking real yields on bonds – a big fraction of typical portfolios – have fallen from approximately 4% to 0% in the last 15 years, the average public pension fund return target has dropped from 8.1% in 2001 to 7.68% today.⁴ Among the tools investors need to bridge the earnings and spending gap is VMPs. Volatility timing and VMPs, even while harboring skepticism around traditional market timing, has several theoretical and empirical underpinnings.

² We do, however note that return forecasts are possible at very short horizons (see for example Lo and MacKinlay, 1999), typically emerging as weak serial correlation in ultra-high frequency returns due to microstructure effects. Additionally, many believe that returns are also forecastable at very long horizons (see for example Fama and French, 1988, 1989, and Campbell and Shiller, 1988), perhaps due to time-varying risk premia. However, the possible presence of very short-run or very long-run mean reversion is of little relevance to this paper, because we focus on return horizons of one month. We do believe that long-run time varying risk premia may be forecastable and leave that discussion for another paper.

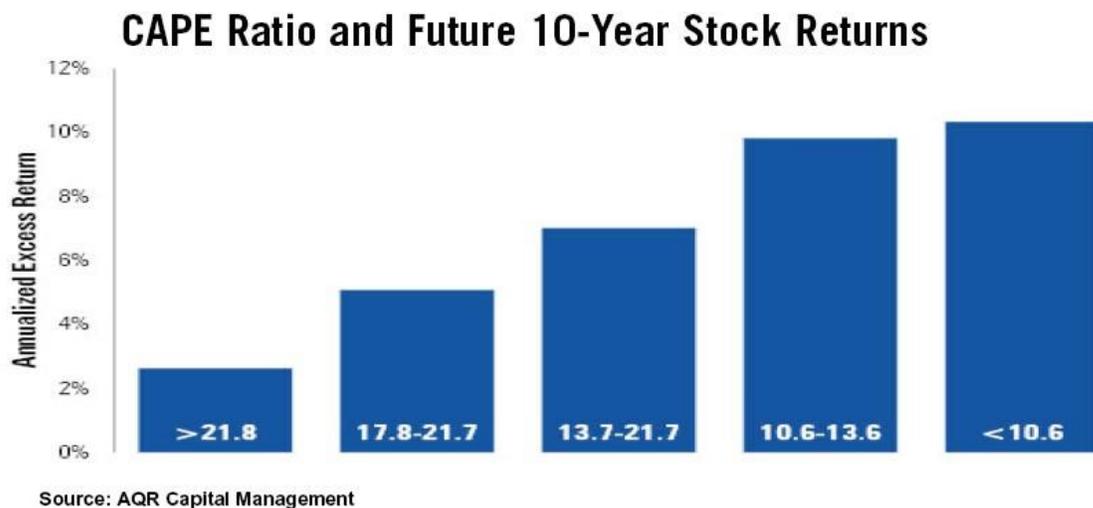
³ Real equity yield is measured by an average of the Shiller earnings yield (using 10-year earnings) and dividend yield plus the long-run real growth (1.5%). Real bond yield is the yield on long-term US Treasury bonds minus the long-term expected inflation. Returns associated with other source of returns such as factors are also likely to be low.

Forecasting Volatility is Possible and Practical

It may seem a contradiction, but it is not. Even as a significant amount of research continues to document the difficulty (if not impossibility) in forecasting market returns, a large body of influential research documents the notable predictability of asset return volatility, which provides important implications for asset allocation. The evidence that volatility is forecastable also comes from different research traditions that include GARCH, stochastic volatility models, realized volatility models, and regime switching models.⁵

Prediction, not correlation, is critical to the implementation of any form of timing – market return or volatility. This may sound obvious, but the implications are often ignored when exhibits such as Exhibit 1 are presented.

Exhibit 1



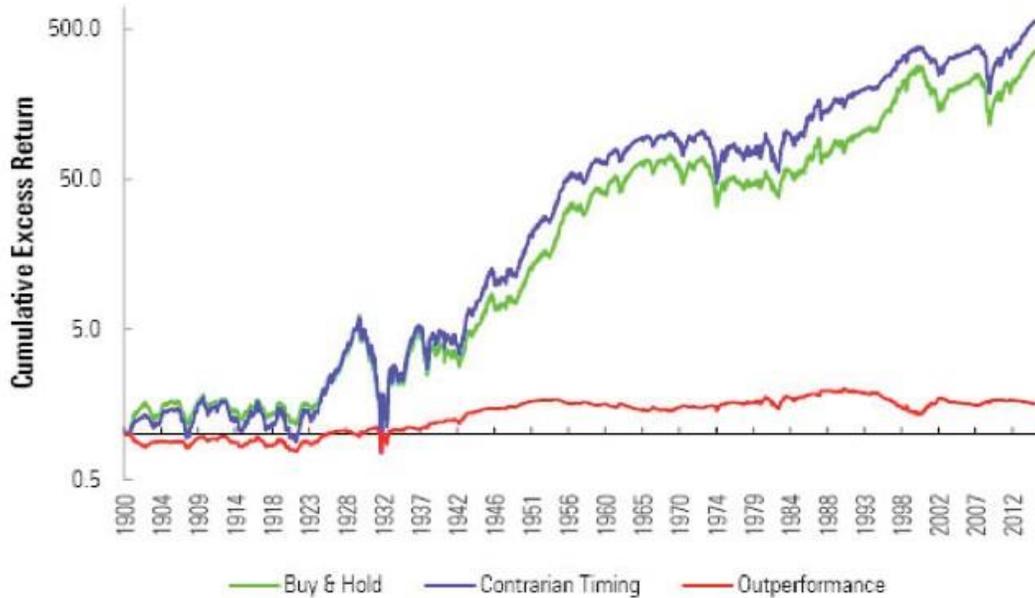
This exhibit suggests that market timing is possible by using the Shiller CAPE (Cyclically adjusted price-to- earnings) ratio. While there are many such graphs and studies that show a correlation between some market timing indicator and future returns, these graphs do not translate into an implementable trading strategy. Asness, Ilmannen and Maloney (Institutional Investor, November 2015) show that a trading strategy using the information in Exhibit 1 translates into a lackluster portfolio – as shown below – and that too without transaction costs.

⁴ See, for example, <http://www.wsj.com/articles/taxpayers-more-pension-burdens-headed-your-way-1441388090>.

⁵ Bollerslev, Chou and Kroner (1992) provide a fine review of evidence in the GARCH tradition, while Ghysels, Harvey and Renault (1996) survey results from stochastic volatility modeling, Andersen, Bollerslev and Diebold (2003) survey results from realized volatility modeling, and Franses and van Dijk (2000) survey results from regime-switching volatility models. The recent literature also contains intriguing theoretical work explaining the empirical phenomena, such as Brock and Hommes (1997).

Exhibit 2

The CAPE Ratio Timing Signal, Put Into Practice



Source: AQR Capital Management

Differences between Exhibit 1 (typical in academic research findings) and Exhibit 2 (practical impact of incorporating a market timing rule) are more profound than commonly realized.

Why is this the case? Statisticians attribute the problems to “estimation risk” and “parameter learning”.

Estimation risk is the uncertainty that is inherently present when estimating any parameter. After all, estimates are drawn from a (hidden) distribution to hopefully reveal the true parameter and there is always some error, or difference, between the true parameter and the inferred estimate. For example, there is, in fact, an average size of all the fish in the ocean, but no matter how many fish we measure, we are never sure how close we are to the actual measurement.

Several aspects govern the magnitude of this error. If the data is limited, the estimation error is larger. We can make statements about whether the error is larger or smaller in a given context, but we cannot know the precise direction and magnitude of the error.⁶

Parameter learning is the dynamic counterpart to estimation risk. In addition to problems identifying the true parameter, the true parameter is also changing over time. Perhaps to stretch the fish analogy, fish are always changing in size due to changes in the ocean. This leads to a different source of error in our sizing estimates. The process of revising beliefs about parameters as more information arrives is the source of uncertainty and parameter learning.

⁶ If we could, we would simply adjust the estimate for the error to get the true parameter!

To provide another analogy, estimation risk relates to errors a marksman makes while aiming for the target while parameter learning relates to errors a marksman makes in identifying the target itself. After all, when the target is moving, the marksman does not know precisely its location.

Since predicting timing indicators, before using them, is the first step in any type of timing, these concepts have significant practical impact on implementing any research on timing. Several papers have highlighted that ignoring estimation risk or parameter learning typically leads to misleading allocations [Brennan (1998), Stambaugh (1999) and Barberis (2000)].

The difference between Exhibits 1 and 2 is the practical representation of this problem.

Much of the criticism around colloquial market timing applies to issues of predicting returns in the presence of estimation risk and parameter learning. Even when theory may suggest that investors should time the market based on expected returns, as it does for long-run returns, in practice this is difficult due to the risk present in estimating the parameters that drive the expected market returns process. Given all past data, we are never certain enough about what will happen next with returns for this to work reliably.

But in the case of volatility timing, the math of estimation risk makes it easier and quicker to see patterns in volatility relative to market returns. It is nothing more than arithmetic at work. Standard errors, a measure of the statistical accuracy of an estimate, are higher for means (returns) than for standard deviations (risk).

The standard error when estimating returns is σ/\sqrt{n} and the standard error of volatility is $\sigma/\sqrt{2(n-1)}$, where n is the number of observations. As a result, the accuracy of volatility estimates is much higher than return estimates.

For example, with 120 observations – 10 years of monthly data – the estimate for volatility is about 40% more accurate than the estimate for returns. Put another way, roughly speaking, it takes a year of returns to learn as much as we do from about two months of volatility calculations.

This improved precision in estimating volatility is further enhanced by increased sampling. For example, instead of n years of monthly data, one could use $\sim 20n$ daily observations to reduce the standard error of the volatility estimate to $\sigma/\sqrt{(40n)}$. Interestingly, increasing the sample size of returns data does not improve the returns estimate. The standard error remains the same as before.⁷ As a result, the estimate for volatility can be almost 6 times better than the estimate for the larger sample of returns!

⁷The standard error for the mean is 20 times standard error of the daily mean which is daily standard deviation/ $\sqrt{(20n)}$). The daily standard deviation is the $\sigma/\sqrt{20}$. As a result, the standard error for returns remains as σ/\sqrt{n} .

Predicting Up or Down Markets: Volatility Timing Makes a Form of Market Timing Feasible

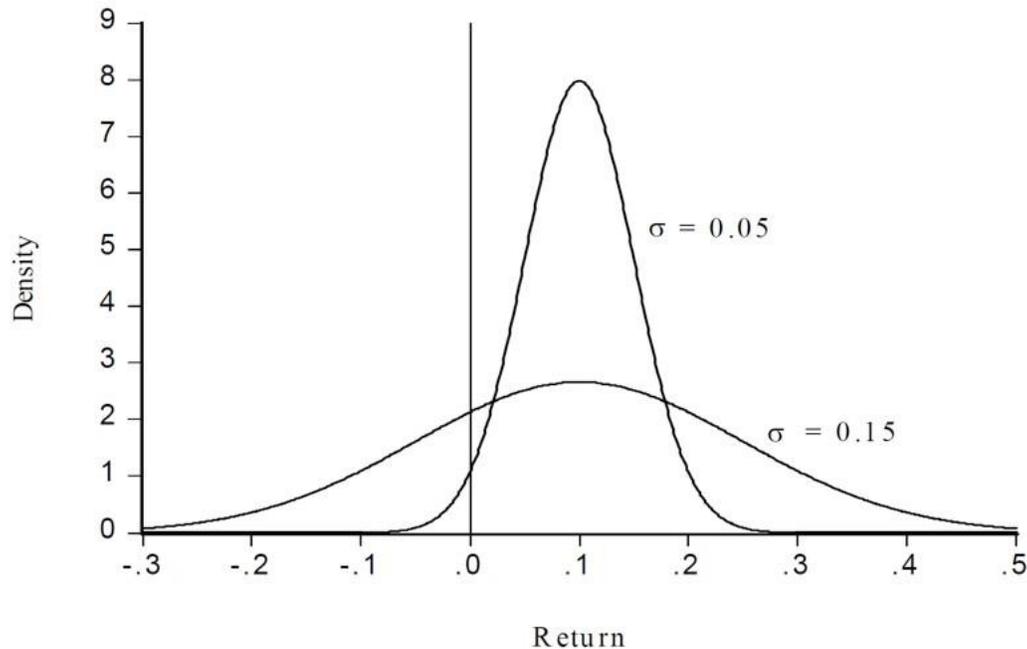
An intriguing implication of forecastable volatility is that the future direction of the market, not the magnitude of the move, is also forecastable. As a result, a specific form of market timing is the natural consequence of volatility timing.⁸ Presented differently, even if future average returns are not predictable, the direction of the future return is predictable if volatility is predictable. To see this, note that the probability of the return being up in the next period is given by

$$\Pr_t(R_{t+1} > 0) = 1 - \Pr_t(R_{t+1} < 0) = 1 - \Pr\left(\frac{R_{t+1} - \mu}{\sigma_{t+1|t}} < \frac{-\mu}{\sigma_{t+1|t}}\right) = 1 - F\left(\frac{-\mu}{\sigma_{t+1|t}}\right),$$

where $F(\cdot)$ is the relevant cumulative distribution function. Hence the market direction is forecastable because the probability of a positive return (> 0.5 in a normal context if mean (μ) is greater than 0) is time varying. So as long as the mean is non-zero and volatility is forecastable, the future direction of the market is forecastable even if the future return (μ) is not. Once again, estimates of volatility are far more, and far more quickly, useful to an investment strategy.

Another way to think about it is shown on Exhibit 3. When volatility is very low (as it is for sovereign bond payments), the returns are likely to be clustered around the positive expected return and most realizations are positive. On the other hand, when the volatility is very high, there are more occurrences of negative returns, but also of very positive returns as well.

⁸ Often results attributed to volatility predictions are interpreted as market timing. A true test of market timing would be to predict future returns without predicting differences in volatility. This is challenging. Goyal and Welch (2008) showed that simple predictability models with constant volatility do not lead to statistically significant out-of-sample portfolio gains.

Exhibit 3

Notes to figure: We show two Gaussian return densities, each with expected return of ten percent. The first return has a standard deviation of five percent and hence is positive with probability 0.98 (the area to the right of zero under the more peaked density function). The second return has a standard deviation of fifteen percent and hence is positive with smaller probability 0.75 (the area to the right of zero under the less peaked density function).

Source: “Financial Asset Returns, Direction-of-Change Forecasting, and Volatility Dynamics,” Peter F. Christoffersen and Francis X. Diebold, National Bureau of Economic Research (NBER) Working Paper No. 10009, October 2003.

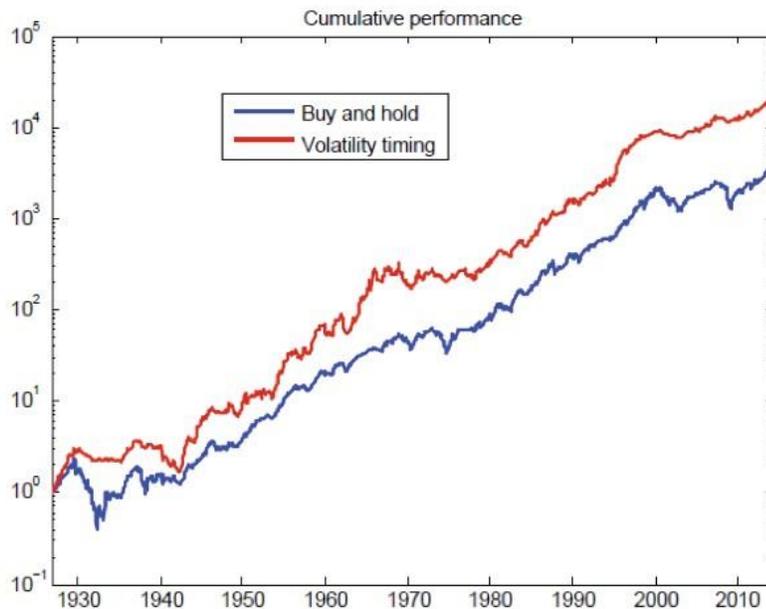
This rather interesting result has been empirically validated by several papers starting from Breen, Glosten and Jagannathan (1989). Breen et. al. show that treasury bill returns can predict the volatility and direction of next month’s stock market index returns in economically meaningful magnitudes. Bollerslev and Christofferson (2010) have a more comprehensive investigation of this phenomena that highlights its robustness and theoretical underpinnings.

A Simple Example of Volatility Timing and Performance

A recent paper by Tyler Muir and Alan Moreira (Muir and Moreira, 2015) highlights the power of VMPs using a very simple predictor of volatility. The predictor they focus on is simply the realized volatility of the last month. Since our objective in this paper is to highlight the potential of VMPs, we focus on the usefulness of this simple volatility forecast rather than contrast different volatility forecasts. Indeed, there are better and more powerful forecasts of volatility available to savvy investors and these more powerful forecasts only strengthen the case for volatility timing and VMPs.

To see the implications of volatility timing for an investor, Muir and Moreira compare a buy-and-hold portfolio with a portfolio that allocates capital to the market based on the variance of the market portfolio in the immediately preceding month. The dynamic portfolio adjusts so that, if the variance of the daily returns in the past month is higher, then the allocation to the market in the next month is reduced. As the (log scale) graph from Muir and Moreira shows in Exhibit 4, this dynamic portfolio outperforms a buy-and-hold portfolio in a rather consistent manner from 1926 to 2015 with the same monthly standard deviation.

Exhibit 4



Source: Muir and Moreira

Transaction Costs

Transactions costs make many an interesting result impractical. Therefore, it is useful to check whether the above result is robust when taking into account practical considerations such as transaction costs. Muir and Moreira show that they are. Specifically, a version of the strategy that trades only when last month volatility is higher than the historical average generates an average monthly turnover of 10% and an annual outperformance of 2.2% before trading costs. For this outperformance to vanish, the transaction costs will have to be 1.83%. To put this in context, current reasonable estimates of transaction costs are 10-15bps of traded value in retail platforms and lower in institutional settings.

Volatility Timing and Portfolio Holdings

To appreciate the implications of this approach, it is helpful to hone in on the portfolio dynamics. When forecasted volatility goes up, this approach decreases the exposure to the market. For example, this approach would have decreased an investor's exposure to the market after the significant rise in volatility in October 2008 (annualized market volatility had risen to approximately 60%). This may seem counterintuitive especially since this increase was accompanied by a significant drop in returns. After all, these are times when we think future expected returns are rising (note how market timing is creeping in!). In fact, the common wisdom among those with a value mindset was not to sell or panic, but rather buy.

Muir and Moreira show in their research that the horizon of these two variables – volatility and future return – behave quite differently. Volatility tends to spike and recede quickly, whereas expected returns are more persistent and realized in a slower manner with a longer horizon. An investor may miss the first leg of the future returns while volatility decreases/normalizes, but returns tend to remain persistent allowing the investor a chance to participate even as they enter with a delay. On a risk-adjusted basis, this reaction generates superior performance and better downside protection.

Understanding this counterintuitive implication is also related to the research on the link between risk and return. The empirical results of a risk-return tradeoff are surprisingly mixed. Various studies have found little to no relation between returns and risk in regressions across different sample periods, specifications and horizons. If this is indeed the case, there is no conflict in decreasing allocations when risk is higher. Even if the link between risk and return is positive, volatility timing can still be beneficial if expected returns do not rise by enough compared to increases in volatility.

A more complete discussion of this topic would also include an appreciation of the type of risk that generates an increase in volatility and associated views about future market returns.⁹ Holding on or buying more, for example, would not be detrimental if the source of the volatility shock is believed to be temporary and will revert soon, implying a view about future market returns (market timing and volatility timing will interact if one entertains both possibilities).

⁹For example, market volatility comes from multiple sources of shocks, and these shocks may or may not be compensated by the market.

Conclusion

This paper highlights the promise of using volatility forecasts to improve the performance of buy-and-hold portfolios. Even when the next period's returns are not predictable, it is possible to forecast the volatility of the market. This in turn also generates the ability to forecast up or down markets.

Using volatility forecasts to generate volatility managed portfolios can improve the Sharpe ratios of buy-and-hold portfolios through a meaningful reduction in downside risk. We believe this approach holds significant promise to investors who are faced with low expected returns and high uncertainty in many asset classes such as bonds and equities.

In combination with the findings about long-term predictability of risk premia, this could be a valuable input to constructing portfolios that are superior to buy-and-hold portfolios. Further enhancements are also possible by distinguishing the nature of the volatility shocks between temporary and permanent shocks, as well as rewarded and unrewarded constituents of risk. At 55ip, we use these insights to generate portfolios with downside protection across equities, asset allocation and macro strategies.



55ip
617-217-2376
www.55-ip.com

BOSTON OFFICE
One Liberty Square, Suite 200,
Boston, MA 02109

NEW YORK OFFICE
12 East 49th Street, 11th Floor New
York, NY 10017

MUMBAI OFFICE
Aston Building, 10th Floor
Shastri Nagar, Andheri (W),
Mumbai 400 102, India

Opinions expressed are current opinions as of the date appearing in this material only. While the data contained herein has been prepared from information that 55ip believes to be reliable, 55ip does not warrant the accuracy or completeness of such information. This communication is for informational purposes only. This is not intended as nor is it an offer, or solicitation of any offer to buy or sell any security, investment or product. This document may not be reproduced or transmitted in any form or by any means, electronic or mechanical, for any purpose, without our prior written permission.

Copyright © 2016, 55 Institutional Partners, LLC ("55ip"). All rights reserved. Advisory services provided by 55I, LLC. 55ip does not guarantee future results.