

Volatility as an Asset Class and Dynamic Asset Allocation

The proliferation of financial tools and instruments in the last decades has increased the flexibility of portfolio managers to shape the return profiles of their portfolios. In this paper we look at volatility as an investable asset class for global investors, explain what it is, describe its characteristics and its role in a global portfolio.

Multi-Asset Solutions Research Papers - Issue 1


For professional clients only



Volatility as a concept

The concept of volatility long predates its use as an asset class, and typically is used as a descriptive characteristic of the risk of an asset. For any financial instrument for which we have historical return data we can calculate the volatility of the returns. Mathematically this is the standard deviation of the return numbers, and the larger this number, the larger the swings in performance of the underlying asset.

Figure 1: Historical Indexed Performance



Source: First Sentier Investors

The concept of volatility is intuitive and obvious, but using it presents numerous pitfalls.

As an example we show the indexed performance of US equities, US bonds and USD cash¹ since February 1992. As is immediately obvious from the graph, the equity returns have seen much wider swings than either bonds or cash. Based on monthly data, we can now easily calculate the standard deviation of each of these time series, and after annualising them we see that the volatilities are 16.5% for US equities, 6.6% for US bonds and 0.6% for USD cash.

So over this period spanning almost 20 years we have observed that the volatility of bonds returns is 10 times that of cash, and equities have a volatility just over two and a half times that of bonds. This gives us a sense of the relative swings in the asset categories under consideration.

¹ Using the S&P 500 total return index, the JP Morgan US Government bond index and the USD one-month cash rate, all in USD.

Changes in Volatility over time

Volatility however is not constant over time. It varies in different periods as the economic cycle goes through ups and downs, and markets are perturbed by events. To show this we calculated the annualised volatility of the S&P 500 index using four different time windows: rolling over 12, 24 and 60 months, as well as an expanding window from 1980 until the last data point ending February 2012.

Figure 2: Annualised Volatility of the S&P 500 index



Source: First Sentier Investors

The chart above shows these four measurements of volatility, with the shorter time windows being obviously subject to the largest changes. We can see that the annualised volatility of the index varied between 8% and 25%, with some outliers outside of that range. It is also clear that volatility is a non-compounding asset as its realised value will always be range bound. In this respect it differs fundamentally from compounding assets such as equities and fixed income, and indeed in our Long-Term Asset Return Model our expected return for volatility is -0.25% .² Studies by Bakshi and Kapadia 2003 and Carr and Lee, Variance Risk Premiums 2009 suggest that the volatility risk premium is negative over longer periods of time, implying that the structural allocation to volatility, if any, should be a short position. However, as we shall show, a dynamic asset allocation approach can provide improvements in the risk/return profile of an overall portfolio.

² This is a zero expected return minus costs of carrying and trading the asset.

Implied volatility is information that is embedded in options prices.

Implied Volatility of Index Options

The level of either realised or expected volatility also plays a role in determining the price of options. There are fairly liquid markets in options on major equity indices, providing tools for investors to hedge their exposures in different ways. The Black-Scholes formula for valuing options contains volatility as a parameter, and as such it is possible to calculate an “implied volatility” of an option. Given the value of an index, the market price of an option and its characteristics (such as strike price, time to expiry, discount rates, etc.) it is possible to calculate the volatility required to end up with a “fair” price for the option in the Black-Scholes sense.

While it is possible to calculate the implied volatility of any instrument’s options, the most famous example of its expression has become the VIX index. This index tracks the implied volatility of options on the S&P 500 index and represents a market estimate of future S&P 500 volatility. This therefore can differ from the realised volatility; there are strategies that seek to exploit this difference, but delving into that would be beyond the scope of this paper; a good overview of various strategies can be found in Warren 2012.

Investing in Volatility

The VIX index however is not directly investible in the way that the S&P 500 is. To get S&P 500 exposure, we could simply buy the underlying 500 stocks in the right proportion, but it is impossible to buy the implied volatility directly. However just as there are futures on the S&P 500 index, there also exist futures on the VIX index, which are investible. Whaley 2009 provides an in-depth look at the construction of the VIX index and various investible instruments related to it.

Apart from the VIX index there are a few other similar indices such as the VDAX, VSMI and the VSTOXX, which measure the implied volatility in index options on the DAX, SMI and STOXX indices respectively. Their characteristics are broadly similar to the VIX and their traded futures based on these indices as well. In addition in early 2012 the VHSI was introduced, based on the Hang Seng Index in Hong Kong.

While VIX futures are our preferred instruments for investing in volatility, there are other ways of gaining exposure as well. The VXX ETF provides exposure to the VIX without having to purchase and roll futures, while there are many mutual funds that seek to optimise volatility exposure with active management, hoping to do better than with plain VIX futures. Finally customised swaps in the over-the-counter market can be structured with specific payoff patterns, and variance swaps are a commonly used tool in this space.

Correlation of Volatility with Equity Markets

As can be deduced from Figure 2 volatility typically goes up in times of stress and falling markets as witnessed by the spikes following the 1987 crash and the 2008 crisis. This characteristic is what makes it such an interesting hedging tool.

Neither volatility nor correlation is static, and that presents

Table 1: Historical Asset Class Characteristics

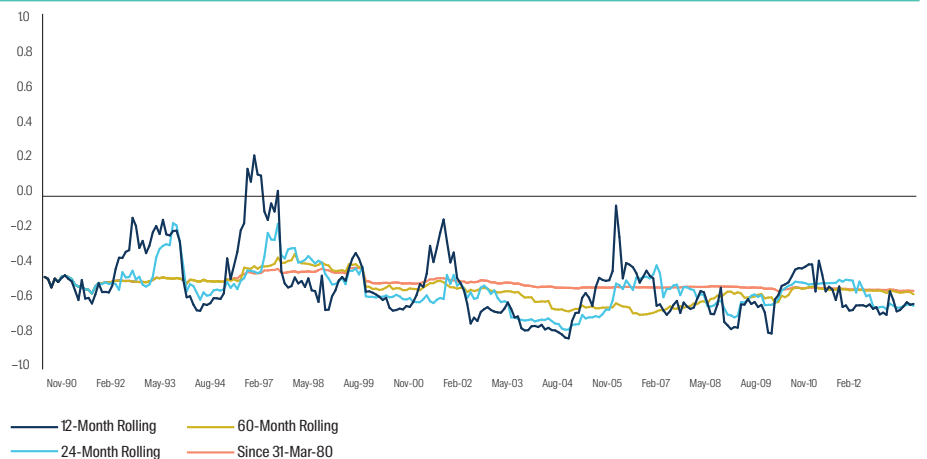
Asset Classes	Historical Return	Historical Volatility	Correlations						
			USD Cash	US Bonds	S&P 500	World Equities	Emerging Markets Equities	VIX Volatility	VDAX Volatility
USD Cash	3.5%	0.6%	1.00	0.07	0.03	-0.02	-0.12	0.08	0.04
US Bonds	6.6%	4.9%	0.07	1.00	-0.17	-0.18	-0.23	0.13	0.25
S&P 500	9.6%	16.5%	0.03	-0.17	1.00	0.93	0.72	-0.57	-0.52
World Equities	8.2%	16.5%	-0.02	-0.18	0.93	1.00	0.78	-0.56	-0.52
Emerging Markets Equities	12.0%	27.0%	-0.12	-0.23	0.72	0.78	1.00	-0.49	-0.47
VIX Volatility	20.1%	82.6%	0.08	0.13	-0.57	-0.56	-0.49	1.00	0.72
VDAX Volatility	24.2%	90.4%	0.04	0.25	-0.52	-0.52	-0.47	0.72	1.00

Source: First Sentier Investors

Table 1 above³ shows that the correlations between the S&P 500's volatility and various equity markets have been negative, which corroborates the hedging behaviour that we alluded to above. In order to get a better sense of what adding volatility does to a portfolio's performance characteristics, we look at just the simple case of adding 10% VIX exposure to either the S&P 500, the MSCI World or MSCI Emerging Markets indices. There is no corresponding volatility index or derivatives market for the world or emerging markets equity indices, but as we shall see, the characteristics of the VIX as a proxy provide acceptable hedging.

Just as volatility changes over time, so does correlation. Depending on the amount of stress and fear in the markets, the implied volatility will change over time, and the way it correlates with equity markets will change too. Analogously to Figure 2 we can calculate the rolling correlation of the VIX index with the S&P 500 index, plotting the results in Figure 3.

Figure 3: Correlation between the VIX and S&P 500 indices




Source: First Sentier Investors

³ All data calculated in USD. We use the S&P 500, MSCI World and MSCI Emerging Markets total return indices for their respective exposures, and the VIX and VDAX for volatility. We use the JP Morgan US Government bond index for US bonds and the USD one-month cash rate for USD cash.

The correlation has generally been negative with only a brief spike into positive territory on a 12-month rolling horizon. The more negative the correlation, the more favourable the hedging behaviour of volatility relative to equities. As the financial crisis in Europe intensified in the second and third quarters of 2011, we saw an increase in volatility and a decrease in correlation, the latter moving closer to -0.8 .⁴

Figure 4: Term Structure of Correlation



Source: First Sentier Investors

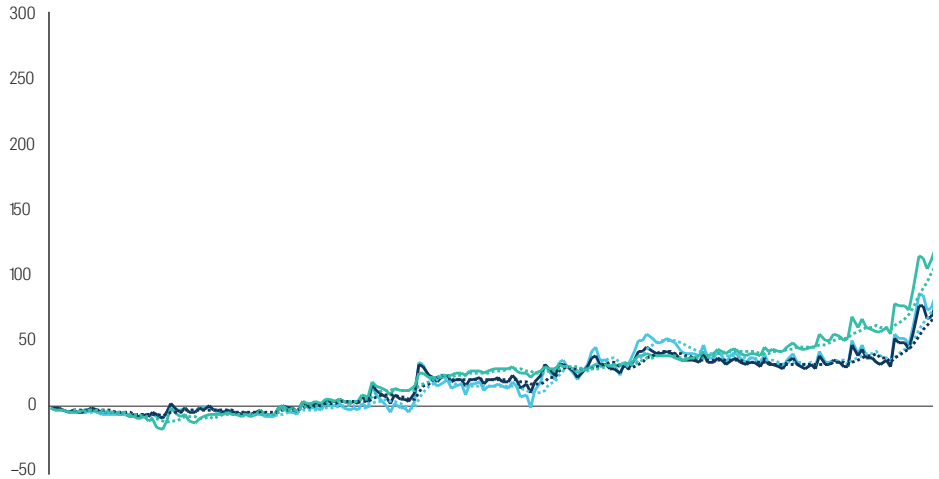
Measuring with a finer granularity of rolling window periods, we can calculate the term structure of the distribution of observed correlations over these window periods. The results are plotted in a correlation cone in Figure 4, with each measurement window period horizontally, and the correlations vertically. As the blue line that indicates the latest observations is hugging the bottom of the gray cone, we can deduce that the latest correlations are low historically for periods of 12 months and more. Only the relatively insignificant three and six month correlations are positive.⁵

⁴ This is based on a separate calculation using weekly data to capture shorter-term movements.

⁵ Essentially, these are correlations calculated using three and six pairs of data, and therefore are not statistically significant. Using weekly data we see negative correlations prevailing, approaching -0.8 .

Portfolio impact of adding Volatility

Figure 5: Relative Performance of Indices versus a static 10% Volatility Addition



Source: First Sentier Investors

We look at the performance of the S&P 500, the MSCI World and MSCI EM indices, as well as at the combination of 90% being invested in the equity index, and 10% in the VIX. Figure 5 shows that the portfolios that did include volatility as an asset class performed better than a pure equity portfolio would have. The graph shows the cumulative relative performance of a static 90/10 allocation between the index and volatility versus the index itself as dashed lines, and the six-month moving averages as solid lines. The static 90/10 allocation outperformed the index in a number of periods, especially in the turbulent times of 2008 to 2012. Extended periods of underperformance were in evidence in the more placid period of 2003–2006, for instance, which is more consistent with a negative risk premium for volatility.

The effects of this behaviour are also manifest in the downside performance characteristics of these asset combinations. In Table 2 we show the maximum drawdown in performance as well as the required recovery times. In the three cases we looked at here there would have been a marked improvement in performance with the addition of volatility exposure. The maximum drawdowns go from over 50% for the equity-only allocations to around 40% for the S&P 500 and MSCI World, and down from 61.4% to 48.6% for the MSCI EM.⁶

Table 2: Drawdown Characteristics Relative to Target 0% 241 Monthly Observations, Feb 1992–Feb 2012

Asset Mix	Maximum Drawdown	Drawdown Period Start	Drawdown Period End	Recovery in Months
S&P 500	50.9%	30 Nov 2007	31 Mar 2009	10
90% S&P 500 + 10% Vol	37.8%	30 Nov 2007	31 Mar 2009	
MSCI World	53.7%	30 Nov 2007	31 Mar 2009	
90% MSCI World + 10% Vol	40.6%	30 Nov 2007	31 Mar 2009	
MSCI Emerging Markets	61.4%	30 Nov 2007	31 Mar 2009	22
90% MSCI EM + 10% Vol	48.6%	30 Nov 2007	31 Mar 2009	41

Source: First Sentier Investors

⁶ A blank entry in the Recovery columns means the relevant asset mix has not recovered from its maximum drawdown.

This is also reflected in the parametrically modelled historical Conditional Value-at-Risk numbers shown in Table 3 and Table 4. Taking the one-year Value-at-Risk at 99% confidence as an example, the table shows us that the maximum expected loss occurring once in a hundred years is 23.4% for the S&P 500, 24.9% for the MSCI World and 37.4% for the MSCI EM.⁷ Adding an allocation of 10% to volatility reduces the Value-at-Risk to 14.7%, 16.2% and 28.6% respectively, which is a significant improvement. This improvement is replicated for different levels of confidence and different time horizons.

Table 3 Value-at-Risk Relative to Target r at Confidence Level c as Percentage of Invested Capital

Asset Mix	1 Year Horizon			3 Year Horizon			5 Year Horizon		
	c=90% r=0%	c=95% r=0%	c=99% r=0%	c=90% r=0%	c=95% r=0%	c=99% r=0%	c=90% r=0%	c=95% r=0%	c=99% r=0%
S&P 500	10.5%	15.2%	23.4%	8.6%	16.8%	30.3%	2.5%	13.7%	31.2%
90% S&P 500 + 10% Vol	4.4%	8.1%	14.7%	0.0%	2.6%	14.4%	0.0%	0.0%	8.9%
MSCI World	12.0%	16.7%	24.9%	12.7%	20.6%	33.6%	9.4%	19.9%	36.5%
90% MSCI World + 10% Vol	5.8%	9.6%	16.2%	0.0%	6.6%	18.2%	0.0%	0.0%	15.2%
MSCI Emerging Markets	19.7%	26.4%	37.4%	23.8%	34.4%	50.5%	22.6%	36.2%	55.6%
90% MSCI EM + 10% Vol	13.0%	18.8%	28.6%	10.6%	20.6%	36.5%	2.8%	16.6%	37.5%

Source: First Sentier Investors

The same applies to Conditional Value-at-Risk, which is the expected magnitude of a loss that exceeds the Value-at-Risk as shown above. Taking again the 99% confidence level as an example on the one-year horizon, we see that the expected loss of the S&P 500 would be 27.1%, in the case that the loss is bigger than the 23.4% Value-at-Risk above. One can regard the Value-at-Risk as a threshold level, and the Conditional Value-at-Risk as the average outcome if that threshold is exceeded.

Table 4 Conditional Value-at-Risk Relative to Target r at Confidence Level c as Percentage of Invested Capital

Asset Mix	1 Year Horizon			3 Year Horizon			5 Year Horizon		
	c=90% r=0%	c=95% r=0%	c=99% r=0%	c=90% r=0%	c=95% r=0%	c=99% r=0%	c=90% r=0%	c=95% r=0%	c=99% r=0%
S&P 500	16.5%	20.2%	27.1%	18.7%	25.0%	35.9%	16.1%	24.4%	38.3%
90% S&P 500 + 10% Vol	9.1%	12.2%	17.8%	4.3%	9.8%	19.6%	0.0%	2.5%	15.9%
MSCI World	17.9%	21.7%	28.5%	22.4%	28.5%	39.1%	22.2%	30.0%	43.1%
90% MSCI World+10% Vol	10.6%	13.6%	19.3%	8.3%	13.7%	23.2%	1.7%	9.1%	21.9%
MSCI Emerging Markets	27.9%	33.1%	42.1%	36.5%	44.2%	56.6%	38.5%	48.0%	62.4%
90% MSCI EM + 10% Vol	20.3%	24.8%	33.0%	22.8%	30.3%	42.9%	19.3%	29.3%	45.4%

Source: First Sentier Investors

Adding 10% in volatility improves the Conditional Value-at-Risk for all three cases we looked at, reducing the expected magnitude of the loss from 27.1% to 17.8% for the S&P 500 for instance. This again also holds for different levels of confidence and different time horizons.

⁷ This is based on historical data and does not include any stress-testing.

Dynamic Asset Allocation with Volatility

In practice the determination as to how much volatility to add is part of a larger strategic asset allocation and portfolio construction exercise, as other objectives, risk constraints and considerations need to be taken into account. It is also important to note that we would not hold a static percentage of the portfolio in volatility, but would vary this depending on market circumstances. For instance, the VIX index stood at 34.5 and the October 2011 future was trading at 33.15 with the entire term structure of VIX futures above 31.5. At such expensive levels we would be hesitant to commit any significant percentage of the portfolio to volatility exposure.

This is but a qualitative example of how one might use volatility. In practice we have a set of quantitative models to help us in our Dynamic Asset Allocation decision making that give buy and sell signals for volatility. Below in Figure 6 we present a simplified version of this model for the purposes of this paper. Running on the basis of monthly data, the strategy generates a buy signal whenever the VIX in the previous month increased by five percentage points or more, or when the VIX fell below 15%. Sell signals are generated when the VIX changed by less than

five percentage points and the VIX closed above 30%. Note that this example does not include trading costs, and simply uses the spot VIX for performance measurement. In actual portfolios one would have to account for both trading costs and the cost of carry for the futures, and pricing mismatches between the spot VIX and the futures. For a more complete overview of implementation and pricing issues, we refer to Carr and Lee, *Volatility Derivatives* 2009 and Bondarenko 2010.

Figure 6: Dynamic Asset Allocation Example using the S&P 500 and the VIX



Source: First Sentier Investors

Historically the return for this strategy would have been an annualised 12.1%, compared to 9.6% for the S&P 500 and 10.6% for the static 90/10 mix that we showed earlier. Volatility would have been 13.7% (versus 16.5% and 12.1% respectively), while the Sharpe ratio would have been 0.63 (versus 0.37 and 0.59).

Using the VIX as a timing signal is not new, and the above example shows that it is possible to create a profitable trading strategy even with simple rules. Other uses for the VIX as a trading signal have been proposed by Satchell and Scherer 2011 for hedge funds to hedge potential asset outflows.

Static allocations to volatility are not the optimal way of using this asset class.

Summary

Volatility as an asset class is based on futures which track the implied volatility of equity index options. Historically it has been a good partial hedge for equity market sell-offs, and with active tactical asset allocation it can be an extremely useful part of a well-diversified global portfolio.

References

Bakshi, G, and N Kapadia. "Delta-Hedged Gains and the Negative Market Volatility Risk Premium."

Review of Financial Studies 16, no. 2 (2003): 527-566.

Bondarenko, Oleg. "Variance Trading and Market Price of Variance Risk."

Working Paper, no. University of Illinois (December 2010).

Carr, Peter, and Roger Lee. "Variance Risk Premiums."

Review of Financial Studies 22, no. 3 (2009): 1311-1341.

Carr, Peter, and Roger Lee. "Volatility Derivatives."

Annual Review of Financial Economics 1 (2009): 319-339.

Satchell, Steve, and Bernd Scherer. "Buy-Side Risk Management: Hedging Hedge Fund Outflows."

The Journal of Alternative Investments 14, no. 2 (2011): 18-23.

Warren, Geoffrey J. "Can Investing in Volatility Help Meet Your Policy Objectives?"

The Journal of Portfolio Management 38, no. 2 (Winter 2012): 82-98.

Whaley, Robert. "Understanding the VIX."

The Journal of Portfolio Management 35, no. 3 (2009): 98-105.

Important Information

This document is not a financial promotion and has been prepared for general information purposes only and the views expressed are those of the writer and may change over time. Unless otherwise stated, the source of information contained in this document is First Sentier Investors and is believed to be reliable and accurate.

References to “we” or “us” are references to First Sentier Investors.

First Sentier Investors recommends that investors seek independent financial and professional advice prior to making investment decisions.

In the United Kingdom, this document is issued by First Sentier Investors (UK) Funds Limited which is authorised and regulated in the UK by the Financial Conduct Authority (registration number 143359). Registered office: Finsbury Circus House, 15 Finsbury Circus, London, EC2M 7EB, number 2294743. Outside the UK, issued by First Sentier Investors International IM Limited which is authorised and regulated in the UK by the Financial Conduct Authority (registration number 122512). Registered office 23 St. Andrew Square, Edinburgh, EH2 1BB number SC079063.