The Financial Times [FT] recently published an article [19/03/2010] titled "Analysts scramble to decipher calmer Vix". In the article they quote some research by Birinyi Associates:

> An analysis of the VIX and moves in the S&P 500 by Birinyi Associates, published this month, reached the same conclusion. "The VIX is alleged to be an indicator and has become a staple of analysts and journalists alike," the report says. "We respectfully disagree and ultimately conclude it is a measure of current volatility with little or no predictive or indicative value regarding the course of the market."

Given that JFCP uses the Australian VIX in its research process to help determine the discount rate used to value Australian listed companies, I thought I would provide a critical review of the article using both US data and Australian data (note the Bririnyi Associates research used just US data).

Firstly on the issue of the predictive power of the VIX: I agree that statistically the VIX appears to have limited, if any, directional predictive power, but it's not as 'black and white' as the article would have you believe.

Looking at the US data first [monthly data from February 1986 to February 2010: ***S&P 100 OEX Volatility Index Change*** and ***S&P 500 Total Return***]. The contemporaneous correlation between the change in the US VIX and the US equity market return is -0.60.

This means that contemporaneously as the US VIX increased, the US equity market return was highly likely to be negative, and vice versa. Intuitively this makes sense: more risk means lower equity prices to deliver higher future equity returns to investors to compensate them for the increased level risk (*ceteris paribus*).

However, this is contemporaneous and tells us nothing about the US VIX's predictive ability. What if we lag the US equity market return by say 1, 2, or 3 months to test the directional predictability? When we do this the correlation falls dramatically, and in fact becomes positive [+0.10, +0.20, +0.11, respectively]. So there is no obvious sign of directional predictability here.

The chart below shows the two data series. The US VIX is the blue line and the US equity market index [i.e. log scale base 10] is the red line.

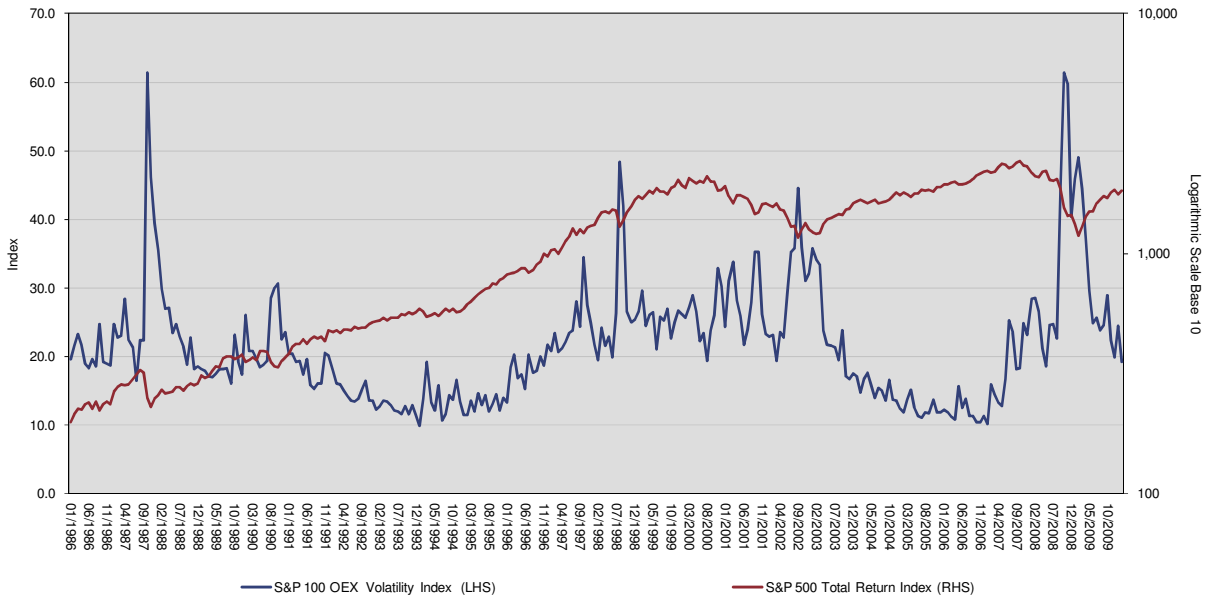

Now let's look at the Australian data [here we have weekly data from January 1997 to March 2010: ***S&P/ASX 200 Options Volatility Index Change***, ***Australia ASX All-Ordinaries Total Return***, and ***All Ordinaries Historic Volatility Change***]. The contemporaneous correlation between the change in the Australian VIX and the Australian equity market return is -0.68, which is slightly higher than the US monthly data. This intuitively makes sense and also argues in favour of market efficiency [e.g. higher risk leads to lower market price and higher expected return to compensate for the increased level of risk].

However, what if we lag the Australian equity market return by say 1, 2, or 3 weeks to test the directional predictability? The correlation fall is again dramatic [+0.01, -0.01, +0.03, respectively]. So once again there is no obvious sign of directional predictability in this data either.

This -0.68 figure is shown in the matrix at the top of the next page.

Correlation Matrix

	S&P/ASX 200 Options Volatility Index Change	Australia ASX All-Ordinaries Total Return
S&P/ASX 200 Options Volatility Index Change	1.00	
Australia ASX All-Ordinaries Total Return	-0.68	1.00
All Ordinaries Historic Volatility Change	0.22	-0.18

What about the other correlations in the matrix above [i.e. ***All Ordinaries Historic Volatility Change*** v. ***S&P/ASX 200 Options Volatility Index Change***, and ***All Ordinaries Historic Volatility Change*** v. ***Australia ASX All-Ordinaries Total Return***]?

On the first of these, as you would expect, there is positive correlation of +0.22 between the change in the Australian equity market's historic volatility[1] and the Australian VIX. This positive correlation is consistent with investors who observe an increase in the historic volatility over the preceding 26 weeks, and thereby increase their estimates of future volatility, or vice versa. Why isn't the correlation higher? Because the VIX is a current measure of risk and the historic volatility has an inherent moving average anchoring effect within it, this will statistically lower the correlation between the two data series.

On the final correlation in the matrix [***All Ordinaries Historic Volatility Change*** v. ***Australia ASX All-Ordinaries Total Return***], this negative correlation also makes intuitive sense [i.e. higher historic risk impacts investor risk expectations and leads to a lower market pricing thereby increasing the expected market return to compensate investors for the higher expected risk]. But note that it's much lower than the correlation between the Australian VIX and the Australian equity market return. **Therefore, faced with the choice of which measure of risk is the best representation of the current risk expectations being priced by investors, then it seems clear that the VIX is far superior to historic volatility.**

Despite the correlation statistics described above [and because it's well known that statistics lie[2]], before we move on from the issue of the VIX's directional predictability, it's worth 'eyeballing' the chart below [i.e. pictures don't tend to lie as much as statistics].

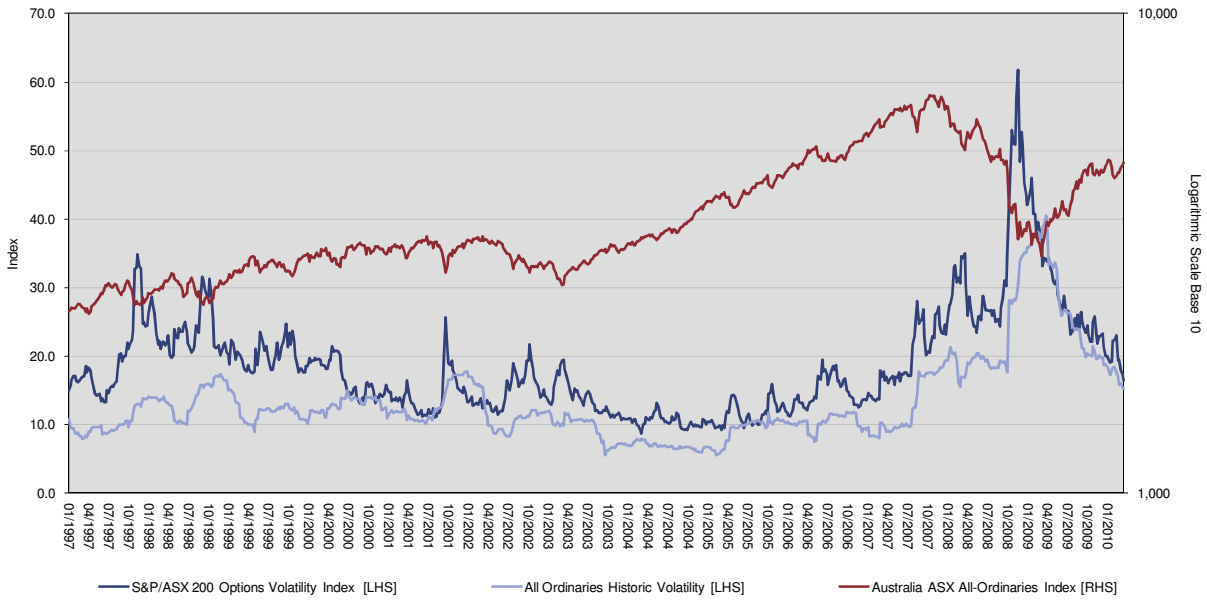

For mine, the Australian VIX provided a clear early warning signal throughout late 2006 and 2007 of the impending equity market crisis that hit in 2008. The Australian VIX rose 114% between 31/3/2006 [from 12.07] and 4/1/2007 [to 25.83]. This early warning is harder to detect in the monthly data for the US, although there was a noticeable rise [39%] in the US VIX between June 2007 [from 16.69] and December 2008 [to 23.19].

[1] Historic volatility is measured as the annualised standard deviation of the previous 26 weekly total returns from the Australian equity market.
[2] "Lies, damned lies, and statistics" is a phrase describing the persuasive power of numbers, particularly the use of statistics to bolster weak arguments.

The FT article also quoted Birinyi Associates as saying that:

> "It [the Vix] might, however, be useful in that it appears that high volatility might actually be a contrarian indicator,"

I believe this is a spurious argument. From the chart on the previous page you can clearly see that the Australian VIX was on average a contrarian indicator between mid 2006 and 2007 because the Australian equity market rose along with the Australian VIX which is counter-intuitive, but we all know how that story ends!

The FT article seems to confuse the contemporaneous increases in the expected return on the market and the likelihood of large negative realised return that is needed to achieve this higher expected return. Until a large negative return is subsequently realised, a high VIX will appear to have been predicting higher market realised returns [e.g. March 2006 to January 2007]. But when the large negative return is eventually realised, the same high VIX will seem also to have predicted the lower realised market returns.

Is the VIX therefore a contrarian indicator? The answer [for sound reasons] is that after analysing any realised return series [such as the equity market]... yes and no, as we can see from the chart on the previous page.

Whether the recent fall in the Australian VIX in a harbinger of a future large negative equity market return [i.e. the contrarian argument], and/or a consequence of excess global liquidity, is a moot point.

On this point, the FT article points out that:

> The decline in volatility – and the gains in stock markets, corporate bond markets and commodity prices in the past year – is very much a reflection of the record amounts of liquidity that central banks have pumped into the global financial markets in the past 18 months and the continued record low official interest rates.

I reckon this point is pretty much 'on the money', and highlights that there are times when the VIX itself can be significantly mispriced as a result extreme fear and falling liquidity [e.g. from August 2007 to November 2008], or ample liquidity [e.g. from December 2008 to March 2010] and 'irrational exuberance' [e.g. from October 2002 to April 2005].

In an excellent research paper by Tobias Adrian and Hyun Song Shin titled "Liquidity and Leverage" (2008), the authors provide empirical evidence of the link between expanding liquidity and financial market volatility. See an extract of this research paper below:

> Our second objective is to pursue the aggregate consequences of such pro-cyclical leverage, and document evidence that expansions and contractions of balance sheets have asset pricing consequences through shifts in financial market volatility. In particular, we show that changes in collateralized borrowing and lending on intermediary's balance sheet are significant forecasting variables for innovations in market-wide risk as measured by the VIX index of implied volatility in the stock market. We also decompose VIX innovations into changes of stock market volatility and changes of the difference between implied volatility and actual volatility (the volatility risk premium). We find that dealer balance sheet changes primarily forecast changes in the volatility risk premium, which has a natural interpretation as the price of risk of aggregate volatility. Previous work in asset pricing has shown that innovations in market volatility are important cross sectional asset pricing factors (see Ang, Hodrick, Xing, and Zhang (2006), and Adrian, and Rosenberg (2008)), and that the volatility risk premium forecasts future equity returns (Bollerslev and Zhou (2007)). Our finding that expansions and contractions of the balance sheets of security dealers forecast volatility innovations shows that intermediary balance sheets matter for the aggregate pricing of risk. Consistent with the conjectures of Gromb and Vayanos (2002), Brunnermeier and Pedersen (2007) and He and Krishnamurthy (2008), we thus document that funding liquidity of financial intermediaries has aggregate pricing implications.
>
> Our findings also shed light on the concept of "liquidity" as used in common discourse about financial market conditions. In the financial press and other market commentary, asset price booms are sometimes attributed to "excess liquidity" in the financial system. Financial commentators are fond of using the associated metaphors, such as the financial markets being "awash with liquidity", or liquidity "sloshing around". However, the precise sense in which "liquidity" is being used in such contexts is often left unspecified.
>
> Our empirical findings suggest that financial market liquidity can be understood as the rate of growth of aggregate balance sheets. In response to increases in prices on the asset side of intermediaries' balance sheets, leverage falls, and intermediaries hold surplus capital. They will then search for uses of their surplus capital. In a loose analogy with manufacturing firms, we may see the financial system as having "surplus capacity". For such surplus capacity to be utilized, the intermediaries must expand their balance sheets. On the liabilities side, they take on more short-term debt. On the asset side, they search for potential borrowers that they can lend to. Financial market liquidity is intimately tied to how hard the financial intermediaries search for borrowers.

Despite the fact that the Australian VIX appeared to be a good predictor of the Australian equity market fall in 2008 and early 2009, and despite the pro-cyclical leverage argument put forward by Adrian and Shin, let's go with the correlation statistics above and the proposition in the FT article, and assume that there is little, if any, directional forecasting ability in the change in the VIX for predicting future market returns or changes in historic market volatility. **What about the ability of the VIX to predict the level of equity market volatility in the future?**

On this issue the statistical story is very compelling: the correlation between the level of Australian VIX and: (1) the forward 26 week volatility of Australian equity market returns is +0.68; (2) the forward 104 week volatility of the Australian equity market is +0.42.

The charts below show a graphical depiction of these correlations.

The first chart shows the VIX and the 26 week forward historic volatility.

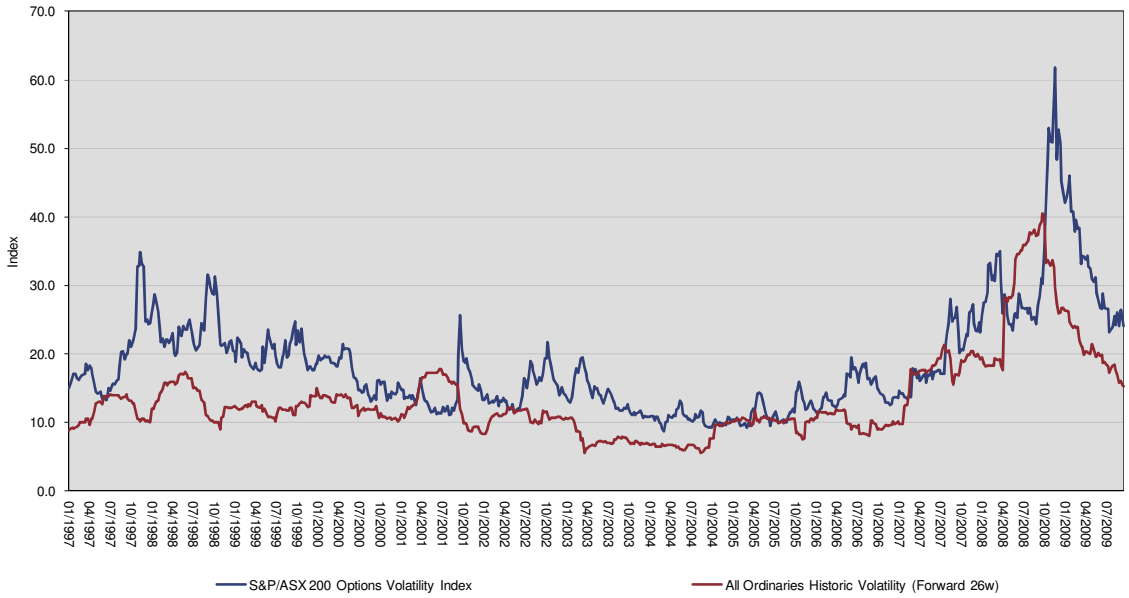

The second chart shows the VIX and the 104 week forward historic volatility.

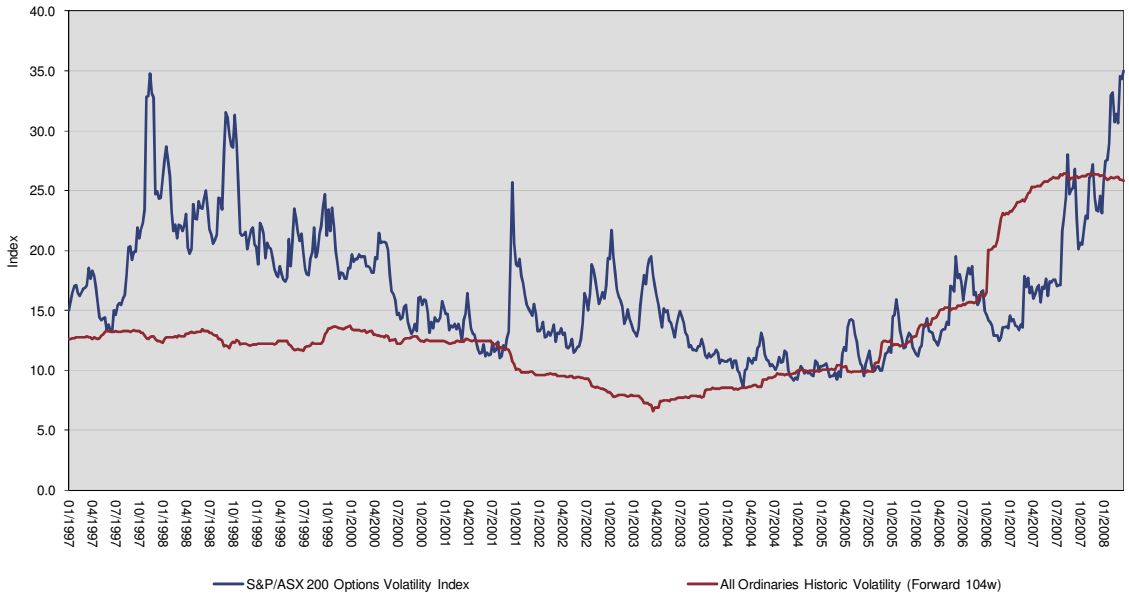

This latter correlation is important because JFCP uses the VIX as a short-term [i.e. 2 year forward] **forecasting tool to estimate the short-term future level of volatility of the Australian equity market.** This short-term forecast of Australian equity market risk is then used to calculate our short-term estimate of the equity risk premium [ERP] over the next 2 years[3], which in turn goes into calculating a company's cost of equity[4] and cost of capital[5] which is used as the discount rate in our discounted cash flow [DCF] stock valuation model[6].

So what does this mean for portfolio management at JFCP? If the VIX is high, and we therefore expect higher equity market risk over the next 2 years, then our cost of capital forecast for the market will be higher, our DCF valuations will be lower, and (*ceteris paribus*) our expected stock returns will be lower; and vice versa if the VIX is low.

For our Absolute Return product this will mean more cash and lower equity beta[7] stocks will be favoured. For our 'fully invested' products, this simply means lower equity beta stocks will be favoured relative to higher equity beta stocks (*ceteris paribus*).

For mine, a higher than average VIX is sending a clear message to investors: the probability of higher future market volatility and a negative market return in the future are both greater than normal. The lower price of the market [i.e. an expected higher return (*ceteris paribus*)] is a necessary inducement to compensate investors for the perceived increased level of risk. Alternatively, a lower than average VIX would send the opposite signal. That is, an increased probability of lower future market volatility and an expected positive market return in the future. The resultant higher market price reduces the expected return needed to compensate investors for the decreased level of risk.

There is little doubt in my mind that JFCP's strong performance surrounding the GFC [i.e. our Benchmark Based equity product outperformed the S&P/ASX300 Accum. Index by +9.4% from August 2007 to March 2009] was the result of not only good individual company key parameter forecasting [i.e. revenue, operating margin, and capital investment] by our analysts, but also our ability to dynamically adjust the cost of capital for increased market risk leading into the GFC [via the ERP], and thereby sell more highly leveraged [i.e. high equity beta] stocks and buy more defensive [i.e. lower equity beta] stocks ahead of many other investors desire to do so.

Therefore, I believe the VIX is a very useful risk forecasting tool when used in conjunction with a risk-dependent valuation model [e.g. the DCF model] **for estimating required rates of return, and can be used to profitably allocate between low and high equity beta investments** [e.g. cash and equities, and/or low and high equity beta stocks].

Michael Fitzsimmons
Managing Director / Chief Investment Officer
JF Capital Partners

[3] **Short-term forecast ERP** = VIX ÷ 100 x Price per unit of risk. Price per unit of risk = Long-term historic ERP ÷ Long-term historic volatility.
[4] **Cost of equity** = Risk-free rate + Equity risk premium x Equity beta.
[5] **Cost of capital** = (Cost of equity x Equity value + Cost of debt x Debt value - Risk-free rate x Surplus cash) ÷ Operating value. **Operating value** = Total value – Surplus cash.
[6] **DCF valuation model** = Free cash flow ÷ (Cost of capital – Growth rate).
[7] A stock's equity beta is defined here as a measure of the estimated future leverage of the stock's equity. This equity leverage is a combined function of the stock's economic leverage, operating leverage, and financial leverage.