Woodhall's Stock Market Forecast Methodology

Ron Bewley PhD, FASSA

An internally-consistent strategy

> ➤ **Sector and market forecasts are based on consensus broker forecasts**
> ➤ **Market entry and exit signals are based on sector forecasts**
> ➤ **Market forecasts of intra-year highs and lows depend upon forecasts of returns and volatility.**

_____

Over-riding principles

Since we want to achieve a number of objectives with our forecasts – such as portfolio construction and market timing – it is important to us that all pieces of our framework fit together in a *logically consistent manner*.

We also want a numerically-based solution – rather than a qualitative approach – since we want to know that any change in any of our forecasts ripple through the whole system in an appropriate way.

Importantly, a formulaic approach also allows us to perform a forensic analysis on any big forecast errors – so that we can learn as quickly as possible from our mistakes.

Finally, we want to use the best publically available data to form our views.

Because of the approach we have chosen we can readily analyse any other market in the world that has the same the same data collection service that we use in a manner that is consistent with our ASX 200 approach. Since 2011 we have also monitored the S&P 500 for US stock, sectors and the market.

Broker forecasts

A number of agencies, such as Thomson-Reuters and Bloomberg, collect forecasts from broking analysts. These analysts usually work in teams to produce a spreadsheet financial model of a company – or the equivalent. Not all broking houses cover all stocks in the ASX 200 but usually there is at least one set of forecasts for each company.

In Chart 1 we show the number of companies covered in the ASX 200 against the number of forecasts submitted to Thomson-Reuters for next year's earnings. On the left we can see that there were only two companies that had only one broker forecast. On the right we can see there was only one company covered by 17 broking houses. The peak (also known as the mode) is 28 companies that were covered by 11 analysts. 197 companies have at least three broker forecasts. 95% of the top 100 stocks have at least eight broker forecasts.

Chart 1: Distribution of coverage of stocks

Source: Thomson Reuters Datastream. Data as at May 18th 2018.

WoodHall
INVESTMENT RESEARCH
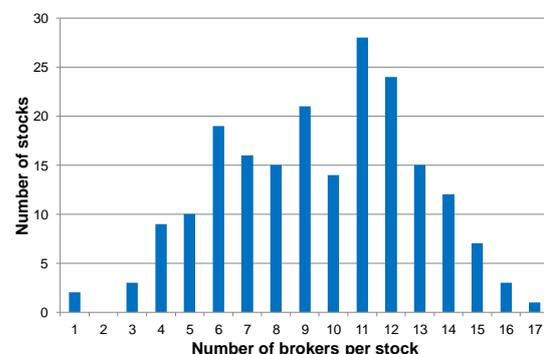
Since each broking house may have several analysts covering any one company, but the same team might cover several companies, it is not feasible to calculate how many analysts collectively contribute to these surveys. However, we can conclude from summing the data over companies and brokers, there are forecasts from over 1,900 models being submitted to Thomson Reuters that enable us to analyse the ASX 200.

Brokers forecast one or more of a number of attributes for each company. The main statistics we use are:-

1) Earnings up to three years ahead.

2) Dividends up to three years ahead.

3) Target stock prices – usually considered to be a forecast of a stock price about one year out.

4) A buy/sell recommendation using the scoring system,

   1 – buy

   2 – outperform

   3 – hold

   4 – underperform

   5 – sell

Earnings forecasts

Companies are required by regulation to provide so-called 'continuous disclosure'. That is, they must swiftly inform the exchange when they are aware of any factor that might materially affect their share prices.

Companies usually give profit guidance – often in the form of a range of possible profits or earnings. Should a company expect that range to have changed, it must inform the exchange. In that sense all brokers have good information on which to base their earnings' forecasts.

By averaging broker forecasts of earnings to form the 'consensus average', we believe that we get a reasonably reliable estimate of earnings for the first year but, of course, forecasts further out than one year are likely to be less precise.

Not only do analysts have access to the continuous disclosure earnings' forecasts provided by the company, they can also add their own views – including information gleaned from site visits by the analysts.

Broking analysts are typically very highly paid – with some earning seven-figure packages. They are incentivised to win 'best analyst' awards from these forecasts. Part of thesis is – as academic research suggests – that the consensus view cannot consistently be beaten by any individual forecaster.

Dividend forecasts

The dollar value of dividends is usually much easier to predict than earnings because each company has some control over the amount of earnings they will distribute in the form of dividends – called the 'payout ratio'.

Moreover, some companies, such as the big banks, are relied upon by investors to provide a steady income stream so that they are incentivised to alter their payout ratios as necessary (and if possible) to maintain a steady or improving dividend stream. Occasionally companies even borrow when earnings are insufficient to maintain the expected dividend stream.

Franking credits might also be distributed to investors. A company that has paid sufficient tax in Australia can pass on a full tax credit to the investor at the company tax rate. That ensures an investor is not taxed twice: once on earnings in the form of company taxes, and secondly as income tax by the investor.

We know of no ready source from which to collect forecasts of franking credits. Historically, there has been an average franking credit rate that has often been in the range of 70% - 80% but the ratio varies across sectors and over time. We do not specifically use franking credits forecasts in our modelling as we have no reliable forecast input.

Target prices

While analysts have a benchmark provided by the company from which to base their earnings and dividend forecasts, there is no such benchmark for target prices.

There are a number of ways in which to value a company but there is a strong incentive – in academic forecast theory – *not* to reveal one's true forecast if it is 'too far' from the current value and those of other forecasters – in case the forecaster is 'wrong and alone'. Analysts can change their forecasts as often as they like – and they do – to inch towards their unrevealed 'best' forecast. This process is called 'rational cheating' in the academic literature.

We have not done a formal analysis with stock prices (but we have published an academic paper

on interest rate forecasting in a multi-country analysis of hundreds of forecasters. This study was also the feature story on a 7:30 Report ABC TV program in 1996 but our general feeling from years in the industry is that consensus target prices contain little useful information. We believe target prices tend to lag actual share price movements by a large degree.

Target prices have no place in our forecast methodology – although they do play a minor, qualitative role in our portfolio construction monitoring.

Consensus recommendations

There is a clear and well-known bias in consensus recommendations from brokers. Outright sell recommendations are very rare indeed. In Chart 2 we show the distribution of recommendations across the ASX 200 for a certain day. Since this paper is designed to be expository, the date is irrelevant – but it happens to be 18th May 2018. We post various statistics and forecasts each Saturday on our website. Interested readers of our *Woodhall Weekly* are invited to refer back to this paper to save repetition in the *Weekly*.

Chart 2: Distribution of recommendations
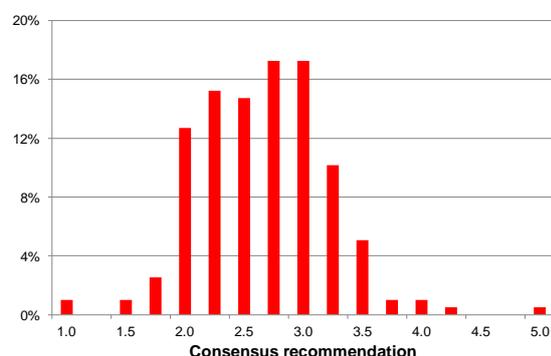
Source: Thomson Reuters Datastream. Data as at May 18th 2018.

Because the consensus recommendations (hereafter, 'recs') are averages of the individual ratings' scores across brokers, the recs are not normally integers (round numbers) as in 1, 2, ... 5. Indeed, it is almost impossible to score a '1' as that would require every reporting broker to give a 'one' on a given stock. On May 18th 2018 LYC and TME each had one broker recommendation that happened to be a 'one'. CNU also had only one broker recommendation and that was a 'five'.

As we experienced in the industry, companies often phone analysts up to complain when they get a bad rating so there is an incentive to disguise one's true

views – a little like the rational cheating we discussed above in reporting target prices.

Another problem in using recs is that the people who make recs on, say, mining stocks are typically not the same people as those who rate the banks, etc. Indeed, there may be no cross-checking of ratings across sectors – and the analysts might not even be in the same city!

The average rec across all 200 stocks was 2.7 and the joint 'best' for sector averages were Discretionary and Property at 2.4. The worst sector was Telcos at 3.2.

We find no *absolute* merit in recs, but we do have a strong belief in within-sector comparisons of recs. In other words, analysts considering, say, CBA and WBC (or, say, BHP and RIO) have a reasonable basis for comparing their recs.

Indeed, we believe that we have had a strong track record for over a decade from choosing stocks in portfolios based only on *within-sector* recs (but sector allocations being based on other criteria) which we call 'High Octane' portfolios. We, at times, have had quite impressive returns with these portfolios over the last decade or so but, of course, with high risk.

Woodhall philosophy

We strongly believe that it is hard to consistently beat consensus forecasts. There is theory and applied work dating back to the 60's and 70's from the late Nobel Laureate, Professor Sir Clive Granger, supporting this notion.

We note that to be successful in investing, it is not necessary to be the best – at any point in time – but being consistently good has a lot of merit. For that reason, we do not expect to be top of the performance rankings at any point in time as that requires taking on too much risk.

We base our sector and market forecasts of capital gains and dividend yields solely on Thomson Reuters' consensus earnings and dividends forecasts and our quantitative analysis of them.

In portfolio construction, we select stocks *within* a sector based on criteria that include recs, expected yield, and market capitalisation – depending upon the 'style' of the portfolio: High Conviction; High Yield; and High Octane. We blend the Conviction and Yield portfolios with time varying weights to construct our flagship 'Woodhall' portfolio. These blending weights only changed twice between February 2014 and May 2018.

Target prices, in our analyses, are nothing but cross-checks that help us think and, possibly, question our quantitative approach.

Sector and market forecasts

In much the same way that we prefer the consensus forecasts to individual forecasts, we strongly believe that it is far easier – or more robust in forecasting parlance – to forecast a sector's performance than an individual company's performance.

As a result, we aggregate data on earnings and dividend forecasts for three years out across stocks in a given sector to produce a capital gains and a total return forecast. These forecasts are then weighted across sectors to produce a market forecast. The horizon for all of these forecasts is for twelve-months ahead.

We will not divulge how we perform this last step as it is a matter of our intellectual property. We stand by our track record and the internal logic of the calculations we perform.

The current system was devised in 2010 and it is a major advance on the work that was done at CBA in the previous six or seven years – but the new approach is not inconsistent with it.

It follows that if one analyst changes his/her earnings and/or dividends forecasts for one company, there is a ripple effect through to the consensus forecasts for that company – then through to that sector's forecasts, as calculated by us – and finally through to our market forecast. *That is what we mean by logical consistency.*

We use the end-of-month forecasts for a number of purposes on a rolling twelve-month-ahead basis. Although we update these forecasts daily, we do not use daily updates for anything other than a cross-check to look for emerging trends and a possible signal for an 'amber alert'.

Because most companies release their annual and half-yearly results in February and August, we tend to find a little more volatility in the daily forecasts just before and during February and August. Analysts often sensibly withhold their forecasts while they are updating their models. Even though their forecasts might not change by much before and after the reports, the simple withdrawal of a forecast from the consensus panel can cause noticeable changes in the consensus forecast.

We show our forecasts for yield and capital gains – along with the index weights – by sector in Table 1 for illustration.

Table 1: Market shares capital gains by sector

	Sector	Index weights	12 month forecasts	
			yield	cap gain
Resource-related	Energy	5.2%	3.1%	5.0%
	Materials	19.0%	3.1%	4.2%
	Industrials	6.5%	4.0%	6.2%
High yield	Fin-x-REIT	34.4%	7.0%	5.3%
	REITS	7.5%	4.8%	4.1%
	Telco	2.5%	7.0%	3.8%
	Utilities	1.8%	5.9%	10.2%
Other	Discretionary	4.8%	3.3%	14.0%
	Staples	7.4%	4.1%	6.8%
	Health	9.3%	1.5%	10.5%
	IT	1.8%	2.0%	14.5%
	ASX 200	100.0%	4.7%	6.3%

Source: Thomson Reuters Datastream; and Woodhall Investment Research. Data to May 18[th] 2018

For some of our work, we group the sectors as shown in column 1 of Table 1 into three aggregated sectors – but this is not a rigorous definition – High-Yield, Resource-related and Other.

Since we do not use the current day's stock prices in these calculations, the capital gains forecasts need to be adjusted for practical purposes. We think of the forecasts in Table 1 to be based on the 'fundamentals' and the forecasts that we produce should be based on the current 'fair price' of the relevant index. That is, we need to measure how much above or below that fair price the current index value is. We call this mispricing 'exuberance'.

Measuring exuberance

If exuberance is, say, -5% the sector is said to be cheap by an amount of 5%. Conversely, a positive exuberance measure implies that the sector, or market, is expensive.

The fundamental assumption that we make in measuring exuberance is that the forecasts are 'credible'. That is, if, say, the index grows faster than the predicted trend line, the market is becoming more expensive.

Since new information may come to light that changes the forecast trend, we cannot rely on just one single forecast origin – such as January 1[st] 2018. Rather, we add each new 12-month rolling forecast on the first day of each month to the analysis.

Although we do not revise our old forecasts we can reasonably project each of our new forecasts backwards at the same trend rate. In this sense, we can tell which of the recent forecasts have been fitting current data better than other forecasts.

In a fairly complicated way, we are able to 'back out' what we think the fundamental value of the index is and, hence, what is the exuberance measure. Since

we do not revise past values of fundamentals or exuberance, these calculations are legitimate.

We show our exuberance index since 2002 in Chart 3 for the ASX 200. What we have noted across the 11 sectors – and the S&P 500 index and its 10 sectors – that +6% is a reasonably reliable indicator for calling a significant overpricing in the sense that a correction of -6% to -10% is imminent, or there will be a protracted sideways movement in the index.

We have successfully used the exuberance indexes in portfolio management for over a decade. We typically choose not to buy in a sector that is more than about +2% overpriced.

Chart 3: ASX 200 mispricing
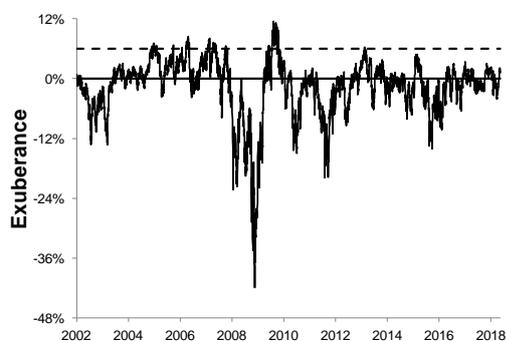
Source: Thomson Reuters Datastream; and Woodhall Investment Research; data to close 18th May 2018

When rebalancing, we prefer to first sell relevant stocks in sectors that are overpriced and buy them in sectors that are cheap. In big portfolios we have sometimes taken two or three months to complete a full portfolio rebalance. We believe that we add value from this process.

Adjusted capital gains

Now that we have a measure of mispricing, we can combine our exuberance measures with the capital gains forecasts from Table 1 to produce an adjusted forecast from the current index value. We show the impact of these calculations in Table 2. Naturally these adjusted forecasts can change markedly from day to day with swings in market prices.

We note from Table 2 that the Energy sector is above the so-called 'magic +6%' level – at +8.7% – indicating that a correction or sideways movement is quite probable. As stated earlier, this signal is not necessarily one to sell but the adjusted capital gains forecast of -4.1% suggests that the dividend yield of 3.1% (Table 1) is likely to be eroded – and more – by capital losses. However, the investor who bought this index when it was cheap can withstand this sort of volatility and still recoup a good yield.

While the large negative exuberance level for Telcos makes the adjusted capital gains very attractive, this signal could be misleading. For example the dividends may be cut rendering a probable downgrade in forecast stock prices.

Table 2: Adjusting capital gains forecasts

	Sector	Exuberance	12 month forecasts	
			cap gain	adj gain
Resource-related	Energy	8.7%	5.0%	-4.1%
	Materials	7.6%	4.2%	-3.7%
	Industrials	-1.4%	6.2%	7.7%
High yield	Fin-x-REIT	-4.8%	5.3%	10.4%
	REITS	0.7%	4.1%	3.3%
	Telco	-15.5%	3.8%	19.9%
	Utilities	-6.7%	10.2%	17.5%
Other	Discretionary	1.6%	14.0%	12.2%
	Staples	2.5%	6.8%	4.1%
	Health	9.6%	10.5%	0.0%
	IT	3.2%	14.5%	10.8%
	ASX 200	1.2%	6.3%	5.0%

Source: Thomson Reuters Datastream; and Woodhall Investment Research. Data to May 18th 2018

As a rule of thumb we have used for much more than a decade, a +6% overpricing has proven to be a reasonable trigger to call a correction of -6% to -10% – or a prolonged sideways market movement while the fundamental 'catches up' to erode the overpricing. However, it can be noted that higher levels were achieved in late 2009.

At that time, the market was recovering from the March 2009 low and, with interest rates very low, investors seemingly were prepared to pay over the odds just to get dividends and some exposure to possible capital gains in the long term. The market moved sideways for a very long time after that period of excessive exuberance.

We have no rule for significant underpricing but we do have a fear index and a disorder index to help interpret significant underpricing. We argue that when there is elevated fear and disorder, markets are likely to fall more quickly and further than under normal conditions. Moreover when there is elevated fear and disorder, cheap markets are more likely to fall further – or stay cheap for longer – than under normal conditions.

Measuring fear

We came up with the concept of measuring fear in February 2008. The ASX 200 had fallen sharply from a high in November 2007 but, importantly, the index was going through massive gyrations during the trading days.

Data are available not only on the open and the closing index values but also the high and the low. Our fear index defines excess volatility over that which might normally be observed for a given open and close. We back-calculated fear from historical data to find what the normal range for our index

should be. We represent this range in Chart 4 along with the index itself.

Chart 4: The ASX 200 fear index
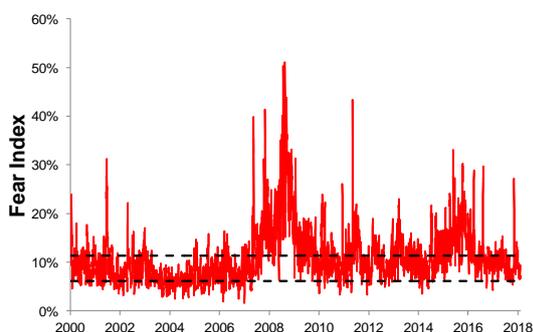
Source: Thomson Reuters Datastream; and Woodhall Investment Research. Data to May 18[th] 2018

When the index is well above the upper 'tramline', or consistently above the upper tramline for a number of days, we conclude that the market is too fearful to act rationally.

At the time of its development in 2008 there was no 'Aussie VIX' index. Of course the CBOE VIX index, which measures implied volatility on the S&P 500, was then available.

It does not follow that our market necessarily follows the US market in terms of volatility. The VIX is based on investors trading options that reflect those investors wanting to take out 'insurance' on future market moves. We argue that activity takes place after market gyrations have occurred in observed market prices. In other words, we argue that our fear index leads the Aussie VIX, and our S&P 500 version of our fear index – the actual VIX.

Measuring disorder

The concept of disorder stems from that fact that, in normal times, the returns on different sectors on the same day move reasonably closely together. In other words, the 'market component' of sector returns dominates.

We have observed that, at some times, sectors more in sharply different directions. This behaviour happened markedly during the GFC when Financials sector returns and Resources sector returns often moved sharply in opposite directions as investors tried to make sense of market information.

As with our fear index, we define normal behaviour by tramlines – as we showed in Chart 4. In both cases there is a two thirds chance of being within the range on one sixth above or below.

Chart 5: The ASX 200 disorder index
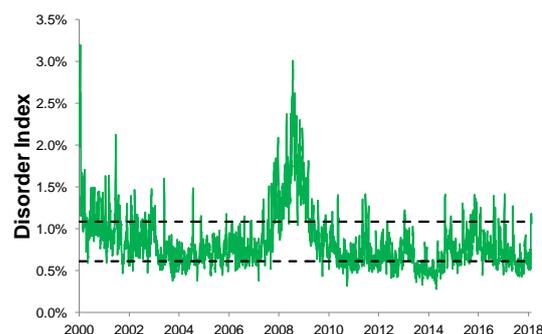
Source: Thomson Reuters Datastream; and Woodhall Investment Research. Data to May 18[th] 2018

We believe that the fear index should dominate the signals created by the disorder index. When both indexes are above normal, we argue that the level of fear has been amplified.

As it happens, the first time both fear and disorder were in the normal range (after we created the indexes in early 2008) – and the market, as measured by exuberance, was cheap – was in early March 2009 or the very bottom of the market!

Measuring volatility

Our main estimate of market volatility comes from our modelling of realised market volatility. That is, we measure volatility separately within each month using the 'standard deviation' of the daily market index returns – which is then scaled up to become an annualised estimate.

Many analysts use an historical measure of volatility over, say, three or five years. We rejected that approach, even well before the GFC, since such long 'windows' of data make the estimate slow to react to a genuine change in the underlying level of volatility.

We developed a method which allows for the average to change from time to time – but not over-react to noise. The method assumes that the underlying average might have changed this latest month, or the month before, etc, back to 12 months before.

In what is an extremely sophisticated statistical procedure, the system 'learns' as new data become available. We measure the strength of any new signal having entered the data. When the signal is sufficiently strong we 'lock it in' but repeatedly test (going forward) whether or not we had the best past location for a 'mean break'.

The theory took well over a decade of intensive academic research to complete. Although we

published the results in an academic journal, we do not know of anyone who has developed the machine-learning part of the system that we now use.

We show the monthly estimates for the ASX 200 long-run volatility in Chart 6 as the thick black line. Short-run forecasts are show in blue. It is clear that the long-run estimate is indeed stable for long periods of time – but it more than doubled as early as March 2008 and sharply returned to 'normal' levels in February 2013. We use this methodology on sectoral data for our portfolio construction. We use this market level index in our production of estimated forecast bounds to accompany our market capital gains forecasts.

Chart 6: The ASX 200 long-run volatility index

Source: Thomson Reuters Datastream; and Woodhall Investment Research. Data to April 28[th] 2018

Forecast accuracy

Any student of the market understands the extent of day-to-day volatility, even in calm times. Therefore, it is a thankless task to try and make a forecast for a year-end value on the ASX 200.

Once one loosens the requirement to one of forecasting a value 'somewhere near' the end of the year – say to within a month either side – the task becomes far more reasonable.

Moreover, we believe that many people when making, or interpreting, end-of-year forecasts are really thinking in terms of the maximum during the year. As a result we supplemented our e-o-y forecasts with 'forecast boxes'.

We show our calendar year forecasts in Chart 7 for the ASX 200. In each case, the diagonal solid line is the 'best' forecast we can make – being the adjusted capital gains forecast discussed above. Each line starts on the January 1[st] value of the index and ends at December 31[st] (we also produce financial year forecasts). Of course we do not expect the market to move in a straight line.

All of these forecasts are as we actually made at the time and reported on each January 1[st]. They have not been edited or revised.

The box forecast methodology was not developed until January 1[st] 2014. The previous forecasts have been filled in using the same formulae as we used going forward from January 1[st] 2014. There was no back-testing to improve forecast accuracy of the highs and lows. These dotted lines do not depend upon any data in the future – i.e. after each January 1[st]. They are not like those used in so-called technical analysis where certain data features are joined together.

Chart 7: ASX 200 forecast boxes

Source: Thomson Reuters Datastream; and Woodhall Investment Research; data to close 18[th] May 2018

The two horizontal dotted lines in each year represent our forecast of the high and the low within each year. They are calculated from simulating possible future paths for the index using the adjusted capital gains forecasts and our estimate of volatility at the first day of the financial year.

The red lines for 2016 show that the e-o-y forecast was almost spot-on and the lower bound was breached before bouncing back.

For 2017 and 2018, the green marker was a published consensus forecast with the diamonds being the panel's high and low.

So far in 2018, the low dotted yellow line was a good predictor of a low and the e-o-y forecast appears to be on track.

It is not normally to be expected that both the high and the low are good forecasts in the same year. They are both equally likely at the beginning of the year but if, say, the low were to be reached halfway through the year, there is then much further to go to the high with only half of the time to get there! The high and low forecasts are not independent.

We show a similar box chart in Chart 8 for the S&P 500. Because of the common knock-on effect from

Wall Street to the ASX 200, we watch this chart for the ASX 200 as well as for our US portfolios.

Chart 8: S&P 500 forecast boxes

Source: Thomson Reuters Datastream; and Woodhall Investment Research; data to close 18th May 2018

One use for our box forecasts is as a sanity check on market movements. While the index stays within the box, all is consistent with our forecast views. Movements outside the box might instigate some action on our part.

In some strategies, we use the dotted lines as buy-sell signals.

While we do not update the forecasts in Charts 7 and 8, we do update the e-o-y and buy-sell signals on a daily basis. These updates are reported in our Weekly but we only use them to check whether or not the original annual forecasts are 'out of control'.

Medium-term forecasting

When a market is significantly mispriced, the task of forecasting out for a period of less than a few months is problematic as we do not know how quickly the mispricing will be resolved.

We believe that it is usually reasonable to assume that mispricing will be worked out over a 6 – 12 month period.

An alternative to our procedure would be to assume mispricing is eroded almost immediately and the forecast becomes the unadjusted capital gains forecast from 'fair value'. The outcome for 12-months out would be the same as we compute.

Short-run forecasts would be very different. However, an assumption of an almost immediate erosion of mispricing is not sensible as a rule.

Since we do not know what might return the market to a 'common sense' fair value, we think it is best not to make end-of-year predictions (calendar or fiscal) in the last three months of a given year. What we do in practice is to switch from calendar to fiscal year and vice versa as appropriate in our box forecasting – but it doesn't hurt to look at the full year in hindsight.

Of course when mispricing is insignificant, these issues are immaterial.

Conclusions

We believe that we have a complete, logical system for navigating the ASX 200 and its sectors. We have a replica system for the S&P 500. None of these methods has been altered since 2010 but the box forecasts were added to the suite n 2014.

Each week we post a report on our website: www.woodhall.com.au. This report contains a commentary on six charts up to the Friday close and comments on major events during the previous week. We also provide a narrative about how we see the future is looking relative to our forecasts.

The ASX 200 papers are posted on Saturday mornings – usually between 10am and 11am because the data supplier does not update its data until around 8am (when Wall Street closes in AEDT). US data for the Friday close is not available until Saturday afternoon, Sydney time.

We also produce a monthly 'Woodhall Wrap' on the first of each month. This paper is more of a discursive, macro, consolidated view of the world as we – and others – see it.