
Extracting market expectations from option prices

Aron Gereben, Financial Markets Department¹

Indicators of market expectations based on option prices are gaining popularity among central banks. The Reserve Bank recently began to use these indicators in financial stability and monetary policy analysis. This article provides a non-technical overview of these techniques and highlights how they might be used, through examples.²

1 Introduction

In most aspects of life, people – consciously or unconsciously – hold expectations about the future. They collect relevant, currently available, information and use it to help form their view about future events. Expectations play an important role in our lives and in the decisions we make.

Expectations about financial market prices are particularly important for market participants – both investors and borrowers. Investors in these markets have to decide what to buy and what to sell, when, and at what price. Investors' decisions are based on their views about what various asset prices will be in the future. If the investor's expectations prove to be right, they make a profit, whereas if they prove to be wrong, the result is a loss. Therefore, active market participants put a lot of effort into gathering information about the future to produce the best forecasts possible.

Most of the techniques that reveal the expectations of market participants focus on the average of the expectations, and therefore collapse a range of potentially different and individually uncertain views about the future to a single figure. This approach is useful to track the changes in the direction of market sentiments. However, in some cases we are also interested in the uncertainty and the dispersion of individual views surrounding the central view.

This article discusses techniques that use information in the prices of financial options to reveal the nature of the uncertainty that surrounds market expectations about financial market indicators such as stock indices and the

exchange rate. First, we introduce concepts such as probability density, standard deviation, skewness, and kurtosis that allow us to quantify and analyse uncertainty. This is followed by a short discussion of options in section 3. Sections 4 and 5 present techniques for measuring market uncertainty: implied volatility and implied probability density functions, respectively. These techniques are illustrated with examples of recent episodes of financial market uncertainty. Finally, we highlight the practical application of these techniques for the Reserve Bank.

2 Expectations and uncertainty

Information about market expectations of financial prices, such as interest rates and exchange rates, can be useful for the Reserve Bank. What markets expect affects the way they respond to the Bank's policy actions. Furthermore, quantifying expectations is useful for comparing the Bank's own forecasts with the market view and in assessing the stability of the financial system – these issues are discussed further in section 6.³

In quantifying expectations about how a financial asset's price will behave, however, there is a complication, as we have to take into account the uncertainty surrounding these expectations. The uncertainty of market expectations comes from two sources. First, different market participants will have different opinions about the future. When we talk about expectations, we generally refer to the average, or mean

¹ The author would like to thank Geof Mortlock, Michael Reddell and Ian Woolford for their comments on earlier drafts.

² For a more detailed and technical discussion of the methods presented here see Gereben (2002).

³ Gordon and Krippner (2001) gives a thorough overview of the use of expectation-based measures in monetary policy, and the impact of expectations on the effectiveness of monetary policy actions.

value, of a number of individual views – underlying the average is a range of potentially disparate views. Secondly, individuals themselves have a sense of how uncertain the future price of the given asset is. An individual may have a “central forecast”, but personal views about the future could probably be better described as if the person were assigning probabilities to a number of different potential outcomes: high probabilities are assigned to the most likely results and lower ones to the less likely cases. This uncertainty at an individual level contributes to the overall dispersion of expectations in the market as a whole.

Although the average (or mean) of market expectations is valuable information in itself, it is often useful to quantify the underlying uncertainty as well – are conditions ‘normal’, or particularly uncertain?

When facing an uncertain event in the future, people generally have views about the possible outcomes. They also have views about the probability of all the different possible outcomes. The different outcomes and the probabilities assigned to them together can be considered as a *probability distribution*.

Figure 1
Probability density function (PDF)





Figure 1 is a graphical representation of expectations about a hypothetical asset's price. On the horizontal axis the potential outcomes – in this case, potential changes in the price – are plotted. The vertical axis shows the probability of the outcomes. Loosely speaking, the value y measured on the vertical axis at point x means that there is y per cent probability of the asset price being in a 1 percentage point wide area around x . This mathematical concept, which is often used to visualise a probability distribution, is called a *probability density function* (PDF). Loosely speaking, a probability density function is a mathematical tool that gives us the expected likelihood of each possible future outcome. The area measured below the density function between two points is equal to the probability of the asset price being between the two points. To illustrate, let us assume that we want to calculate the probability of a price increase between 10 and 20 per cent. The answer is given by the size of the area below the density function between the 10 per cent and the 20 per cent points as shown by the shaded area in figure 1.

The shape of the curve in figure 1 tells us about the nature of the uncertainty that surrounds expectations about the asset's price. For example, the wider the density function, the higher the uncertainty, and conversely, a narrower probability density function means that expectations are more tightly bunched around the most likely event. Economists and statisticians have precise definitions and tools to quantify the shape of PDFs. The three most widely used concepts for characterising the shape of probability density functions are *volatility*, *skewness*, and *kurtosis*.

Volatility reflects the width of a density function. The first panel of figure 2 shows the difference between two density functions with different volatilities. Lower volatility (shown by the grey line) means that the final outcome of the financial

Figure 2
Shape characteristics of a PDF



asset's price is expected by most observers to be relatively closer to the mean of the expectations. Higher volatility (shown by the black line) on the other hand indicates that the chance of a deviation from the mean is relatively higher. Volatility can be thought of as a general measure of uncertainty about the final outcome. The most commonly used measure of volatility is the *standard deviation*.

Probability density functions are not necessarily symmetric: for example, there may be cases when individuals assign higher chances to large price changes in one direction than in the other. The asymmetry of a PDF is called "skewness". The second panel of figure 2 displays a (negatively) skewed density function relative to a symmetric one. The negative skewness on the chart indicates that in this case large price falls are perceived to be much more likely than large price increases. Note that on the right side of the panel the probability – the area under the PDF – is concentrated close to the centre of the distribution, whereas on the left side of the panel the probability of outcomes away from the centre is not negligible. The most commonly used measure is called the *skewness coefficient*.

Another shape characteristic of a PDF is called *kurtosis*. Kurtosis is an indicator of the likelihood of large deviations from the mean of the expectations. Unlike volatility, which indicates the average size of potential deviations from the mean, kurtosis indicates the extent to which deviations themselves can be different from each other. A density function with a high degree of kurtosis indicates that although in most cases the deviation is likely to be small, there is a perception of a non-negligible chance of extremely high or extremely low outcomes. Panel 3 of figure 2 illustrates the difference between distributions with low kurtosis (shown by the grey line) and high kurtosis (shown by the black line). The most widespread indicator is called the *coefficient of kurtosis*.

The coefficients of standard deviation, skewness, and kurtosis are generally expressed as percentages. The standard deviation and kurtosis coefficients take only positive values, while the coefficient of skewness can be either positive –

indicating larger probability at the right tail of the density function, or negative, indicating larger probability at the left tail of the density function. To give a precise definition of the coefficients of standard deviation, skewness, and kurtosis goes beyond the scope of this article.⁴ However, the precise interpretation of their numerical scale is not relevant from the point of view of our analysis. What is more important for us is the *dynamics* of these indicators – eg how these indicators evolve over time and how their current levels compare to their values in the past.

Now we have a toolkit that allows us to analyse the distribution of market expectations. The challenge is to obtain the expectations themselves. This is where options are particularly useful.

3 What is an option?

Options are one form of financial derivatives. They are 'derivative' in the sense that their final payoff depends on the price of some other product.⁵ This other product is often called the option's 'underlying asset'. Options can be written on a large number of assets, such as agricultural commodities, shares, bonds, foreign currency, and so on.

As a central bank, the Reserve Bank is particularly interested in following the evolution of the New Zealand interest rate and currency options markets. However, other option markets, such as equity or commodity options can also provide useful information about the domestic economy or the international environment.

Due to their special features, options are particularly useful for recovering the market participants' view about the uncertainty that surrounds the price of the asset on which the option is written. As the option's value is heavily dependent on the distribution of the probabilities of various possible outcomes, the market participants reveal their view about uncertainty through the price at which various options are traded. Options are essentially like insurance products, and we are looking at the relative prices of various different

⁴ An introduction into probability density functions and other statistical concepts mentioned in this article can be found in Ramanathan (1993).

⁵ For a detailed explanation of derivative securities, and of the way in "hedgers", "speculators" and "traders" participate in markets for these products see Hawkesby (1999).

insurance policies. To get a broader insight into this process, we have to take a closer look at option contracts.

The two basic types of options are "calls" and "puts". They are contracts between two parties, the option's holder and the option's writer. In a *call* option contract, the holder pays a fee to the writer to obtain the *right (but not the obligation) to buy* a certain amount of the underlying asset at a pre-specified day (called the expiration date or maturity) and at a pre-specified price (called the strike or exercise price). On the other hand, the writer of the option is obliged to sell – on the request of the holder – the underlying asset at the pre-specified strike price when the option expires.

A *put* option, on the other hand, gives its holder the *right (but not the obligation) to sell* a certain amount of the underlying asset at a pre-specified day and at a pre-specified price. On the other hand, the writer is obliged to buy the underlying asset at maturity on request. Similarly to the call option, the holder pays an option fee to the writer.⁶


To get a better understanding of options it is useful to look at their payoff to the holder at maturity. Figure 3 shows the payoffs and the final profits – which are equal to the payoff minus the option's price – of call and put options at maturity the underlying asset's final price changes. As noted above, a call option gives the right to buy an asset at a pre-specified strike price. If the price of the underlying asset at maturity is

lower than the strike price, it is not profitable to exploit this right by exercising the option: therefore the option is worth nothing. On the other hand, if the price at maturity is higher than the strike, the option holder can exercise the option, and gain the difference between the market price and the strike price.

Similarly, as noted above, a put option confers the right to sell an asset at a pre-specified price. If the market price of the underlying asset is higher than the strike price, it does not make sense to sell at a lower price; thus the option is worthless. However, if the market price is lower than the strike price it is worth exercising the option. Here, the option holder's payoff is the difference between the two prices.

A closer look at figure 3 gives us an insight into the relationship between the option's value and the price uncertainty of the underlying asset. The potential loss to the holder of the option is bounded: the option's final payoff is never worth less than zero. If we take into account that the option's holder has to pay a fee – the option's price – to enter into the option contract, the total profit of the option can be negative (see the grey line on figure 3). However, it is never less than the initial price of the option, the upfront "insurance premium". On the other hand, the potential gain is practically unbounded. The higher is the underlying asset's price at maturity, the higher is the gain from holding

Figure 3
The payoff and profit diagrams of call and put options



⁶ There are several different types of options. The ones described here are called European-style options. Another popular class of options, called American-style options differs from the ones described here in that they can be exercised not only on a single day, but at any time between the purchase and the maturity date.

a call option. Similarly, the lower is the asset's price the higher is the gain from holding a put option.⁷ If the uncertainty surrounding the price in the future is higher, so too is the chance that the price will deviate significantly from the strike price from the option. Under these conditions, the chance for a large profit is higher and therefore the buyer of the option is willing to pay a higher amount to enter into a contract. From the viewpoint of the option's writer, higher uncertainty means higher probability of the option being executed, in which case the option's writer would make a loss. Thus to compensate for this risk, the writer would require a higher initial option fee to enter into a contract. All in all, there exists a positive relationship between the uncertainty (or volatility) of the underlying asset's price and the price of an option written on the asset.

4 Implied volatility

As we have seen, option prices appear to be closely related to the market's expectation about the uncertainty that surrounds the price of the underlying asset. Because of this relationship, we can use option prices to extract information about market expectations. Several different methods for unfolding this information from option prices exist.

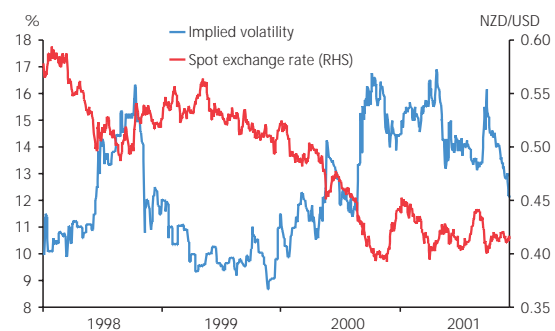
The simplest method is to calculate the so-called *Black-Scholes implied volatility* measure. The Black-Scholes formula is a pricing tool that gives an approximate fix on the option's price in terms of several parameters: the current price of the underlying asset, the strike price, the interest rate of a risk-free investment, the time until maturity, and the standard deviation of the underlying asset's price.⁸ This last parameter is generally referred to simply as volatility or 'vol'. One use of the Black-Scholes formula is to calculate option prices from given parameters. However, unlike the other four that are precisely given at any time, the volatility parameter is not directly observable, as it is related to the future uncertainty of the underlying asset's price. Thus we can only rely on estimates of it.

Another way of using the formula is to 'reverse engineer' the volatility parameter from existing option price data. To

do this, one has to find the level of the volatility parameter that makes the option's price calculated by the formula equal to the option price observed in the market. The resulting measure of uncertainty is called implied volatility, which can be considered as a proxy for the market's expectation of the volatility of future price changes. It is usually expressed in annual terms. Measures of implied volatility can be calculated for all financial assets, provided that options written on them are regularly traded. Therefore, the prices of these options are available and observable on a regular basis.

Given that there exists a clear-cut relationship between the implied volatility and the option's price, on some option markets the quotes are given directly in terms of implied volatilities rather than in terms of prices expressed in currency units. On these markets the implied volatility is directly observable through the quotes of option writers and option buyers.

Figure 4
Implied volatility of the NZD/USD exchange rate




Source: UBS Warburg

To illustrate, figure 4 displays the implied volatility of the NZD / USD exchange rate over the last four years, together with the spot exchange rate. The data are over-the-counter currency option quotes provided by UBS Warburg. The level of implied volatility varied significantly over the period. At first sight, there is a noticeable co-movement between the exchange rate and the implied volatility. Periods characterised with a weak exchange rate seem to go together with a perception of higher volatility, therefore higher risk. Other changes in the implied volatility can be associated with domestic and international events. For example, we can identify the effects of the Asian and Russian crises that caused significant fluctuations in the implied volatility in the second

⁷ In case of a put option, the profit is not unlimited, as the underlying asset's price cannot go to below zero.

⁸ Black and Scholes (1973).

Figure 5
Implied volatility of the S&P 100 and Nasdaq
stock indices



Source: Chicago Board Options Exchange

half of 1998. An upsurge in negative sentiment about the New Zealand economy prompted a sharp rise in the perceived uncertainty and fall in the spot exchange rate in September 2000. Finally, the terrorist attack in the United States caused another temporary increase in the implied volatility in September 2001.

Figure 5 provides another example, this time of the implied volatility of two stock indices, the S&P 100 and the Nasdaq. The S&P 100 index tracks the price of the largest and most actively traded stocks in the United States. The Nasdaq has a large proportion of IT and telecommunications-oriented companies in it. Options written on these indices are actively traded on the Chicago Board Options Exchange (CBOE). The CBOE publishes implied volatility indices for both the S&P 100 and the Nasdaq. As with the NZD/USD spot rate, the Russian crisis caused a significant increase in market uncertainty. Similarly, the burst of the equity price bubble of the IT (or 'tech') sector starting around the first half of 2000 caused a large increase in the gap between the implied volatility of the Nasdaq and the S&P 100 indices. According to the market's perception, the technology-oriented Nasdaq index (the 'new economy' sector) became increasingly risky compared to the S&P 100, which represented the more "traditional" sectors of the economy. Finally, a spike in the implied volatility of both indices could be observed after the terrorist attacks against the US on 11 September 2001.

5 Implied probability distributions

Implied volatility indicates the range of uncertainty surrounding market expectations about financial prices. However, implied volatility in itself does not allow us to analyse fully the general shape and the nature of the distribution of these expectations. To obtain a more precise characterisation of the shape, more advanced methods are needed.

A widely used method, originally developed by Malz (1998), allows us to extract a probability density function that describes the structure of the uncertainty surrounding market expectations (see the box for more details on the Malz method). After deriving the density functions, we can calculate the measures that characterise its shape: the standard deviation, the skewness, and the kurtosis.

The Bank uses this method to derive the density functions for the NZD/USD exchange rate. We use a daily data-set of currency option prices, covering the last two years. UBS Warburg provided the data on option quotes. As price data for options with different maturities is available, it is possible to calculate the expectations for 1 month, 3 months, 6 months, and one year ahead.

As the market for options written on the NZ dollar is relatively small by international standards, there is a risk that some of the option contracts that were used in the process of estimating the implied PDFs are not traded on a regular basis. If there is no actual trade behind the quoted prices, there is no guarantee that those prices represent the market participants' expectations.

The analysis of our data-set reveals that at the first half of the period covered there were some short periods characterised with only a low level of market liquidity. Despite these signs, however, the information derived from options seems to match other data sources on expectations and market sentiment. This suggests that the level of market liquidity, while not always high, is sufficient for fluctuations in quoted option prices to give us meaningful information on changes in the distribution of expectations.

Box: Malz's method for extracting implied density functions from currency options

Malz (1998) developed a method to derive the uncertainty surrounding the market expectations about exchange rates using the prices of currency options. The global currency option market is a highly liquid marketplace of options written on different currencies. Furthermore, the standard contract types of this market are remarkably well suited for deriving implied probability density functions.

One particular feature of the currency option market is that the majority of trades involve standard option combinations. Option combinations are two or more options traded together as a single product. The rationale behind the combinations is to create financial products with special payoff structures that match the market participants' needs.

The most popular option combinations traded on the market are *straddles*, *risk reversals*, and *strangles*.

A straddle is a simultaneous purchase of a put and a call option with the same strike prices. The first panel of figure B1 shows the payoff function of the straddle. As the chart displays, the larger the deviation from the strike price, the higher the payoff for the straddle's owner. This means that the straddle becomes more lucrative in times of high

volatility. As a consequence, the market price of a straddle contract reflects the currency option market's guess about the future volatility of the exchange rate.

Entering into a risk reversal contract means buying a call option with a high exercise price and selling a put option with a low exercise price. Its payoff is displayed on the second panel. The structure of the contract is such that if risks are symmetric, the price of the call is equal to the price of the put option, and the price of the risk reversal is therefore zero. In the case of asymmetric expectations, however, the risk reversal price may be either negative or positive. Negative risk reversal prices indicate negative skewness, whereas positive risk reversal prices indicate positive skewness.

A strangle is a combination of a buying a call option with a high strike price and a put option with a low strike price. As can be seen on the third panel, the strangle turns profitable in the case of highly positive or highly negative price changes in the underlying asset. Thus, if the price of a strangle contract is high, the market assigns relatively high probability to these extreme outcomes. Strangle prices can therefore be used as an indicator of kurtosis.

Quotes on straddles, risk reversals and strangles can be obtained on a daily basis. Malz's method uses the daily observations of these instruments to deduct the shape of the probability density functions, which in turn describe market expectations, using mathematical techniques.

Figure B1
Option combinations






Figure 6
Uncertainty measures of the NZD/USD
exchange rate
(calculated from 3-month NZD/USD options)


Spot rate



Standard deviation



Skewness



Kurtosis




Figure 6 displays the path of the NZD/USD spot exchange rate together with the standard deviation, skewness, and kurtosis measures calculated from the implied probability density functions. Two significant episodes in the exchange rate history of the last two years – the sharp decline of the New Zealand dollar over August and September 2000, and the September 11 attack on the World Trade Center and the Pentagon – are highlighted.⁹ These episodes are well reflected in the standard deviation, skewness, and kurtosis statistics. Over the first period (August-September 2000) implied volatility and kurtosis were strongly increasing, indicating an increase in uncertainty in general and a higher perceived chance of extreme price movements in particular. Also, the skewness measure went deeply negative, indicating ‘bearish’ expectations on the exchange rate, ie the market was much more worried about further significant price falls than price increases.

The second highlighted episode is the aftermath of the terrorist attack against the World Trade Center, which showed similar, although rather more temporary fluctuations in the standard deviation, skewness, and kurtosis measures. The analysis in the rest of this section focuses on this event, given that the impact on market uncertainty was dramatic.

Figure 7 shows the shape of the options-implied density functions of the NZD/USD before and after the 11 September attack. The distributions were calculated from options with a one month maturity. It can be observed that the market turmoil caused by the attacks resulted in a more dispersed distribution of the one month-ahead expected returns. Moreover, the increase in the distribution was more marked at the lower tail, implying an increase in the skewness that suggests expectations being biased towards greater downside risk for the NZD/USD cross rate.

⁹ Movements leading up to the May 2000 *Monetary Policy Statement* (MPS) are also interesting. The sharp fall in the skewness indicator over the early months of 2000 highlighted the vulnerability of the market to a marked fall in the exchange rate, for which the widely expected rise in the OCR at the May MPS appears to have been the trigger.

Figure 7
Probability density functions for the NZD/USD 1 month ahead (expected percentage change relative to the spot rate)




Figure 8
Standard deviation and skewness around 11 September (calculated from options with a maturity of three months)

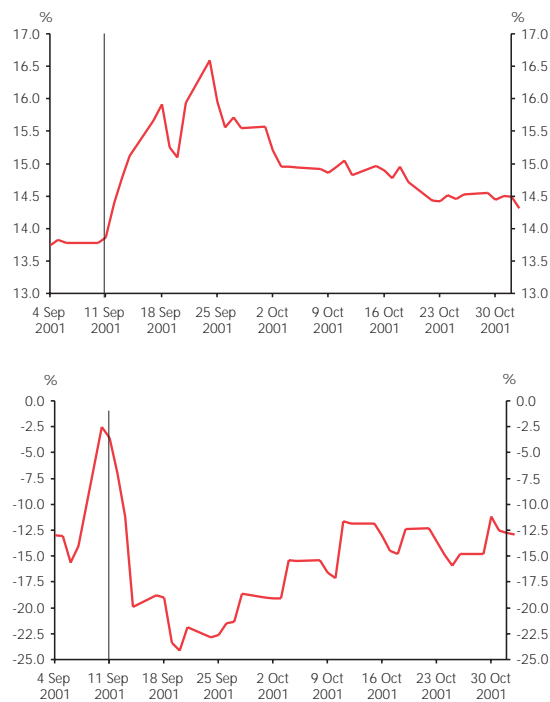


Figure 8 displays the evolution of the standard deviation and the skewness measure of the implied probability distributions in the aftermath of the crisis. Although initially the standard deviation increased quickly and the skewness turned highly negative (indicating depreciation expectations) the market view began to retrace from the extreme uncertainty recorded to more 'normal' levels by the last week of September. By the end of October both the standard deviation and the skewness had returned to around their pre-crisis levels.

A very similar pattern could be observed in the expectations surrounding the exchange rate of the Australian dollar. This suggests that the observed fluctuations in the options-based indicators were indeed due to the international financial turmoil, rather than being New Zealand specific.

6 How can the Reserve Bank use the information derived from option prices?

Option-based information about market uncertainty can be useful for a central bank in several ways.¹⁰ However, as with many other statistical or econometric methods used in economic analysis, these indicators come with a health warning. Data quality, fluctuations in market liquidity – especially on small markets such as the New Zealand option market – and the assumptions behind the techniques may influence the results obtained from these measures.¹¹ Options-based indicators should therefore always be used in conjunction with other information sources (such as other financial prices, market surveys and opinions, discussions with market participants etc) that can act as a crosscheck on the results.

The Reserve Bank has recently started to analyse and use options-implied probability indicators. These indicators are included in our standard set of macro-prudential indicators that are monitored by the Bank's Financial System Oversight Committee and which will be discussed in the annual macro-

¹⁰ For applications of option-implied probability measures in monetary policy analysis see for example Bahra (1997), Aguilar and Hördahl (1999), Eitrheim, Frøyland, and Røisland (1999) or Bank for International Settlements (1999). Examples of use of implied volatility indicators and options-based probability density functions as indicators of financial stability can be found in Bank of England (2001) and IMF (2002).

¹¹ One simplifying technical assumption used in the majority of the techniques that use derivative prices to infer market expectations is that market participants are risk neutral, ie they do not require a yield premium for risky investments. This assumption is not intuitively obvious and is generally rejected empirically. Although the risk neutrality biases somewhat the results that we obtain from options-based indicators, this bias is generally found to influence the shape of the obtained probability density functions only negligibly (see Rubinstein (1994)).

financial stability article in the *Bulletin*.¹² They also form part of the analysis of financial market conditions that help inform our monetary policy deliberations.

Options-based indicators can enrich our understanding of market sentiment. These methods focus on the uncertainty, and heightened uncertainty might be a warning sign of potential vulnerabilities. Together with other market based indicators such as different credit, bond and swap spreads, and equity price indices, options-related indicators are particularly useful in the Reserve Bank's monitoring of the stability of the financial system and complement other macro-prudential data sources such as macroeconomic and banking system data. The main advantage of indicators based on market prices is their forward-looking nature and their availability on a daily basis. Also, we can follow both domestic and overseas option markets simultaneously, which can help us to separate global market shocks from domestic disturbances.

Besides the analysis of the stability of the financial system, there are other areas where the Reserve Bank can make use of the information derived from option prices. For example, interest rate options can provide information about the distribution of interest rate expectations before OCR decisions. Market-based measures of expectations can also indicate whether the Reserve Bank's monetary policy measures influenced the market in the intended way. For example, did market expectations change in the desired direction after a change in monetary policy settings? Option-based indicators can also show whether central bank measures reduced or increased the uncertainty in the market. They can also explain the behaviour of other financial variables, such as risk premia.

References

Aguilar, J and P Hördahl (1999), "Option prices and market expectations", *Sveriges Riksbank Quarterly Review*, 1/1999, 43-70.

¹² A detailed description of the Reserve Bank's approach to financial stability indicators and of the role of the Financial System Oversight Committee can be found in Woolford (2001).

Bahra B (1997), "Implied risk-neutral probability density functions from option prices: theory and application", *Bank of England Working Paper*, 66.

Bank for International Settlements (1999), *Proceedings of the workshop 'Estimating and interpreting probability density functions'* held at the BIS on 14 June 1999.

Bank of England (2001), 'The financial stability conjuncture and outlook', *Financial Stability Review*, December 2001.

Black, F and M Scholes (1973), "The Pricing of Options and Corporate Liabilities," *Journal of Political Economy*, 81, p 637-59.

Eitheim Ø, E Frøyland and Ø Røisland (1999), "Can the price of currency options provide an indication of market perceptions of the uncertainty attached to the krone exchange rate?" *Norges Bank Economic Bulletin*, 3/99, 266-77.

Gereben, Á (2002), 'Extracting market expectations from option prices: an application to over-the-counter New Zealand dollar options', *Reserve Bank of New Zealand Discussion Paper* (forthcoming).

Gordon, M and L Krippner (2001), 'Market expectations of the Official Cash Rate', *Reserve Bank of New Zealand Bulletin*, 64, 2, June, pp 25-39.

Hawkesby, C (1999), 'A primer on derivatives markets', *Reserve Bank of New Zealand Bulletin*, 62, 2, June, pp 24-43.

IMF (2002), *Global Financial Stability Report*, (forthcoming).

Malz, A M (1997) "Option-implied probability distributions and currency excess returns", *Federal Reserve Bank of New York Staff Report* 32, November 1997.

Ramanathan, R (1993), *Statistical Methods in Econometrics*, Academic Press.

Rubinstein, M (1994), "Implied Binomial Trees", *Journal of Finance*, LXIX/3, 771-818.

Woolford, I (2001), 'Macro-financial stability and macroprudential analysis', *Reserve Bank of New Zealand Bulletin*, 64, 3, September, pp 29-43.