
Life Insurance Practice Committee

Technical Paper: The Development and Use of Volatility Assumptions

January 2023

Contents

1. Status of Technical Paper	3
2. Background	3
2.1 What is volatility	3
2.2 Determining volatility assumptions	4
2.3 Cautionary observations	4
2.4 Need for volatility assumptions	5
2.4.1 Valuing asymmetric pay-offs	5
2.4.2 Capital reserving models	6
2.4.3 Setting asset allocations	8
2.5 Accounting standard guidance on setting volatility assumptions	8
2.6 Actuarial standards and guidance on setting volatility assumptions	10
3. Modelling considerations	10
3.1 Realised volatility	10
3.2 Implied volatility	11
3.2.1 Interpreting implied volatility	12
3.2.2 Observed term structure of implied volatility	13
3.3 Smiles and smirks	14
3.4 Implied volatility using call options and put options	16
4. Applications	16
4.1 Setting a volatility assumption for a parametric option pricing model	16
4.2 Setting a volatility assumption for stochastic models	18
5. Practical considerations	19
5.1 Volatility surface	19
5.2 Monitoring	20
5.3 Obtaining market prices	20
5.4 Uncertainty	21

5.5 Other considerations	21
5.6 Communication	21
6. Conclusion	22
7. References	22
Appendix A – Black-Scholes	24
Appendix B – Modelling realised volatility in more detail	25
Appendix C – Asset returns	27

1. Status of Technical Paper

This Technical Paper was prepared by the Life Insurance Practice Committee ("LIPC") of the Institute of Actuaries of Australia ("Institute").

It does not represent a Professional Standard or Practice Guideline of the Institute and it has been prepared for the purpose of informing members of issues in developing and using volatility assumptions for different purposes, and with a view to developing a more complete understanding of the considerations involved.

This Technical Paper does not constitute legal advice. Any interpretation or commentary within the Technical Paper regarding specific legislative or regulatory requirements reflects the expectations of the Institute but does not guarantee compliance under applicable legislation or regulations.

Accordingly, Members should seek clarification from the relevant regulator and/or seek legal advice in the event they are unsure or require specific guidance regarding their legal or regulatory obligations.

This IN does not override the requirements in this Professional Practice Document or in any other Professional Standards or Practice Guidelines that are relevant to this area of work.

This is the second version of this Technical Paper, with references to AASB 1038 Life Insurance Contracts replaced with AASB 17 Insurance Contracts (applying to annual periods beginning on or after 1 January 2023). A full review of the developing practice has not been completed since the original Technical Paper was drafted in 2011. Consequently, practices may have developed since that time.

Feedback or comments on the Technical Paper can be submitted to LIPC at the following address: ppd@actuaries.asn.au or directly to the current Chairperson and Secretary using the contact details on the Institute's [Committee webpage](#).

2. Background

2.1 What is volatility

Volatility is a common measure of risk. It is most often used to describe the risk associated with equity investments or equity indices and is calculated, for this purpose, as the annualised standard deviation of continuously compounded returns.

If volatility is defined to mean the standard deviation of any random variable, volatility can be calculated for economic variables (for example, interest rates and inflation) and

demographic variables (for example, mortality, morbidity and lapses). For the purpose of this Technical Paper, the focus is mainly on equity volatility.

2.2 Determining volatility assumptions

There are two common methods of determining volatility assumptions. The first approach involves determining implied volatility. Implied (or imputed) volatility is derived from the market prices of options using an option pricing model or formula. In other words, it is the volatility that, when inputted into an option pricing model, produces an option value equal to the current market price of that option.

The second approach involves calculating historical realised volatility. For equity volatility this would involve calculating the standard deviation of historical log equity returns.

2.3 Cautionary observations

This Technical Paper is intended to be an introduction to the key theoretical modelling and practical considerations involved in setting volatility assumptions.

Volatility modelling is particularly complex and it is easy to produce misleading results by incorporating spurious relationships that are unlikely to be reproduced, or to miss important relationships that may otherwise assert themselves at inconvenient times. As a consequence, due care should be taken when setting volatility assumptions.

Moreover, volatility itself is only a complete measure of risk for a narrow range of statistical distributions (for example, normal and lognormal distributions). The tractability of these distributions is so powerful that many financial models assume normality in situations where this assumption cannot be justified. The focus on volatility in this Technical Paper may implicitly encourage the use of inappropriate models. Members are encouraged to address this issue in the first instance. If the underlying distributions cannot be adequately estimated with mean and variance alone, then more complex models or distributions should be used (for example, using economic models). This Technical Paper occasionally touches on this issue, but its focus is on estimating volatility in cases where the underlying models are broadly appropriate to the situation at hand. Derman and Wilmot (2009) provide a salutary warning:

"Building financial models is challenging and worthwhile: you need to combine the qualitative and the quantitative, imagination and observation, art and science, all in the service of finding approximate patterns in the behavior of markets and securities. The greatest danger is the age-old sin of idolatry. Financial markets are alive but a model, however beautiful, is an artifice. No matter how hard you try, you will not be able to breathe life into it. To confuse the model with the world is to embrace a future disaster driven by the belief that humans obey mathematical rules."

2.4 Need for volatility assumptions

For actuaries practising in life insurance and wealth management, volatility assumptions are generally required for a range of purposes including:

- ▶ valuing asymmetric pay-offs;
- ▶ capital reserving models;
- ▶ setting asset allocations; and
- ▶ dynamic hedging.

This Technical Paper focuses on setting volatility assumptions for the purposes of valuing asymmetric pay-offs and setting capital reserves. However, a number of the considerations outlined in this Technical Paper also apply to other situations where volatility assumptions are required.

2.4.1 Valuing asymmetric pay-offs

Volatility assumptions are typically required when valuing obligations or assets with asymmetric pay-offs. This reflects that, when valuing asymmetric risks, the distribution of potential pay-off outcomes needs to be allowed for in the valuation calculations.

The LIPC Technical Paper regarding asymmetric risks (the “Asymmetric Risk Technical Paper”) outlines in Section 3.1 and Appendix 1 some examples of asymmetric risks associated with life insurance.

The Asymmetric Risk Technical Paper states that the following techniques could be used to determine the value of obligations/assets with asymmetric pay-offs:

- ▶ replicating portfolios. This approach is commonly adopted when it is easier to value the assets in the replicating portfolio than it is to value the obligation/asset. It involves choosing a portfolio of assets (fixed interest bonds, options etc) whose cash flows are identical to the magnitude and the timing of the cash flows of the obligation/asset being valued. Using arbitrage arguments, the value of the obligation/asset is equal to the value of the replicating portfolio;
- ▶ stochastic models (risk neutral and real world); or
- ▶ combinations of the above.

Volatility assumptions may be unnecessary if a replicating portfolio technique is adopted, and the market prices of each asset in the replicating portfolio is known (for example, as each asset is traded on an exchange).

In general stochastic models, however, require a volatility assumption, of which the best known is the Black-Scholes option pricing formula set out in Appendix A.

Asymmetric risks may need to be valued for a number of purposes, including:

- ▶ financial reporting purposes - a number of accounting standards refer to the valuation of asymmetric risks. Section 2.5 below outlines the key accounting standards referring to asymmetric risks and their requirements regarding the setting of volatility;
- ▶ pricing for products that contain asymmetric pay-offs (for example, investment account contracts);
- ▶ valuation of employer share schemes; and
- ▶ economic valuations, for example:
 - the valuation of real options. A real option is the right but not the obligation to undertake a specific business decision that has asymmetric pay-offs. For example, the right but not the obligation to develop a mining site is a real option; and
 - the valuation of life insurers using Market Consistent Embedded Value ("MCEV") techniques. MCEV may involve the valuation of options embedded in life insurance contracts and using option pricing techniques to value the limited liability put option associated with a limited liability company structure.

Volatility estimates used for valuing asymmetric risks should generally produce market consistent values. Implied volatilities will do so, as they are derived from market price of options. However, as reflected in the middle column of the table in section 4.1, some adjustments to implied volatilities may be required before they can be used to value more complicated asymmetric risks.

Historical realised volatility can also be used as the starting point for determining market consistent volatility assumptions. The final column in the table of section 4.1 tabulates the adjustments that may need to be made to historical realised volatility in order to determine a market consistent volatility assumption.

2.4.2 Capital reserving models

Economic capital or target surplus models ("Capital Models") that use parametric models¹ to model asset and liability risks (for example, lapses, expenses, mortality and morbidity) may require explicit volatility assumptions. Examples are:

¹ Parametric models are models based on mathematical equations.

- ▶ parametric value-at-risk models ("VAR") that assume asset values follow Brownian motion² or geometric Brownian motion; and
- ▶ some complex parametric VAR models that assume asset returns exhibit fat tails.

Not all Capital Models that use parametric models to model asset and liability risks will require an explicit volatility assumption. Examples of parametric models where the volatility assumptions are implicit are:

- ▶ parametric models that assume lapse or mortality risk is explained by a Bernoulli distribution, where the volatility is inherent in the distribution; and
- ▶ some economic models used to model asset risk. A model that assumes share price returns are a function of economic random variables may not require a share price volatility estimate if the volatility arises from the movement of the explanatory variables.

Capital Models that use non-parametric models to model asset and liability risks generally will also not require an explicit volatility assumption. An example of a non-parametric Capital Model is a VAR model that simulates asset risks by randomly sampling from a historical share price return data set.

Volatility estimates used for capital reserving models should reflect a number of factors. These include:

- ▶ the actuary's expectations of future realised volatility over the time frame consistent with the capital reserving period. For example, if the capital reserve is designed to reserve for potential losses over the next 12 month period, then the volatility estimate should reflect best estimate expectations over the next twelve months. Implied volatilities calculated using options with a term to maturity commensurate with the capital reserving period can be an unbiased estimate of the market's view of expected future realised volatility over the capital reserving period if certain conditions apply. These conditions are considered in the middle column in the table in section 4.2. Historical realised volatility may also be an unbiased estimate of expected future volatility if certain assumptions apply. These conditions are considered in the final column in the table in section 4.2;
- ▶ approximations and limitations associated with the parametric model adopted. For example, if the parametric model assumes volatility is constant when, in fact, volatility changes over time and may mean revert or shift to a new regime after a period, it may

² Brownian Motion in its most basic form is a continuous distribution with initial value and mean value of 0. Each increment is independent of the history and follows a normal distribution as set out below:

$$W_0 = 0$$

$$W_t - W_s \sim N(0, t-s)$$

be appropriate to adopt a volatility assumption that is higher than the actuary's best estimate expectations of future realised volatility; and

- ▶ limitations and approximations associated with the data used to determine the volatility assumption. If there is limited data available to set the volatility assumption, it may be appropriate to adopt a volatility assumption that is higher than the actuary's best estimate expectations of future realised volatility. There are a variety of methods available to address the estimation error in the assumptions in order to improve the stability of the results. See, for instance, Michaud and Michaud (2008). Further details regarding determining volatility assumptions for Capital Models is outlined in section 4.2.

2.4.3 Setting asset allocations

Constructing optimal investment portfolios, or determining the asset allocation of a life insurer, involves trading-off expected returns against expected risks and may require the use of stochastic models that need volatility assumptions.

Simplistic mean/variance modelling depends on the volatility assumptions for each asset class and implicit or explicit assumptions as to the correlation between classes. Similar considerations apply to those mentioned in the previous sections.

Care should be exercised in using simplistic mean variance models for these purposes, especially if optimisation is undertaken. A typical feature of such modelling is that optimal portfolios are extremely sensitive to the assumptions used. The extent of this sensitivity should be assessed and measures taken to make the outcomes robust under a range of differing input assumptions. Some measures involve using Bayesian techniques to adjust volatility to reflect uncertainties in the assumptions, both means and variances.

Moreover, the non-normal return distributions (for example, leptokurtic / fat tailed) may make the use of simplistic mean variance approaches inappropriate. Alternative approaches using different parametric return distributions, or even non-parametric distributions, are outside of the scope of this Technical Paper.

2.5 Accounting standard guidance on setting volatility assumptions

A number of accounting standards make reference to the valuation of asymmetric risk, including:

- ▶ "AASB 2 Share-based Payment", which covers the valuation of executive share options;
- ▶ "AASB 17 Insurance Contracts" which covers the valuation of insurance contracts (including policyholder options); and

- ▶ “AASB 139 Financial Instruments: Recognition and Measurement” for the valuation of financial instruments not specifically covered in other standards (including AASB 1038 and AASB 2) as outlined in section 2 of AASB 139.

AASB 17 does not provide any specific guidance on determining volatility assumptions.

Appendix B (paragraph 25) of AASB 2 provides detailed guidance regarding the issues to consider in setting volatility assumptions for executive share options. Factors that should be considered include:

- ▶ implied volatility from traded share options or other traded instruments that have option features (generally there is limited market data available for liquid securities particularly after 1 year's duration);
- ▶ historical realised volatility over the most recent period consistent with the expected term of the option;
- ▶ the length of time an entity has been publicly listed. That is, a newly listed entity may be expected to have higher volatility of returns than an entity that has been listed for several years. If an entity has insufficient historical price information, then historical realised volatility of similar entities should be considered as well as the length of time from listing. For example, the historical volatility of a similar entity over the first five years after being listed should be considered when valuing an option with a five year expected life granted by a company that has just listed;
- ▶ factors that indicate that future volatility will differ from past volatility, for example potential mean reversion and other extraordinary factors affecting historic realised volatility; and
- ▶ for unlisted entities, the volatility could be estimated by analysing historical volatility for unlisted entities that have created an internal market for their shares. In the absence of an internal market, the methodology used to value the shares should be considered to set the volatility assumption. For example, an approach to valuing unlisted shares can be based on the share prices of similar listed entities. In this case, an analysis of the historical volatility of similar listed entities should be considered in setting the volatility assumption for the unlisted entity.

AASB 139 states that the volatility assumption can be derived based on historic market data and implied volatilities in current market prices. This is consistent with the more detailed guidance offered in AASB 2.

Neither AASB 17 nor the APRA prudential standards applying to life insurers provide any guidance on setting volatility assumptions.

2.6 Actuarial standards and guidance on setting volatility assumptions

The Institute's Guidance Note 510 (Valuation of Share Based Payments) details the items a Member should consider when setting assumptions (including volatility) under AASB 2. These items are consistent with the factors outlined in AASB 2, albeit not as detailed.

The Board of Actuarial Standards of the UK Financial Reporting Council has issued technical actuarial standards on data, modelling and reporting that address some of the issues that arise in developing and using volatility assumptions. In summary, they require:

- ▶ documentation of all data, the mathematical model used and the underlying assumptions – including justification for their choice and any adaptations;
- ▶ reports that enable a competent reader to understand the model and assess the judgments made; and
- ▶ specific references to the degree of conservatism in the results and other limitations of the model.

3. Modelling considerations


This section considers the factors that drive realised volatility and additional issues that impact implied volatility.

3.1 Realised volatility

As suggested in chapter 9 of the Part II actuarial text book³, it is useful to graph the data before developing the model. Rolling 12 month measures of monthly share price volatility are shown below for the US and Australia.

³ Smith A D, "Modeling" (2010) in *Understanding Actuarial Management: The Actuarial Control Cycle*, 2nd edition, ed Clare Bellis, Richard Lyon, Stuart Klugman and John Shepherd, Institute of Actuaries of Australia

Figure 1



Poon and Granger (2003), in a review of 93 papers, summarise the research and confirm what can be observed from the graph above:

“There are several salient features about financial time series and financial market volatility that are now well documented. These include fat tail distributions of risk asset returns, volatility clustering, asymmetry and mean reversion, and co-movement of volatilities across assets and financial markets. More recent research finds correlation among volatility is stronger than that among returns and both tend to increase during bear markets and financial crises. Since volatility of financial time series has complex structure, Diebold *et al* (1998) warn that forecast estimates will differ depending on the current level of volatility, volatility structure (e.g. the degree of persistence and mean reversion etc.) and the forecast horizon.”

Appendix B discusses each of these issues in more detail, considering the causes and how they can be modelled. It also notes additional issues that arise from measuring the volatility of infrequently valued assets.

3.2 Implied volatility

As noted above, implied (or imputed or notional) volatility is the volatility that, when inputted into an option pricing model, produces an option value equal to the current market price of that option.

Generally the Black-Scholes option pricing formula is the option pricing model used to determine implied volatilities.


3.2.1 Interpreting implied volatility

Implied volatility can be interpreted as the “market view” of expected future realised volatility over the life of the option if the assumptions underlying the option pricing model apply. For example, if the structure of the Black-Scholes model was used, it would be assumed that:

- ▶ equity and option markets are complete and liquid: efficient prices are available for all options;
- ▶ transaction costs are low;
- ▶ market price makers invest and borrow at the same “risk free rate”; and
- ▶ asset prices are log-normally distributed, the parameters being constant over the period concerned.

As none of the above assumptions always apply in practice, implied volatilities calculated using the Black-Scholes formula may be, but are unlikely to be, unbiased estimates of the “market view” of expected future realised volatility over the life of the option. Haug and Taleb (2011) suggest that market makers in options set prices based on a variety of heuristics that more accurately reflect the views of the market, and that the implied volatility is one way of expressing the price.

Some empirical research suggests that implied volatility is on average higher than realised volatility, while other research suggests that implied and realised volatility may take on similar values. For example, analysis by Li and Yang (2008) found that imputed volatility was an efficient predictor of future volatility in the US and Australia over the period 2001 and 2006. In contrast, the chart below suggests that over the 1 year period to 30 June 2011, implied volatility was higher than realised volatility.



3.2.2 Observed term structure of implied volatility


When the Black-Scholes option pricing formula is adopted to determine implied volatility, the implied volatility calculated for options over the same underlying asset or interest rate will generally vary depending on the term of the option and the exercise price of the option.

This is similar to interest rates varying by duration and the existence of a term structure for interest rates. The term structure of interest rates allows market participants to determine spot rates and forward rates at points in time in the future, which are then commonly used in valuing future cash flows. Historically, the term structure of interest rates has shown that there is a correlation between short-term interest rates and long-term interest rates. However, the short- and long-term interest rates are not fully correlated and moved differently. This leads to the yield curve changing shape, and sometimes to inverting (that is, short-term interest rates being higher than long-term interest rates).

The same concepts apply to implied volatility and a term structure for implied volatility can be derived from option market prices. Like short-term and long-term interest rates, short-term and long-term implied volatility show a degree of correlation. It can be noted however that, while forward rates for any period can be uniquely determined from the yield curve, the term structure of volatility can arise from anticipated future changes in the level of volatility or from the correlation structure. (For example, the volatility curve normally rises and then falls over time. This may reflect an expectation that the future volatility will rise and then fall, or positive autocorrelation (momentum) in the short run that reverses – possibly due to mean reversion – at a later date.) The volatility of the implied volatility is also generally greater than that of

interest rates, that is, implied volatility can vary significantly from week to week as shown (for equities) in the chart below.

Figure 2: Volatility is variable



3.3 Smiles and smirks

The volatility smile is the pattern usually observed by plotting the implied volatilities of options on a particular instrument against different strike prices. They are lower at the strike prices close to spot prices and higher at prices that are much lower or much higher. The smile is skewed – or turned to a smirk – as higher volatilities are observed at strike prices that are lower than the spot price.


Figure 3: The smile and smirk



Smiles and smirks represent violations of the Black-Scholes assumption that the underlying volatility does not vary with how close or how far the option is in or out of the money. The volatility smile can be explained by asset returns exhibiting fat tails.

The progress of the smile in the Australian equity market is shown in the chart below.

Figure 4: The smile over time



Generally the volatility is lower for options that are in the money than are out of the money. The chart shows that, for the Australian share market, this can vary between 4.5% and 5.5% for a put option that is 10% out of the money and the same put option 10% in the money.

The smirk will be aggravated if the demand for downside protection is particularly high. Xing *et al* (2008) found that the smirk pattern on individual shares can predict a falling share price

and may represent selling by better informed investors. It has also been ascribed to the growth in investment guarantees under variable annuities, and the need for insurance companies (or their hedging counterparties) to purchase matching options.

3.4 Implied volatility using call options and put options

Empirical analysis of put options and call options written over the same underlying security has typically found that the implied volatility calculated using the put options is higher than the implied volatility calculated using call options. This result contravenes put/call parity and reflects that the assumptions underlying the Black-Scholes model do not apply in practice.

For further details refer to Ahoniemi and Lanne (2007).

4. Applications

4.1 Setting a volatility assumption for a parametric option pricing model

The above sections highlight that there are two different starting points (implied volatility and historical realised volatility) commonly adopted for determining the volatility assumption for the use in option pricing models like the Black-Scholes option pricing model or related models.

Irrespective of the starting point adopted, a number of adjustments may be required in order to set the volatility assumption. The following table summarises these adjustments.

Table 4.1 Adjustments for valuation purposes

Reason for adjustments	Implied volatility	Historical realised volatility
Future realised volatility differs from historical realised volatility	No adjustment required. Implied volatility reflects market's expectations of future realised volatility.	An adjustment may be required. This reflects that it appears that volatility levels are not constant and tend to mean revert.
Implied volatility is a biased estimator of expected future realised volatility (that is, assumptions underlying the pricing model do not hold in practice)	No adjustment required.	An adjustment may be required, if the assumptions underlying the pricing model do not hold in practice. Some empirical research suggests that implied volatility levels are higher than the emerging realised volatility.

Reason for adjustments	Implied volatility	Historical realised volatility
Moneyiness of option	Adjustment generally required if moneyiness of option being valued deviates materially from the moneyiness of the options used to determine implied volatility. This reflects that implied volatility generally exhibits a "smile" or "smirk".	An adjustment may be required for options that are significantly "in the money" or significantly "out of the money". This reflects that implied volatility generally exhibits a "smile" or "smirk".
Term of liability	May be required if term of options used to determine implied volatility differs from term of asset/liability being valued. This reflects that implied volatility generally exhibits a term structure.	Different calibrations may be required for assets/options with different terms. This reflects that implied volatility generally exhibits a term structure.
Type of option (call or put)	Adjustment generally required if valuing a put (call) and implied volatility calculated using call (put) prices. This reflects that implied volatility is generally higher for put options than call options.	Different calibrations may be required for calls and puts. This reflects that implied volatility is generally higher for put options than call options.
Assumptions underlying option pricing model used to determine implied volatility deviates from assumptions underlying option pricing model used to value asset or liability	May be required.	Not applicable.

4.2 Setting a volatility assumption for stochastic models

The same different starting points (implied volatility and historical realised volatility) can be adopted for determining the volatility assumption for stochastic models whether used for capital reserving or asset allocation purposes.

Both approaches may require adjustments. The adjustments required are generally similar to those outlined above.

The following table summarises the two key adjustments that generally require consideration.

Table 4.2 Key adjustments for modelling purposes

Reason for adjustments	Implied volatility	Historical realised volatility
Future volatility differs from historical realised volatility	An adjustment may be required. The need for an adjustment will depend on the capital reserving methodology adopted. For example, if the model requires a long term estimate of future volatility, using an implied volatility assumption that reflects short term expectations of future realised volatility would not be appropriate	An adjustment may be required. The considerations are similar to those applying when using implied volatility except further consideration may need to be given to patterns of heteroskedasticity: mean reversion or regime shifting
Asset return modelling approximations and other modelling approximations	An adjustment may be required. The need for adjustment will depend on whether: ▶ the option pricing model provides an unbiased estimate of future volatility. If it does not, an	An adjustment may be required. For example, if the stochastic model assumes asset returns are normally distributed, it may be appropriate to adjust the volatility assumption to reflect that asset returns may exhibit skewness or kurtosis⁴. Alternatively, if the stochastic model assumes there is no uncertainty associated


⁴ Appendix C provides some brief observations on asset returns.

Reason for adjustments	Implied volatility	Historical realised volatility
	adjustment may be required. ► The asset return distribution underlying the model is accurate. If not, an adjustment may be required. Refer to the cell to the right for examples.	with the level of future volatility, it may be appropriate to adjust the volatility assumption adopted to reflect that, in practice, the level of volatility is uncertain. It may be inappropriate to increase estimates of future volatility in times when it is currently at historically higher levels given its propensity to mean revert. It is noted that in both of the above examples an adjustment to volatility would most likely not be required if the methodology allowed for these approximations through another mechanism (for example, adjustment to the confidence level adopted). Additional examples are outlined in section 5.6.

5. Practical considerations

5.1 Volatility surface

As discussed in section 4, implied volatility varies with the duration of the option and with the in-the-moneyness of the option, giving rise to the existence of a volatility surface. Depending on the purpose of the volatility assumption, consideration will be given to the use of a volatility surface assumption. An example is shown in the chart below.



This chart shows an estimated volatility surface on the ASX 200 as a point in time. The implied volatility is plotted against duration (time to maturity) and moneyness. There is no smile, but a volatility smirk, which is demonstrated by the fact that the implied volatility increases as the strike price decreases, that is, the put option moves further out of the money. The lines show that the volatility also varies by time, although the in-the-moneyness appears to be the main driver in this case.

5.2 Monitoring

As illustrated above, volatility is variable and changes with time. The frequency of monitoring and the time period over which it is measured will depend on the needs of users, particularly the risks to which they are exposed as a consequence of the volatility and their capital resources.

5.3 Obtaining market prices

The prices of equity and commodity options listed on the ASX are available on their website, and on those of a number of online brokers. The price of interest rate options and implied volatilities are available from larger brokers and market makers, but they normally require an account. Information providers such as Bloomberg and Reuters provide a wide range of current and historical prices and volatilities to subscribers. Some prices and volatilities are also reported in the Australian Financial Review.

Some investment banks will issue more unusual and exotic options as over-the-counter ("OTC") options. They are likely to charge higher margins (that is, higher options prices) leading to higher implied volatility.

While a range of put and call options are traded on a large number of Australian stocks, the market is not that deep or liquid. Prices are therefore not necessarily mutually consistent. Depending on the need for accuracy, implied volatility may need to be determined using a range of prices over a number of days in order to smooth away any obvious anomalies.

The same applies to interest rate volatility. Whilst the market in long-dated bonds is liquid and long-term interest rates can reasonably reliably be observed, there is not a liquid market to reliably determine a market view on long-term expected volatility. In practice, quoted long-term implied volatilities are generally derived from the liquid short-term option market with profit margins and risk allowances added. This means that, in a market where volatility is driven by more short-term events (for example, the earthquake in Japan in March 2011), the implied volatility for the next 20 years, to the extent that any data are available, increases significantly. This seems unlikely to be reflective of genuine expectations of increased volatility in the long-term and instead is a reflection of the fact that long-term options are not traded frequently and no liquid market exists for long-term options.

In all Australian markets, liquidity of traded options is limited past the one-year duration mark. An implication of this is that actuaries may need to use a combination of implied volatility and historic volatility to determine an appropriate volatility assumption. For example, implied volatility data may be used to set short duration volatility, while historic volatility may be used to determine long term volatility.

There is no market price by which to determine the volatility of insurance claims, expenses and discontinuances. Volatility needs therefore to be based on historical analysis, with care being taken to adjust for any smoothing in the measurement, and for the appropriate autocorrelations when looking at longer terms.

5.4 Uncertainty

Estimation methods will result in some level of uncertainty on the assumption used for volatility. In order to address this uncertainty it may be necessary to consider estimation error and what impact that would have on results. One way this may be demonstrated is through sensitivity analysis.

5.5 Other considerations

When determining prices from implied market prices, there are potentially other risk premia at play, for example, liquidity premiums and credit risk. Depending on the purpose of the assumption, it may be necessary to strip these risk premiums out of the volatility assumption or to adjust their magnitude to reflect the nature of the asymmetric risk being valued.

5.6 Communication

A report documenting the volatility assumption may also need to include:

- ▶ a definition of the assumption and whether it is a best estimate, or it contains an element of prudence or adjustment, and if so how much?
- ▶ documentation on data, including the source and the extent that it has been modified (grouped or outliers removed) and / or checked;
- ▶ implicit or explicit assumptions that were made in order to determine the volatility assumption (for example, distribution of data); and
- ▶ any other information in order to allow a suitably qualified professional to reproduce the results.

The users of the report may not necessarily appreciate statistical terminology, and therefore the structure of the technical information would need to be considered carefully to avoid confusion for non-technical readers.

The concept of volatility may rely on a statistical training to fully appreciate the implications of the assumptions. Perhaps providing scenarios (or other simplified explanations) will assist in fulfilling the communication role to non-technical audiences.

6. Conclusion

An actuary may use volatility assumptions for a range of purposes including valuing asymmetric risks and calculating capital reserves. The volatility assumption is complex and care must be taken in understanding the nuances of volatility when setting an assumption. Limitations in any approximation method need to be understood to ensure end results are not misleading.

The method used to determine the volatility must consider the end use of the assumption, particularly when considering whether to use implied volatility or historical volatility as a basis for setting the assumption. Irrespective of the starting point (implied or historical) used to determine volatility, adjustments may be required.

Finally an appropriate form of communicating the volatility assumption and assumption setting process is required, and typically this would involve a report where a suitably qualified professional could replicate the results. Communication to non-technical users may be addressed through less technical means such as providing scenarios.

7. References

Ahoniemi K and Lanne M, 2007, Joint Modeling of Call and Put Implied Volatility. Discussion paper 198 of the Helsinki Centre of Economic Research. Available at SSRN: <http://ssrn.com/abstract=1033804>

Derman, Emanuel and Wilmott, Paul, The Financial Modelers' Manifesto (January, 08 2009). Available at SSRN: <http://ssrn.com/abstract=1324878>

Haug, Espen Gaarder and Taleb, Nassim Nicholas, Option Traders Use (very) Sophisticated Heuristics, Never the Black–Scholes–Merton Formula (February 26, 2009). *Journal of Economic Behavior and Organization*, Vol. 77, No. 2, 2011. Available at SSRN: <http://ssrn.com/abstract=1012075>

Li, Steven and Qianqian Yang, 2008, The relationship between implied and realized volatility: evidence from the Australian stock index option market, *Review of Quantitative Finance and Accounting* 32.4: 405-419

Michaud, Richard and Michaud, Robert, 2008, Estimation Error and Portfolio Optimization: A Resampling Solution. *Journal of Investment Management*, 6 1, pp 8–28

Yuhang Xing, Xiaoyan Zhang and Rui Zhao, 2008, What Does Individual Option Volatility Smirk Tell Us About Future Equity Returns?
http://papers.ssrn.com/sol3/papers.cfm?abstract_id=1107464

Appendix A – Black-Scholes

The Black-Scholes⁵ option pricing formula is a closed form solution to a stochastic model.

$$Call = N(d_1) * S - N(d_2) * K * e^{-r*(T-t)}$$

$$d_1 = \frac{\ln(\frac{S}{K}) + (r + \frac{\sigma^2}{2}) * (T-t)}{\sigma * \sqrt{T-t}}$$

$$d_2 = d_1 - \sigma * \sqrt{T-t}$$

Where:

$N(*)$ is the cumulative distribution function of the standard normal distribution

$T-t$ is the time to maturity

S is the spot price of the underlying asset

K is the strike price

r is the continuously compounded risk free rate

σ is the volatility of returns of the underlying asset

⁵ A range of option valuation models exist, either as enhancements to the Black-Scholes model or as separate models.

Appendix B – Modelling realised volatility in more detail

Introduction

This appendix provides further details on the behaviour of realised volatility.

Changes in volatility

In expanding the discussion in section 3.1 (using the graph for illustration purposes), one would expect that realised volatility would increase:

- ▶ in less certain times when market information is subject to more widely differing interpretations;
- ▶ when investors are forced buyers or sellers (short selling can create forced buyers in a bear squeeze: the 1980 peak in volatility, for instance, seems to have been aggravated by the notorious Bunker Hunt squeeze in the silver market).

Asymmetry

Falls in the equity market and in the level of interest rates tend to be faster than increases. Observed volatility is therefore often larger on the downside than the upside. (This is consistent with implied volatility, where the smile is often more of a lop-sided volatility smirk.) The frequency distribution of volatility itself is asymmetric with a very long tail in respect of volatility increasing.

Clustering

Volatility levels change over time. That is volatility appears to cluster, in that it appears to jump to different levels for a period and then mean reverts. It is noted that some of these periods are relatively short, while others are longer. For example, the spike in US volatility following the tech crash in 2001 lasted longer than that coming out of the 1987 crash.

There are two main approaches to reproducing the clustering observed in the time series:

- ▶ the first is to use G/ARCH (Generalised /auto-regressive conditional heteroskedasticity) type models where the expected volatility varies with recent experience (of both mean and variance); and
- ▶ the second is to use regime shifting models where the volatility switches randomly between two or more regimes. Expected returns can also be linked to the different regimes.

These models produce equity return distributions that exhibit fat tails when analysed over longer periods.

Mean reversion

Volatility is likely to be mean reverting. There appears to be an upper limit as it is inconceivable that share prices should vary infinitely. With over 80 years of data, it may well be that the limit is not far from the 20% observed in 1932/3 in the US, when the market rose from its nadir by more than 30% in three separate months – in that case, higher volatility occurring during a short but very strong bull market.

It is therefore not likely to be appropriate to base capital requirements on an assumption that volatility will increase by a fixed proportion of the current level.

Different markets may revert to similar means. There appears to be an increasing convergence between the US and Australian markets – particularly in the past four years – apparently as a result of the ongoing globalisation of markets.

Infrequently valued assets

Infrequently priced assets, such as unlisted property, display artificially low volatility under benign circumstances because the published prices reflect a moving average of the underlying market prices. Extreme care needs to be exercised if option pricing models are used for infrequently traded assets, even if volatility estimates have been appropriately “de-smoothed”. This is because the arbitrage-free assumption of a replicating strategy is invalid for such assets.

It can be noted, however, that in circumstances of reduced market liquidity, published prices show fatter tailed behaviour under pressure. This can arise because:

- ▶ the underlying assets are all revalued simultaneously to give a better estimate of market value; or
- ▶ the underlying assets are sold at “fire sale” prices because of their inherent illiquidity.

Appendix C – Asset returns

The distribution of observed equity log returns and interest rates have fatter tails than the normal distribution, which means that asset return models that assume returns are normally distributed can produce misleading results. The fat tails can be modelled by:

- ▶ non-normal distributions (such as the Student's t distribution) which have the advantage of analytical tractability;
- ▶ compound distributions, which are sometimes tractable, but are more easily simulated;
- ▶ clustering models as described above;
- ▶ extreme value distributions (such as the Pareto) that are only used for the tail. This approach is commonly used in banking for operational risk modelling, but can leave unanswered the question of consistency between the events modelled in the body of the distribution and probabilities in the often arbitrary tails;
- ▶ using nonparametric distributions derived from historical bootstrapping or from ex-ante scenarios analysis based on judgment; and
- ▶ adjusting the volatility assumption. It is noted that if this approach is adopted, results in the tail of the distribution will be accurately modelled, but results in the middle of the distribution will not be accurately modelled.

End of Technical Paper