



INTERNATIONAL VALUATION STANDARDS COUNCIL

THE VALUATION OF EQUITY DERIVATIVES

EXPOSURE DRAFT

Comments on this Exposure Draft are invited before 30 September 2013. All replies may be put on public record unless confidentiality is requested by the respondent. Comments may be sent as email attachments to:

CommentLetters@ivsc.org

or by post to IVSC, 68 Lombard Street, LONDON EC3V 9LJ, United Kingdom.

Copyright © 2013 International Valuation Standards Council. All rights reserved. Copies of this Exposure Draft may be made for the purpose of preparing comments to be submitted to the IVSC provided such copies are for personal or intra-organisational use only and are not sold or disseminated and provided each copy acknowledges IVSC's copyright and sets out the IVSC's address in full. Otherwise, no part of this Exposure Draft may be translated, reprinted or reproduced or utilised in any form either in whole or in part or by any electronic, mechanical or other means, now known or hereafter invented, including photocopying and recording, or in any information storage and retrieval system, without permission in writing from the International Valuation Standards Council. Please address publication and copyright matters to:

International Valuation Standards Council
68 Lombard Street
LONDON EC3V 6LJ
United Kingdom
Email: ivsc@ivsc.org
www.ivsc.org

The International Valuation Standards Council, the authors and the publishers do not accept responsibility for loss caused to any person who acts or refrains from acting in reliance on the material in this publication, whether such loss is caused by negligence or otherwise.

Introduction to Exposure Draft

One of the contributory causes of the 2008 financial crisis was that the rapid growth in the complexity of financial products over the previous decade and in the complexity of the models developed to value them had created a situation where senior management, independent directors, investors and regulators had insufficient understanding of the products and processes to determine the appropriate degree of confidence that should be placed on the valuations that were provided.

It has also become apparent that some of the accounting requirements for measuring instruments would have benefited from being informed at the time of their creation by an established source of recognised valuation best practice. In the absence of this many consider that there are mismatches between the way in which the accounting fair value of an instrument has to be determined and the way in which it would be valued by market participants.

The IVSC produces valuation standards and technical guidance with the objective of improving consistency and transparency in valuation, which in turn builds confidence in valuations by those who rely on them. This proposed Technical Information Paper is the first of a planned series of papers designed to bring greater consistency and transparency to the valuation of derivatives based on different asset classes, ie fixed income, credit, foreign exchange, commodities and hybrid products.

This Exposure Draft has been produced in order to invite input and comments from valuation specialists, investors, analysts, auditors and anyone else who needs to produce or rely on valuations of equity instruments. Before responding, readers are recommended to read the description of the purpose and status of an IVSC Technical Information Paper on page 1.

Questions for Respondents

The IVSC Standards Board invites responses to the following questions. Not all questions need to be answered but to assist analysis of responses received please use the question numbers in this paper to indicate to which question your comments relate. Further comments on any aspect of the Exposure Draft are also welcome.

Notes for respondents:

In order for us to analyse and give due weight to your comments please observe the following:

1. Responses should be made in letter format, where appropriate on the organisation's letter heading. Respondents should indicate the nature of their business and the main purpose for which they either value, or rely upon the value of, equity derivatives.
2. Comments should not be submitted on an edited version of the Exposure Draft.
3. Unless anonymity is requested, all comments received may be displayed on the IVSC website.
4. Comments letters should be sent as an email attachment in either MS Word or an **unlocked** PDF format and no larger than 1mb. All documents will be converted to secured PDF files before being placed on the web site.
5. The e mail should be sent to commentletters@ivsc.org with the words "Equity Derivatives" included in the subject line.
6. Please be sure to submit comments before the 30 September 2013.

Questions

1. Under the heading of "Equity Derivative Products" (para 11-22) the main types of equity derivative are listed. Do you believe there are any material omissions? If so, please indicate what they are.
2. Do you believe the descriptions provided for each of the listed products are sufficiently detailed?
3. Do you think more complicated derivatives and strategies should be included? For example where products are combined, such as in straddles and strangles?
4. The discussion on forwards (para 23 to -27) includes a number of formulae. Do you find the inclusion of formulae to be helpful in understanding the principles or would you prefer a purely descriptive narrative?
5. Would you prefer to see greater use being made of formulae to illustrate principles in other parts of the TIP?
6. The discussion of various models types includes the key assumptions and other inputs required. The objective is not to provide detailed instruction on the use of the model, but do you think the information on these inputs is sufficiently detailed to provide an understanding of the principles involved by someone relying on the valuation?

7. Do you believe the model section of this paper should discuss each model's relative applications and when it is appropriate to use one rather than another, for example, by mapping each model to a list of products?
8. "The Greeks" are summarised with brief descriptions in this paper. Do you believe it would be helpful if there were a more detailed discussion of sensitivities?
9. Please list the departments within your organisation that you believe would find this document useful, e.g. Executive Management, Treasury, Risk, Financial Reporting, Product Control etc.
10. Do you consider that the overall level scope and level of detail in this proposed TIP is sufficient to meet its objective of reducing diversity of practice and raising awareness of the principle methods used for valuing equity instruments among the wider financial community, and in particular investors?

Exposure Draft

The Valuation of Equity Derivatives

Technical Information Papers

The principal objective of an IVSC Technical Information Paper (TIP) is to reduce diversity of practice by identifying commonly accepted processes and procedures and discussing their application. A TIP is designed to be of assistance to professional valuers and informed users of valuations alike. A TIP will do one or more of the following:

- Provide information on the characteristics of different types of assets that are relevant to their value.
- Provide information on appropriate valuation methods and their application.
- Assist the consistent application of an International Valuation Standard (IVS) by dealing with matters identified in the Standard in greater detail.
- Provide information that is helpful to valuation professionals in exercising the judgements they are required to make during the valuation process in specific situations.

A TIP does not:

- Provide valuation training or instruction.
- Direct that a particular approach or method should or should not be used in any specific situation.

While a TIP may provide simplified examples to illustrate points discussed, these are not intended to be used as templates that can be applied to real life situations. Neither does a TIP set out to cover every variation of a method or technique that may be appropriate in practice. Responsibility for choosing the most appropriate valuation method or methods is the responsibility of the valuer based on the facts and circumstances of each valuation task.

The guidance in this paper is intended to support the application of the principles in the International Valuation Standards (IVSs) to the valuation of equity derivatives. The following IVSs are of particular relevance to the valuations discussed in this TIP:

The IVS Framework,

IVS 101 Scope of Work

IVS 102 Implementation

IVS 103 Reporting

IVS 250 Financial Instruments

Contents

| | Paragraphs | Page |
|--|------------|------|
| Introduction and scope | 1-6 | 3 |
| Definitions | 7 | 4 |
| The Equity derivatives market | 8-10 | 4 |
| Equity derivative products | 11-23 | 5 |
| Valuation of forwards and futures | 24-28 | 7 |
| Equity swaps | 29-30 | 8 |
| Option pricing models | 31-38 | 9 |
| Risk Neutral Valuation | 39-41 | 10 |
| Black Scholes method | 42-45 | 10 |
| Black Scholes extensions | 46 | 11 |
| Black Scholes and exotic derivatives | 47-52 | 11 |
| Limitations of Black Scholes | 53-60 | 12 |
| Alternative stochastic processes | 61 | 13 |
| Jump diffusion models | 62-71 | 13 |
| Alternative diffusion models | 72-83 | 14 |
| The Greeks | 84 | 16 |
| Model resolution | 85 | 16 |
| Implementation | 86-88 | 17 |
| Calibration | 89-96 | 18 |
| Conclusion | 97-98 | 19 |
| Appendix 1 - Closed form solutions under the Black-Scholes framework | | |

Introduction and Scope

1. The objective of this TIP is to produce a high level guide as to the accepted principles of best practice for valuing equity instruments, ie derivative contracts where the underlying asset is equity or an equity based product. It is intended to:
 - Reduce diversity of practice by identifying the commonly accepted methods and procedures for valuing various types of equity derivative.
 - Promote greater transparency in the valuation process by increasing the awareness of those who use or rely on valuations of the principles involved.
2. This TIP is concerned with the valuation of derivatives in either an actual or assumed exchange, such as is required for financial reporting, fund management or supporting transactions. It is not concerned with the pricing of derivative products, ie the analysis of the costs of creating and holding an instrument in order to determine the price that needs to be achieved by the creator in order to show the required return or other objective.
3. *IVS 250 Financial Instruments* applies the principles contained in the IVS General Standards (IVSs 101, 102 and 103) to the valuation of financial instruments. The Commentary to IVS 250 contains an overview of factors that need to be considered when valuing financial instruments and identifies the main categories of valuation method that are used. The objective of the TIP is to assist those providing valuations to identify best practice and those who rely on them to better understand the key assumptions and inputs on which common methods and models rely. This TIP does not provide a detailed examination or critical analysis of different types of model and neither does it provide instruction in their use.
4. The fact that a particular type of method or model is not referred to in this TIP does not mean that its use may not be appropriate in specific circumstances. The purpose of this paper is to identify methods and models that are generally accepted as being appropriate for a particular category of instrument, but it is not intended to limit the scope of methods where the facts of a particular situation differ from those described.
5. This paper indicates the most relevant methods and key principles of model resolution. The purpose of this TIP is not to present detailed algorithms or computational programs but to explain the principal steps needed to solve the specified models. This TIP is not intended to be a comprehensive guide to every known resolution method and others may be appropriate if this is indicated by the facts.
6. Input selection and estimation are key factors in the valuation of derivative instruments. The paper describes the most important procedures and sources to estimate inputs that are used globally.

Definitions

7. The following definitions apply in the context of this TIP. Similar words and terms may have alternative meanings in a different context. The IVSC's International Glossary of Valuation Terms provides a comprehensive list of defined words and terms commonly used in valuation, together with any alternative meanings.

| | |
|----------------|--|
| Cash flow | An exchange of monies between counterparties |
| Coupon | Interest payment made throughout the life of an investment, similar to a <i>dividend</i> |
| Deterministic | A process in which no randomness or variables are involved. |
| Dividend | A regular payment from a company to shareholders |
| Leg | A <i>cash flow</i> on one side of a swap agreement. |
| LIBOR | London inter-bank offered rate. |
| Maturity | The last day that an options or futures contract is valid. |
| Notional | The face value of an instrument used to calculate payments between the counterparties |
| Path-Dependent | An option where the strike price is based on fluctuations in the value or price of an underlying asset during the contract term. |
| Payoff | The final payment or other financial gain that accrues to the beneficiary. |
| Stochastic | A process involving a random variable or series of variables. |
| Strike Price | The price at which an option can be exercised. |

The Equity Derivatives Market

8. The global market for equity derivatives has been estimated by the Bank of International Settlements as having a gross value of many hundreds of billion dollars, with the outstanding *notional* amount measured in trillions of dollars. The equities underlying these instruments can be single stocks, baskets of stocks, or indices of standardised baskets of stocks traded on exchanges.
9. Equity derivatives are either exchange traded or dealt in the over the counter (OTC) markets.
- Exchange traded contracts (also called "listed options") tend to be the simplest products, and often are standardised contracts with common strike and *maturity* provisions. This minimises the number of bespoke contracts and maximises liquidity. By virtue of being exchange traded, counterparty credit risk is mitigated, as counterparties are required to post collateral and/or settle margins regularly.
 - OTC contracts are traded directly between counterparties, and tend to be bespoke; whilst also existing in simple forms, they are often more complicated than those traded

on exchanges. These are subject to the risk of the counterparty defaulting, unless collateralised in one or both directions

10. Equity derivatives are contracts based on price changes in an underlying asset that can be a single stock, a basket of stocks or on movement in an index based on a basket of stocks. Equity derivatives can be settled in cash, physically or a combination of the two. Derivative contracts will, as standard, contain details for cash or physical settlement of a derivative, along with the criteria under which either will occur. Physical settlement entails a transfer of the relevant underlying equity between parties. Fractional share amounts tend to be settled in cash.

Equity Derivative Products

11. The main genres of derivative product are described below in order of increasing complexity. Standard contracts that are regularly traded are often referred to as “vanilla” and the more complex as “exotic”. The evolution of derivatives’ markets means there is no concise rule to classify a product as either vanilla or exotic, so what is now considered vanilla may well historically have been considered as exotic, for example American options.

Forwards and futures

12. A forward contract is an agreement to purchase or sell an underlying equity at a future date at today’s predetermined price; these are OTC contracts. Futures are their exchange-traded equivalents with standardised maturities and lot size, where daily fluctuations in contract price must be settled by the relevant counterparties (marked-to-margin).
13. Forwards and futures are also called “delta-one products” because the change in their price is equal to the change of underlying asset’s price. For market risk hedging purposes, this implies that they may be hedged with one unit of the underlying asset.

Equity Swaps

14. An equity swap is a contract where a set of future *cash flows* are agreed to be exchanged between two counterparties at set dates in the future. In an equity swap one of the *legs* is usually pegged to a floating rate such as LIBOR and the other *leg* is based on the price of either a specific share or a stock market index. In some cases an equity swap can have two equity *legs* instead of one being a floating rate.

Options

15. An option is a contract which gives the holder the right, but not the obligation, to buy or sell an underlying asset at a specified price, called the *strike price*, on or before a specified date. The buyer of an option pays a premium to the seller for this right. An option which gives the right to buy something at a specific price is called a call; an option which gives the right to sell something at a specific price is called a put. An option may be in respect of a single asset or a basket of different assets when an investor wants exposure to a particular industry, or to hedge or diversify across several industries.
16. The option holder will choose to exercise a call option when the market price of a stock is above the *strike price* of his contract and exercise a put option when the market price is below his *strike price*. An example is where an investor pays \$1 for a call option to buy one share in

a company with a *strike price* of \$20 in three months' time. If the market price is \$25 in three months' time the investor will exercise the option and the *payoff* will be \$4 ($\$25 - \$20 - \1). If the market price were below \$20 the investor would choose not to exercise and will have lost the premium of \$1. If the price is greater than \$20 and less than \$21, for example \$20.50, exercising the option will reduce the loss to \$0.50.

17. Where the option relates to multiple assets, there are a numerous permutations of possible *payoff* but some straightforward examples include:
- Best-of or worst-of – where the *payoff* depends on the best/worst relative performance of the assets in the basket.
 - Knock-out components – not to be confused with barrier options, where certain assets are removed from featuring in the *payoff* over the life of the option.
 - Exchange options – where the holder has the right to exchange one asset for another in predefined proportions.
18. There is a wide variety different option types. Some of the more common are described below. The majority of OTC options traded will either be one of types listed or a combination of two or more.
- European – the holder may exercise the option only on its date of *maturity*.
 - American – the holder may exercise the option at any point up to and including *maturity*. These are generally more valuable than European options as the holder has more rights but holder has to determine when it is most advantageous to exercise.
 - Bermudan – exercise is only possible at a certain number of dates up to *maturity*.
 - Asian – instead of depending on the underlying asset price at a single time, the *payoff* depends on the average price. The averaging period may be over the entire life of the contract or just a portion.
 - Digital/Binary – the *payoff* is either a fixed sum or zero.
 - Barriers – here a standard *payoff* is coupled with a separate knock-in/out barrier that may be observed at particular points or throughout the life of the trade, which triggers whether or not the *payoff* is received.
 - Lookbacks – the *payoff* depends on the minimum and/or maximum performance of the equity over the life of the option.
 - Range accruals – variations on barrier options, where the *payoff* depends on the cumulative period an underlying asset has spent within a certain range.
 - Forward-starts – certain features of the option (for example, the strike) are fixed at a future point in time, dependent on the performance of the underlying equity.
 - Cliquets – these are portfolios of consecutive forward-starting options, for example, eight three-month call options, covering a two-year period. The strike for each subsequent contract is set once the previous contract matures, and there may be local and/or global caps or floors on the total *payoff*.

Structured Products

19. At the more complex end of the spectrum are structured products, which are usually funded (meaning the buyer often deposits a *notional* with the issuer). These can generally be viewed as a bond coupled with an equity option containing any number of features. Commonly the features may consist of:
- *Coupons* linked to the equity performance.
 - The whole or part of the *notional* may be at risk.
 - The *notional* may amortise, i.e. be paid back over time.
20. These often appeal to investors because they have a large *coupon* which can give exposure to a specific equity. The risks, on the other hand, are often complex.
21. Frequently, OTC products contain an embedded interest rate swap, which is constructed to pay for the optionality.

Corporate issues

22. There are a number of derivatives that are originated predominantly to serve a purpose other than insurance, hedging or leverage. Examples include
- Warrants – These are very similar to call options, often containing early-exercise features. Unlike options, the issuer is usually the holder of the underlying equity, so settlement is undertaken by issuing new shares rather than existing ones.
 - Convertible bonds – used for raising capital. Corporations issue a bond that may be converted by the holder into a quantity of equity should certain performance criteria be met, and often these are also callable, meaning the issuer may also repay the debt early. Convertible bonds are complex, containing early-exercise features, and are really hybrid products of equity, interest rate, and credit

Volatility derivatives

23. Absorbing volatility is the predominant purpose of derivatives, so it is therefore natural that a market in trading volatility and variance itself has grown. A common contract is known as a variance swap, which references a particular underlying equity. Two parties agree to swap payments at a future time, where one pays a fixed amount, and the other pays an amount equal to the annualised realised daily variance of the equity over the period.

Valuation of forwards and futures

24. The value of forwards and future contracts is not affected by the stock volatility, making the valuation process more straightforward than for other products.
25. An example is a forward contract with *maturity* at time T and delivery price K. The buyer of the contract, the long position holder, has the obligation to buy the stock at time T at the predefined delivery price K. At any point in time the value of the contract will be the present value of the difference between the forward price of the underlying asset and the delivery price. If the forward price of the asset is higher than the delivery price, the contract will have a positive value.

26. For this type of derivative, the valuation process relies on two steps:
- An estimate the future price of the underlying asset.
 - The present value of the difference between the future price and the predefined delivery price.
27. The rational pricing assumption is that asset prices, and hence asset pricing models, will reflect the arbitrage-free value of the asset as any deviation from this value will be "arbitraged away". Under this assumption, given the spot price of an asset, its forward price is explained in terms of what it is known as the cost of carry. In general terms, this measures the storage cost, plus the interest rate paid to finance the asset, minus the income earned with the asset. For a non-dividend paying stock, since there are no storage costs and no income is earned, the cost of carry is simplified to the interest rate paid to finance the asset. For a *dividend* paying stock the cost of carry is given by the interest rate minus the *dividend* rate.

At time t , the forward price of an asset, F_t , whose spot price is S_t is given by¹:

$$F_t = S_t(1+r-d)^{(T-t)}$$

Where r is the interest rate, T is the *maturity* date and d is the *dividend* rate.

28. Once the forward price of the stock is estimated, the value of a long forward contract at time t , V_t , is given by the present value of the difference between the forward price F and the delivery price K :

$$V_t = (F_t - K) (1+r)^{-(T-t)}$$

Equity Swaps

29. Rational pricing also underpins the logic of equity swaps valuation. The value of a swap is the net present value of both sets of future *cash flows* "netted off" against each other. To be arbitrage free, the terms of a swap contract are such that, initially, the net present value of these future *cash flows* is equal to zero. However, once traded, its value may become positive or negative depending on the evolution of the underlying variables. At any time, the swap is valued by computing the present value of both *legs* and netting them.
30. In an equity swap, the floating rate *leg* is valued by, first, decomposing it into a series of forward rate agreements and, second, using the implied forward rates in the zero *coupon* curve to value them. The equity *leg*, which is based on the performance of a stock or a stock market index, is priced using the techniques explained in the discussion on option pricing models below. Depending on the structure of the equity *leg*, it will be appropriate to use one or another valuation model. Once the two *legs* are valued, the value of the swap is computed by simply netting the value of both *legs*.

¹ Discrete time is considered in the formulas.

Option Pricing Models

31. The value of any option contract depends on the price of the underlying asset today and the prices of the underlying asset in the future. Unlike forwards or futures, option prices are affected by the volatility of the underlying assets.
32. A key element in option's valuation is to set the model of behaviour of stock prices and to understand its properties.
33. The weak form of the efficient market hypothesis is based on the assumption that prices fully reflect the information implicit in the sequence of past prices. The semi-strong form of the hypothesis assumes that prices reflect all relevant information that is publicly available, while the strong form of the efficient market hypothesis assumes that all information that is known to any participant is reflected in market prices.
34. In weak-form hypothesis, future prices cannot be predicted by analysing prices from the past. Excess returns cannot be earned in the long run by using investment strategies based on historical share prices or other historical data. Share prices exhibit no serial dependencies, meaning that there are no patterns to asset prices. This implies that future price movements are determined entirely by information not contained in the price series. Hence, prices must follow a random walk. The weak form of efficient market hypothesis does not require that prices remain at or near equilibrium, but only that market participants not be able to systematically profit from market inefficiencies.
35. Any variable whose value changes over time in an uncertain way, as stock prices do, is said to follow a *stochastic* process. A Markov process is a particular *stochastic* process where only the present value of a variable is relevant to predict a future value. Stock prices are usually assumed to follow a Markov process. This states that the present price of a stock gathers all the relevant information contained in past prices. It is important to note that the Markov process is consistent with the weak form of market efficiency. Market competition tends to ensure that this form of efficiency holds. The evidence that there are many investors in the market trying to make profit from it leads to a situation where a stock price at any given time collects all the relevant information in past prices.
36. Otherwise, if a past pattern of a stock offered relevant information to predict future prices, market participants would immediately react neutralising the effect of said information as well as any profitable trading opportunity.
37. Stock prices are often assumed to follow a particular Markov process called a Geometric Brownian Motion (GBM). Practically speaking, this means that stock prices change at random, follow a Normal Distribution and can be decomposed into two parts: a *deterministic* side, which is the expected return of the stock, and a *stochastic* component, which is proportional to the volatility of the stock.
38. The GBM of stock was introduced by Samuelson in 1965 and it has been widely used. However, it is worth mentioning that some models do not use this process to capture price movements. Alternative models to the GBM are presented later.

Risk neutral valuation

39. An important general principle in option valuation is the assumption of risk neutrality. In a risk neutral world, investors require no compensation for risk and, therefore, the expected return of all the securities is the risk-free interest rate. The assumption of risk neutrality simplifies considerably the analysis of derivatives, as it allows using the risk free rate as the expected return of the stock.
40. Investors in the real world do not ignore risk, but the concept of risk neutrality is often misunderstood. When moving from a risk neutral world to one that is risk averse, two things change. Firstly the expected return of the asset increases because an additional premium is required by the investors to compensate for the risk. Secondly the discount rate used to discount the *payoffs* of the derivative also increases. These two offset each other, which is why the solutions under a risk neutral assumption are valid under any set of risk preferences.
41. In a risk neutral world, the value of an option is given by the expected present value of its *payoff* using the risk free rate to discount the flows

Black Scholes method

42. Black-Scholes (and Merton) presented a simple model in 1973 to price European options. The model uses a technique that involves the creation of a “replicating portfolio”, which mimics the *cash flows* of an asset by combining other financial instruments together into a portfolio. In this case the instruments involved were a European option, a riskless bond and the underlying stock of the option. The value of the option is then equivalent to the value of its replicating portfolio.
43. The Black Scholes (B-S) model is built upon the following assumptions:
- The underlying equity process follows the Geometric Brownian Motion explained previously.
 - The risk-free interest rate is constant and equal for all maturities.
 - There are no *dividend* payments during the life of the derivative.
 - There are no transaction costs or taxes.
 - There are no riskless arbitrage possibilities. In other words, there are no possibilities of making a risk-free profit.
 - All securities are perfectly divisible.
 - The short selling of securities is permitted. This is used to delta hedge the derivative by buying or selling the underlying equity in fractional amounts.
44. Under these assumptions Black and Scholes produced the B-S partial differential equation in order to indicate the value of an option. It is a formula which is relatively simple to apply to European options because it utilises observable inputs, with the exception of volatility.
45. It is worth noting that the Black-Scholes differential equation can be used to value more than just European options on stocks. For every different derivative that can be valued using the B-S methodology, there will be a unique differential equation of value. This is called a closed-form solution, where the value of a derivative can be calculated through a series of formulae.

This allows a valuer to determine how a derivative's value varies over time by solving the equation with updated inputs, which is very quick and convenient. Unfortunately, the B-S model cannot value all types of derivatives, meaning that valuers cannot obtain a closed-form solution for all derivatives. When this occurs, it is necessary to rely on calculation techniques that provide approximations to the solution of this equation. This is very often the case of complex derivatives which require numerical approximations to estimate their value as a pricing formula does not exist.

Black-Scholes extensions

46. Several assumptions of the original model were either removed or relaxed in subsequent extensions of the model. Some of the most important extensions are the following:
- Merton (1973) extended the Black Scholes model to allow for a continuous *dividend* yield. Black's formula (1976) gives the price of European options when the underlying security is a forward or futures contract.
 - Ingersoll (1976) allowed transaction costs and taxes.

Black-Scholes and exotic derivatives

Closed form solutions

47. As already mentioned, the Black Scholes formulae are most applicable to vanilla options. Nevertheless, a lot of research has been done into closed-form pricing formulas for exotic options on one or two underlying assets such as: Asiatic options, barriers, forward starting options, etc. Appendix 1 contains examples of some of these models.

Monte Carlo Simulation

48. Monte Carlo simulation is a numerical method that is considered in option pricing when no closed-form solution is available. It is a particularly useful technique when the derivative *payoff* is *path-dependent* and when it depends on more than one underlying asset.
49. This method relies on risk neutral valuation where the price of the option is its discounted expected value. Because Monte Carlo can be used to simulate a wide range of *stochastic* processes, a more detailed description of the method is provided in the resolution section. However the principle of the method is; first, to use the equation for the behaviour of stock prices to generate a large number of price paths; second, for every path, calculate the *payoff* of the option; and, finally, compute the average and discount it to the valuation date: the resulting number is the value of the option.
50. The Monte Carlo method is frequently used in exotic equity derivative valuation alongside a GBM in order to compute the possible paths that the underlying stock price can take.

Finite Difference Method

51. The B-S partial differential equation models the evolution of the option value as a function of time, the prices of the underlying assets, volatility, and some other parameters depending on the structure of the option. The finite difference method is used to find a numerical solution to the differential equation, which yields the estimated value for the option.

52. As in Monte Carlo simulation, the method is valid to solve not only the B-S partial differential equation but a wider range of differential equations. In the B-S framework this method is less used than Monte Carlo simulation as it is more limited in terms of the number of underlying variables it can deal with. This method is considered in more detail in the resolution section.

Limitations of Black-Scholes

53. The Black Scholes method can be considered the standard model both in terms of approach and applicability. However, despite its popularity and wide spread use, the model is built on some assumptions that under some circumstances are unrealistic. As the model is a simplification of reality, almost all the assumptions can be considered to be limiting and unreal. However, not all assumptions have the same impact on the indicated price.
54. As discussed previously, different extensions of the Black-Scholes models have already overcome some of these unrealistic assumptions, eg allowing for *dividends*, transaction costs, etc. These extensions should be used in circumstances where the assumptions in the original model would result in a significant distortion of reality when considered collectively.
55. There are other assumptions that have proven to be inappropriate in certain situations because they are not observable in stock prices data. These assumptions are mainly concerned with the *stochastic* process of the stock prices, volatility and interest rate hypothesis.
56. First, the GBM implies that stocks move in a manner such that investors cannot consistently predict the direction of the market or an individual stock. This is known as a random walk. A random walk means that at any given moment, the price of the underlying stock can go up or down with the same probability. Some argue that the GBM assumption in the B-S model is inconsistent with many documented patterns in stock prices, e.g. the tendency for prices at the beginning of January to be higher than at the preceding year end.
57. Second, asset returns are not always normally distributed. The model assumes that stock prices are log normally distributed and returns are normally distributed. However, experience has shown that returns have much more of a tendency to exhibit outliers than would be the case if they were normally distributed. There is overwhelming evidence that the returns are not always normal, but instead they can exhibit fat tails.
58. Third, volatility is not constant. The model assumes a constant volatility. However, various market crashes over time have shown this assumption to be unrealistic. While volatility can be relatively constant for short periods, it is not constant in the long term. Often, periods of high volatility follow immediately after a large change (often downward) in the level of the original series. For example, the plot of differences of the logs of the S&P 500 (an index of the largest 500 US Companies stock prices) shows very long periods of high volatility interspersed with periods of relative calm. This pattern is often referred to as volatility clustering.
59. Finally, in the B-S model interest rates are assumed to be constant and known. The risk-free rate is used to represent this rate. In the real world, there is no such a thing as risk-free rate but it is possible to approximate it by using, for example, Treasury Bills with high credit quality.

However, these treasury rates do change over time, especially during periods of high volatility.

60. It is, therefore, widely accepted in the financial industry that the above identified limitations are important and it is desirable to come up with models that will take into consideration some of the assumptions not addressed by Black-Scholes models. There is considerable academic literature which proposes alternative models, all attempting to mimic the characteristics of the market more accurately. However, recent history shows that every aspect of the market cannot be considered in any given model, as every factor affecting the price of a financial security cannot be captured mathematically. Mathematical models, by definition, do, and can only attempt, to capture some of the aspects of market behaviour. The convenience of using of one or another type of model will depend upon the valuation circumstances.

Alternative stochastic processes

61. There are many alternatives that can be assumed to the B-S model. This paper only describes the most widely used variations. For the sake of clarity these alternatives are grouped in two categories:

- a) Jump-diffusion models: Models which allow asset prices to exhibit jumps
- b) Alternative Diffusion models: Models where asset prices change continuously but assume a process different than GBM. These differences have to do mainly with the volatility assumptions.

Jump-Diffusion Models

62. Documentation from various empirical studies shows that the asset price evolution can be adequately described by pure diffusion-type processes (such as the B-S model) for most of the time, but from time to time larger jumps may occur and this cannot be captured by these models. Jump-diffusion models address this feature by causing the equity price to jump on a random basis.
63. In a jump-diffusion model a jump term is superimposed on the GBM. In these models prices change continuously, most of the time driven by Brownian diffusion term but they also take into account the fact that from time to time larger jumps may occur, driven by the jump term.
64. There is increasing interest in jump models in finance for the following reasons:
65. First, in a model with continuous paths like a diffusion model, the price process behaves locally like a Brownian motion and the probability that the stock moves by a large amount over a short period of time is very small, unless one fixes an unrealistically high value of volatility. Therefore, in such models the prices of short term options would differ from those observed in real markets.
66. Second, from the point of view of hedging, continuous models of stock price behaviour generally lead to a "complete market"; this is where derivative *payoffs* can be replicated by existing instruments. Since every terminal *payoff* can be exactly replicated through perfect hedging, options can be considered redundant assets, and the existence of traded options becomes a puzzle. This issue is easily solved by allowing for discontinuities: in real markets,

due to the presence of jumps in the prices, perfect hedging is impossible and options enable the market participants to hedge risks that cannot be hedged by using the underlying only.

67. From a risk management perspective, jumps allow the risk of strong stock price movements over short time intervals to be taken into account and quantified, which is not possible in the diffusion framework.
68. Finally, and probably the strongest argument for using discontinuous models is simply the presence of jumps in observed prices especially in small time scales, where prices do not change continuously but rather in discrete jumps in response to trading sentiment or significant new information.
69. When using a jump model an important consideration is the selection of the appropriate jumps. The options are numerous: such as whether to use downward jumps only, or upward jumps as well; many small jumps, or infrequent large jumps; whether the jumps are of a *deterministic* size or they arise from a distribution; or whether combinations of jumps are appropriate. Different choices or combinations of the Poisson processes (which is a type of continuous *stochastic* probabilistic process) will generate those different patterns in the jump side.
70. As mentioned, jump-diffusion models are appropriate where it is inappropriate to assume a smooth diffusion of equity prices, so they are particularly suitable for pricing barrier options and options with short time maturities, but less so for products with a long time horizon.
71. It is important to note that because there are at least two sources of randomness, there is no longer a “complete market” and the risk cannot be perfectly hedged. Consequently, in contrast to the B-S framework, the principle of an absence of arbitrage does not lead to a uniquely defined value of the option. An entire range of values may be obtained and the risk preference structure of the investor has to come into play. Although there are several approaches to overcoming this difficulty, all of them introduce serious additional complexity to the valuation process.

Alternative diffusion models

72. One alternative to the B-S model is to keep the assumption that asset prices change continuously but assume a process other than GBM. Innovation in this area has to do mainly with modelling volatility.
73. In the B-S model, volatility remains constant. Over time, empirical data shows constant volatility to be inconsistent with market behaviour. This has led to the development of dynamic volatility modelling. Volatility modelling may be classified in three categories:
 - Time dependent volatility,
 - Local volatility,
 - *Stochastic* volatility.


Time dependent volatility models:

74. It has been empirically observed that implied volatility varies with an option’s expiration date and strike. Consequently, a straight forward extension proposed to the constant volatility

model was time dependent volatility modelling, where volatility is a function of time: $\sigma = \sigma(t)$. Merton (1973) was the first to propose a formula for pricing options enabling time dependent volatility to account for empirically observed implied volatility increasing with time

Local volatility models:

75. In explaining the empirical characteristics of volatility, a time dependent volatility can be found insufficient as it does not explain the volatility “smile”. A graph of implied volatility vs. *strike price* of an option shows that implied volatility is not a constant function of the strike.



76. Therefore, to address this issue, a local volatility model has been developed to express volatility as a function of price and time where $\sigma = \sigma(S,t)$, for, S, the price of the underlying asset.
77. The advantages of local volatility models are that, first, they are able to account for a greater degree of empirical observations, and theoretical arguments than time dependent and constant volatility models. Second, they can still be calibrated to fit empirically observed data, enabling consistent pricing of derivatives. And, third, no additional source of randomness is introduced into the model. Hence it is theoretically possible to perfectly hedge claims and, thus, a unique value for the option exists.
78. The first local volatility model came into existence when Dupire (1994) showed that, in the presence of volatility skews, consistent models can be built in a risk neutral world if in the asset price process volatility is assumed to be a *deterministic* function of time and the asset price. Dupire derived a formula which allows the local volatility surface to be extracted from the prices of traded call options.
79. Other examples of local volatility models are the Constant Elasticity of Variance Model (CEV) where the process for the stock price is such that, depending on its specification, it allows the volatility of the price to increase or decrease as the stock price moves; Mixture Distribution Models and Implied Local Volatility Models.

Stochastic volatility models:

80. Although local volatility models are an improvement on time dependent volatility, they possessed certain undesirable properties. For example, volatility is perfectly correlated (positively or negatively) with stock price yet empirical observations suggest no perfect correlation exists. Stock prices empirically exhibit volatility clustering but under local volatility this does not necessarily occur. Consequently after local volatility development, models were proposed that allowed volatility to be governed by its own *stochastic* process

81. The key advantages of *stochastic* volatility models are that they capture a richer set of empirical characteristics compared to other volatility models. *Stochastic* volatility models generate return distributions similar to that which is empirically observed. For example, the return distribution has a fatter left tail and “peakedness” compared to normal distributions, with tail asymmetry controlled by the correlation between the spot *stochastic* process and the volatility process. Also, historic volatility may show significantly higher variability in periods of market turmoil than would be expected from local or time dependent volatility, which could be better explained by a *stochastic* process.
82. The disadvantages of *stochastic* volatility are, firstly, that these models, like jump-diffusion models, introduce a non-tradable source of randomness. Hence the market is no longer complete and we can no longer uniquely price options. Therefore, the practical applications of *stochastic* volatility are limited. Secondly, *stochastic* volatility models tend to be analytically less tractable. In fact, it is common for *stochastic* volatility models to have no closed form solutions for option prices. Consequently option prices can either be calculated by simulation or by numerical expansions.
83. There is no generally accepted *stochastic* volatility model and a large number exist. Among most significant ones are the Johnson and Shanno Model, Scott Model, Stein and Stein Model, Heston Model and SABR Model.

The Greeks

84. The different types of sensitivity of a derivative’s value to small changes in the underlying factors are generally annotated by letters from the Greek alphabet, and known collectively as the “Greeks”. The most common Greeks are:
 - Delta - value sensitivity with respect to spot. This is the amount of the underlying equity required to hedge a derivative position.
 - Gamma - sensitivity of the delta to spot. This is a measure of how often a position must be re-hedged to maintain a delta-neutral position.
 - Theta – value sensitivity to time.
 - Rho - How the value changes with linear shifts in the interest rate.
 - Vega - sensitivity of the value to volatility.

Model resolution

85. There are a large number of numerical techniques available for pricing which each have their benefits and drawbacks, and may be limited to particular model implementation and product combinations.
 - Analytic solution - without question the preferred approach for pricing, as these take no time to price and give no numerical error. Unfortunately, there are very few model/product combinations for which closed formulas exist and, therefore, that can be priced using these methods.
 - Monte Carlo - the most frequently used numerical approach. Random numbers are generated and used to construct a possible path of the underlying asset (and therefore the *payoff*) according to the model used. This is repeated numerous times

and an average is taken across all of the *payoffs*, giving the expected *payoff*, which is discounted to give the price. This is computationally intensive and can be slow. However the increase in computational time is linear with the increase in dimensions (e.g. multiple assets), which, in terms of computational time, means that the more assets are analysed the faster the Monte Carlo performs relative to other methods. Unfortunately, the resulting average of a number of random paths is also a random number, so the result will always contain an error. However the likely order of this error can be ascertained and reduced by increasing the number of paths, but it will always remain. There are other techniques that may be used to reduce the error, but these are difficult to apply generally across different products. Computing the Greeks can be time-consuming, and early-exercise features are tricky to model well. An advantage of this method, from a programming perspective, is that it is straightforward to plug in alternative *payoffs* or contracts, meaning that adding to code libraries is simplified.

- Trees - the future path of an underlying asset is modelled using a recombining tree, using discrete upward and downward steps over short time periods to approximate a continuous path. Probabilities are assigned to the up and down movements. The terminal price of the option is known, given the underlying asset price at *maturity*, and therefore it is possible to work back down the tree using the probabilities to obtain the current price of the contract. These methods are good for pricing early exercise features, but finite difference schemes tend to be superior in most respects.
- Finite difference - uses a rectangular grid of specified points of asset price and time to model prices. This yields more stable, better behaved pricing than a tree, with the similar ability to price early-exercise contracts and products such as barriers. These schemes can be very fast, and accuracy can be controlled in a *deterministic* manner, meaning extrapolation techniques can be employed to obtain more accurate prices. Evaluating the Greeks for hedging and risk management is similarly efficient. All of these benefits come at a price, which is limited dimensionality; practically speaking it is unfeasible above three dimensions (two asset processes plus time). Finite difference models are often used for vanilla American options but are less suitable for other exotic products such as *path-dependent* options.
- Numerical integration - this covers two approaches that work in a similar manner, which are considered briefly:
 - Fourier/Laplace transform approaches, - for certain *stochastic* volatility models this is invaluable, as the pricing of vanilla instruments can be reduced to a numerical integration, which is very fast, accurate, and essential for calibration.
 - Quadrature methods - directly evaluate derivative values by integrating probability distributions against *payoffs*. Much like trees and finite difference schemes this is used primarily for low dimensionality problems, and for early exercise.

Implementation

86. There are a vast selection of implementation approaches available, the choice of which depends entirely on the institution's remit, resources and the end user.
87. Initial testing and model development may be performed using high-level languages and software, as these have ready-to-use libraries of numerical tools, and debugging is normally

straightforward. The disadvantage of these methods is that they are slow and therefore generally not usable on a release or portfolio.

88. For release code, a low-level programming language can be used to build a pricing library. These languages do not have much standard functionality so they are harder to program but are faster at the execution level.

Calibration

89. The calibration of models involves ascertaining values for the various input parameters given liquidly traded instrument prices. These instruments are predominantly simple derivatives - such as vanilla calls or puts, but may be more complex if the model in question can capture additional features.

Risk-free Rate

90. The option's time-to-*maturity* risk free interest rate should typically be calculated through a bootstrapping process using market instruments such as deposits rates, forward rate agreements (FRAs) and interest rate swaps (IRS).

Volatility

91. The volatility of a stock is a measure of the variability (or risk) of returns. The volatility cannot be directly observed, it is instead calculated through historical data analysis or an implied via a model. As implied volatility represents market expectations about future volatility, this is generally the preferred option, although is not always possible to compute.

Correlations

92. The calibration of correlations will depend on which equities are being priced. There is a growing market in correlation and covariance trading, which is most developed for indices versus their constituents. Contracts in these markets monetise the correlation from which the correlation value can be implied.
93. When market-implied correlations are unavailable, historical correlations may be used. Unfortunately this does not provide an implicit hedging tool and often trading desks are exposed to this risk, though they may well hedge this on a macro level.

Dividends

94. To estimate stock *dividends* a distinction must be made between the short and long term. In the short term, as companies usually have a stable *dividend* policy, it can be assumed that the amount and timing of the *dividend* can be predicted.
95. In the long term, it is preferable to estimate *dividend* yields implicitly on the basis of market quotations. For example, implicit *dividend* yields can be estimated using the put-call parity relationship for European options or by using *dividend* swap quotes.
96. In the absence of market quotes historical *dividend* yields can be considered.

Conclusion

97. This TIP provides an overview of some of the different types of equity derivative products and of the principles on which some of the more commonly used mathematical models used for their valuation are based. It will be seen that different methods and models have strengths and weaknesses in different situations. Judgement is always required in choosing the most appropriate method or model based on the facts and circumstances and over reliance on one should be avoided. Particular care is required to ensure that the assumptions on which any chosen model is based are as realistic as possible having regard not only to the nature of the product but also to the prevailing market conditions and the purpose for which the valuation is required.
98. The following provisions of the *IVS Framework* are particularly relevant in this context:
- The exercise of judgement and objectivity,
 - The expectation of competence,
 - The need to consider the use of more than one valuation approach or method.

Closed form solutions under the Black-Scholes framework

The most popular Black Scholes model analytical formulas correspond to European call and put options. Nevertheless, these model assumptions have been considered in the derivation of closed-form pricing formulas for exotic options on one or two underlying assets, as the following table summarizes.

| Option Type | # Assets | Model |
|---|----------|--|
| Forward Start Option | 1 | Rubinstein (1990) |
| Chooser Option | 1 | Rubinstein (1991) |
| Option on Option | 1 | Geske (1977, 1979); Selby (1987); Rubinstein (1991) |
| LookBack Option | 1 | Goldman, Sosin and Gatto (1979); Conze and Viswanathan (1991); Heynen and Kat (1994) |
| Mirror Option | 1 | Manzano (2001) |
| Barrier Option | 1 | Merton (1973); Reiner and Rubinstein (1991) |
| Binary Option | 1 | Reiner and Rubinstein (1991); Cox and Rubinstein (1985) |
| Asian Option | 1 | Kemna and Vorst (1990) |
| Exchange-One-Asset-For-Another Option | 2 | Margrabe (1978) |
| Two-Asset Barrier Option | 2 | Heynen and Kat (1994) |
| Relative Outperformance Option | 2 | Derman (1992); Zhang (1998) |
| Options On the Maximum or the Minimum of Two Risky Assets | 2 | Stulz (1982) |