

The Derivatives Academy

Maxime de Bellefroid

2022-10-31

Contents

1	Prerequisites	9
2	Introduction to the Structured Products Market	11
2.1	The Structured Products	11
2.2	The Stakeholders	13
3	Back to Basics!	19
3.1	Interest Rates	19
3.2	Bonds	21
3.3	Equities	23
3.4	Forwards and Futures	24
3.5	Swaps	26
3.6	Options	28
4	A deeper understanding of Options	31
4.1	The Black-Scholes model	31
4.2	European Call Options	33
4.3	European Put Options	38
4.4	Intrinsic Value and Time Value	41
4.5	The Cost of Hedging	43
4.6	The Call-Put Parity	44

5	The Greeks	47
5.1	Introduction to the Greeks	47
5.2	Delta	49
5.3	Gamma	58
5.4	Theta	67
5.5	Vega	74
5.6	Rho	78
5.7	General Practical Example	82
5.8	Second-Order Greeks	84
5.9	Multi-Asset Greeks	86
5.10	Computational Methods	87
6	All about Volatility	91
6.1	The many flavours of volatility	91
6.2	The Volatility Surface	95
6.3	Review of Volatility Models	106
7	Classic Options	123
7.1	European Options	123
7.2	American Options	123
7.3	Digital Options	124
8	Options Strategies	135
8.1	Call Spreads	135
8.2	Put Spreads	135
8.3	Straddles	135
8.4	Strangles	135
8.5	Risk Reversals	135
8.6	Butterfly Spreads	135
8.7	(Iron) Condor Spreads	135

9 Asian Options	137
9.1 Asian-In and Asian-Out	137
9.2 Geometric Average vs Arithmetic Average	138
9.3 Numerical Application	139
9.4 Risk Analysis : The Greeks	139
10 Quanto Options	143
10.1 Description	143
10.2 Additional sensitivities : correlation risk and FX volatility	144
10.3 Hedging FX Exposure	146
11 Compo Options	147
11.1 Description	147
11.2 Additional sensitivities : correlation risk and FX volatility	148
11.3 Hedging FX Exposure	148
12 Barrier Options	151
12.1 Description	151
12.2 Knock-Out Options	152
12.3 Knock-In Options	153
12.4 In-Out Parity for barrier options	154
12.5 Review of Payoffs	154
12.6 Deeper into Risk Analysis	154
12.7 Double Barrier Options	161
13 (Barrier) Reverse Convertibles	163
13.1 Reverse Convertibles	163
13.2 Barrier Reverse Convertibles	164
14 Certificates	167
14.1 General Description	167
14.2 Tracker Certificates	167
14.3 Discount Certificates	169

14.4 Bonus Certificates	171
14.5 Capped Bonus Certificates	173
14.6 Airbag Certificates	174
14.7 Outperformance Certificates	175
14.8 Outperformance Bonus Certificates	176
14.9 Twin-Win Certificates	178
15 Multi-Asset Options	181
15.1 Introduction	181
15.2 Correlation	181
15.3 Basket Options	185
15.4 Rainbow Options	187
15.5 Worst-Of options	188
16 Autocallables	193
16.1 Description	193
16.2 Payoff	194
16.3 Risk Analysis	194
16.4 Memory Feature	196
16.5 Put Feature	197
16.6 Smoothing the autocall barrier	198
16.7 Multi-Asset: Worst-Of	200
16.8 Other variants	200
17 Variance Swaps	203
17.1 Introduction	203
17.2 Liquidity	203
17.3 Uses of Variance Swaps	204
17.4 Mechanics	205
17.5 The Contract	206
17.6 Convexity	207
17.7 Mark-to-Market	208

17.8 Forward Variance	208
17.9 Contract Specifications	209
17.10 Observations from historical prices	210
17.11 Variance Swaps and Option Implied Volatilities	211
17.12 Pricing Rules of Thumb	211
17.13 Drivers of Variance Swaps levels	213
17.14 Are variance swaps good predictor of future volatility?	213
17.15 Is Variance Swap Convexity fairly priced?	214
17.16 Variance Term Structure	214
17.17 From Options to Variance Swaps	215
17.18 Sensitivity to Skew and Convexity	219
17.19 Greeks	221
17.20 Setting up a replicating portfolio	224
17.21 Replicating and Hedging in Practice	225
17.22 Effects of Variance Swaps Hedging	226
17.23 Why not Volatility Swaps?	227
17.24 Third Generation Products	228

Chapter 1

Prerequisites

This material aims at serving the Equity Derivatives community by sharing knowledge and thoughts about different topics in this field. More specifically, it aims at providing graduates with semi-practical knowledge through the use of a derivatives pricer that I developed for pedagogical purposes.

This material is made of notes taken from my readings and my personal experience as an Equity Derivatives Structurer. I have tried to provide links of the books so that you can read them if you want to deepen your knowledge. As mentioned in the material, some parts and examples are entirely taken from these reference books and papers.

I added valuable insights that I am sure you could take advantage of by reading this material attentively. But the real added-value clearly comes from the Exotic Options pricer made available throughout this book and that has a dedicated section on the website. Using the pricer, you can develop a much better intuition about Equity Derivatives. The list of products that the pricer allows you to price has been carefully chosen and is already quite large. Depending on the server capacity, I might add some more exotic structures in the future.

I believe mastering this material is necessary to secure an internship or entry-level position in the field. Necessary but not sufficient as you will also have to show interest and knowledge on financial markets...

Everything will be written the way I speak as I do not intend to publish it anywhere else than on my own website. This is a perpetual work in progress book as improvements and new subjects will come along the way. If you want to contribute, feel free to reach out to me. If you see any mistakes, please do so as well ;).

As I have worked on structured products, they are the main focus of this book. Therefore, I will start speaking shortly about the world of structured products.

Chapter 2

Introduction to the Structured Products Market

2.1 The Structured Products

2.1.1 Introduction

It all started back in the early 90's. Investment banks were looking for ways to attract investors into the Equity markets.

What better ways than providing innovative solutions with advanced payoffs on various assets?

This is what structured products offer. They allow investors to have access to innovative payoffs through different issuing wrappers. By doing so, they provide an excellent solution to invest in a product with a tailor-made risk-return profile in a cost-effective manner.

As structured products are sold over-the-counter (OTC), there must exist a sound legal contract between the two parties. This contractual agreement explaining the features of the structured product is referred to as a Term Sheet (TS).

For the OTC business of structured products to expand, the development of a secondary market was necessary to provide liquidity. The banking system has always been based on confidence. Investors will be more inclined to invest in structured products if they know they can partially/fully unwind their positions easily in the future.

To offer such liquidity, traders estimate the market value of structured products based on the current market data and make bid-ask quotes available to investors. Usually, notes trade in the secondary market with a 1% bid-ask spread. If the traders are struggling in risk-managing a specific product, they might increase their bid-offer spreads (except if legally defined otherwise in the term sheet).

Banks usually have valuation teams that help the salespeople to provide their clients with quotes.

This is more important than it looks! Clients may want to decrease/increase their exposure by selling/buying at the bank's bid/offer price. The fair price of the structured product lies somewhere between the bid price and the offer price but it does not have to be right in the middle. Depending on the estimated probability of the clients to increase/decrease their exposure, the valuation team will shift the bid-ask spread. Therefore the fair price may be closer to the bid price if the valuation team expects the client to increase his exposure. In this manner, the bank would make more money when the client buys at the offer price. In the same way, the fair price may be closer to the offer price if the valuation team expects the client to decrease his exposure. In this manner, the bank would make more money when the clients sell at the bid price.

The valuation job is not a cakewalk. Some illiquid parameters can hardly be implied and will need to be monitored closely. These parameters can considerably impact the whole trading book. Some market consensus data providers such as Totem can be used to mark some illiquid parameters at some market consensus between institutions active in the market.

2.1.2 Issuing Wrappers

As stated, structured products can be launched in different wrappers that have their particular legal status.

I am not an expert on the regulations and laws around the different wrappers and will only speak about their financial side. In the same way, I will only cover the wrappers I have been used to deal with.

2.1.2.1 Warrants and OTC Options

The warrant/option's holder is granted the right to buy a specific amount of shares in a company at an agreed price.

2.1.2.2 Notes and bonds

They are debt instruments issued by the banks usually as senior unsecured debt that can be listed or not.

Investing in such notes is not risk-free since the investor still faces the risk that the issuer of the note defaults. Furthermore, the investor could be selling an option inside a note, putting his capital at risk.

Example of a typical structured note The structured note is composed of:

- a **non-risky asset** providing a percentage of protected capital
- a **risky asset** that adjusts the risk-return profile.

The “non-risky” part is typically a zero-coupon bond (ZCB). When saying “non-risky” here, I mean that the capital is guaranteed as long as the issuer does not default. As a ZCB is bought as a discount, you can consider its value to increase ‘relatively’ linearly through the life of the product to amount to 100% of the notional at maturity.

The initial price of the non-risky part is linked to the level of interest rates. Higher interest rates decrease the initial value of the ZCB and increase the amount available to the structurer to create an attractive payoff enabling the investor to increase his upside in the equity.

The risky part is typically composed of options on single or multiple assets. These options will enable the investor to make additional profit using leverage. This possibility of additional return comes with additional risk as their value can fluctuate considerably.

As we will see, options’ price is convex and subject to many market parameter.

2.2 The Stakeholders

As in every market, the stakeholders can be classified into two categories: the sellers and the buyers.

2.2.1 The Sell Side

2.2.1.1 Sales

The salespeople are in charge of selling structured products to existing or new clients. Useless to say that they get commission for the transactions they settle so that their remuneration depends on:

- The number of deals
- The size of the deals

- The nature of the deals

The size of transactions is important for every sell-side stakeholders.

For the salespeople, their commission on a transaction is a percentage on the deal's notional. They might therefore be tempted to decrease their commission margin to increase their chances to settle the transaction and get a large dollar amount of P&L. Furthermore, there is less costs involved in terms of hours spent on the deal.

For traders and structurers, the size of transactions is important in the valuation of the structured products. The larger a deal with unhedgeable risks, the larger the risks for the traders to manage, the more dispersion in the prices offered to the investor.

Selling structured products is not like selling shoes. Before issuing them, salespeople will spend quite some time with potential buyers reviewing the payoff mechanism and potential return. They will use a large panel of marketing methods to explain the product's risk-return profile: back-testing, stress-testing, etc..

Their role is often underestimated as they usually have a slightly less technical background than the other stakeholders on a trading floor. I personally have a lot of esteem for them. They must understand a large panel of complex products, demonstrate strong adaptability and be able to build relationships of trust with their clients.

They must also be strategist and know their clients. For example, in competitive auctions, they must sometimes accept not to charge much to get the deal and be able to charge smartly on the secondary-market depending on their clients' habits to unwind or increase their positions.

An excellent salesperson is a rare thing and is a pleasure to work with for a structurer.

2.2.1.2 Structurers

Their role involves pricing existing structures and creating new ones. When creating new products, the structurer must be innovative and creative.

This innovation is a collaborative innovation as the structures must have a close interaction with salespeople to understand what the investors are willing to buy.

The capability of innovating is a key asset for an investment bank to stand out in the competitive market of structured products.

When pricing structured products, the structurer analyses all the risks. He also works closely with traders to agree on the levels they will charge for taking those risks.

2.2.1.3 Traders

Structurers must discuss and negotiate with traders on the value of the market parameters used to price the products. This is because, after a deal has been closed, the product is booked in the trader's portfolio. The trader will therefore be in charged of hedging the products' risks.

An exotic trader in an investment bank is therefore more a hedger than a speculator. They usually have a more dynamic view of the products than the other stakeholders as they will have to manage the risks during the entire life of the products. Those risks are quite elaborated and it is not always possible to hedge them completely. The trader will manage them as best as possible within the specified limits set by the risk department.

2.2.1.4 The balance of power

It is important to understand the balance of power at play on a trading floor with different stakeholders having diverging interests.

Sales are focus on increasing their commissions as much as possible and therefore closing as many deals as possible. Traders also want business since they want to increase the size of their books but not at all cost. They will not want to take unhedgeable risks except if the risk is well rewarded, otherwise they might end up losing money on such positions.

Sales will tend to be aggressive and traders defensive. Structurers stand in the middle trying to reconcile at best these diverging interests.

2.2.2 The Buy Side

Buy side clients can be classified into two categories: retail and institutional.

2.2.2.1 Retail Investors

Most of them are asset management firms that invest in structured products to redistribute them to individual end clients.

The payoffs are highly simplified and marketed accordingly.

The process usually works as follow. Retail investors request indicative prices from several banks before moving to a live auction. Based on the indicative prices received, retail investors will select the most competitive counterparties to participate at the live auction.

At the live auction, they will request the selected counterparties to update their indicative prices into tradeable prices. Those live prices might be similar or different than the indicative ones based on two things:

- the evolution of the market parameters between the indicative request and the live auction.
- the aggressivity of the bank might have changed based on the client's feedback after the indicative prices.

Retail investors typically select the bank that offers the most competitive price but may also spread a large notional over several investment banks.

Spreading a large notional over several banks enables the investor not to be fully dependent on the bid-ask of a unique bank if he wants to partly/fully unwind his position. In the same way, it will also allow him to spread the credit risk of having a unique counterparty.

The categories of yield enhancement products are the most popular among retail investors.

2.2.2.2 Institutional Investors

They include financial institutions such as hedge funds or mutual funds.

Transactions with institutional investors tend to be larger in terms of notional sizes and more sophisticated in terms of payoffs than with retail investors.

Solutions offered to institutional investors are business-tailored. As a result, it is less competitive than the retail business but the development of the solution is a much larger part of the job.

As payoffs can be more sophisticated, the room for innovation is greater.

2.2.2.3 A word on credit risk

I said above that the investors will typically select the bank that offers the most competitive price for the structured product. It is not always the case as the buyer of a structured note will also pay careful attention to the seller's credit rating. An investor might decide to trade a product with a better credit-rated issuer even if he finds a more attractive price with a lower credit-rated issuer.

The price difference between several issuers can be quite substantial. This is not surprising as the zero-coupon bonds of low credit-rated issuers are worth less. Indeed the rate used to price the zero-coupon bond part of the structured note can be seen as a reference rate plus some rate that the bank's treasury offers on deposits. This funding rate varies greatly based on the issuer's credit-rating. The lower the issuer's credit-rating, the higher this rate, the lower the value of the issuer's zero-coupon bond. If the non-risky part costs less, it means more money remaining to spend in the risky part of the note. So the riskier issuer can potentially offer a more attractive payoff for the same price or a more attractive price for the same payoff.

Investors should know about the impacts of credit ratings on the structured note investments.

When markets are falling and volatility getting higher, it is quite usual to observe a flight for quality, meaning that investors favour the creditworthiness of their counterparties rather than their price competitiveness.

The credit risk from the structured note's issuance can be hedged using credit default swaps (CDS).

A few words can also be said about collateralized line investments that do not involve counterparty risk. When considering swaps, you can structure the swap along with collateralized lines with a third party to avoid problems in case of default events. It might sound complicated at first but it is not. It is simply about computing the value of the swap and setting aside the equivalent amount as collateral with a third party. This comes at the investor's cost as it mitigates his counterparty risk. I might adress a small chapter of these notes on credit-related adjustments in derivatives valuations even though it is not my domain of expertise.

Chapter 3

Back to Basics!

This chapter aims at reviewing superficially the different asset classes. It is just a quick and dirty reminder. Clearly, without a financial background, it surely won't be enough to understand the rest of the material.

3.1 Interest Rates

3.1.1 Introduction

Interest rates represent the amount charged by a lender to a borrower for the use of assets. The amount of money depends on several factors including the credit risk, which is the risk of loss due to the non-payment of the borrower's duty.

Interest rates are involved to a very large extent in the pricing of all derivatives.

For any given currency, you will find many types of rates. It is essential to apprehend the differences between them. Which rate to use to price this specific financial instrument? What impact in terms of valuation?

That's the difference between practice and theory. In theory, you just use 'r' in your model as being the risk-free interest rate. In practice, you have to ask yourself which specific rate to use as a good proxy for the risk-free interest rate for this particular complex product? And the answer to this question is not always trivial...

Whatever the maturity, interest rates are typically expressed as annual rates, allowing them to be compared easily.

3.1.2 LIBOR and Treasury rates

Treasury rates

Treasury rates are the rates earned on debt instruments issued by governments. Regulatory issues can impact the value of Treasury rates and cause them to be persistently low. Accordingly, derivatives traders rather use LIBOR as a better proxy for short-term risk-free rates.

LIBOR

The London Interbank Offered Rate (LIBOR) is the interest rate at which a bank offers to lend funds to other banks in the interbank market. Depending on the length of deposits, LIBOR rates come in different maturities (overnight, 1w, 1m, 2m, 3m, 6m, 12m) and are associated with all major currencies (EUR, USD, GBP, JPY, CHF).

TED Spread

The TED Spread is the difference between 3-month Treasury Bills and 3-month LIBOR. It is often used as a measure of liquidity in the interbank market. Unlike Treasury rates, LIBOR rates involve some credit risk. Therefore the TED spread serves as a measure of credit risk in the interbank market.

3.1.3 Yield curves

Traders closely watch interest rates from different financial instruments such as bonds, swaps and futures. These interest rates are usually plot against their maturities to form what is called a yield curve. Well, you do not know interest rates for all maturities. You only know interest rates with certainty for specific maturities. You will have to use some interpolation methods to build the rest of the yield curve.

If you are asked about yield curve construction, you can speak about some cubic spline interpolation methods but try not to go too simple with a linear interpolation that would never be used in practice.

Again, there are yield curves for any major currencies.

Yield curves are typically upwards sloping, meaning that longer-term rates are higher than shorter-term rates. Under specific market scenarios, yield curves could also be flat or even downward sloping.

3.2 Bonds

3.2.1 Introduction

A bond is a fixed income instrument that represents a loan made by an investor to a borrower. Bonds are used by governments and companies to raise capital. By lending money, the bond's holder is entitled to receive periodic coupons as well as the initial investment back at the bond's maturity.

The coupon rate can be fixed or floating. In the first case, the coupon rate is constant throughout the life of the bond while, in the second case, the coupon rate is linked to another index.

Bonds are usually categorized based on their maturities:

- **Short-term:** maturity less than 2-3 years
- **Medium-term:** maturity between 3-4 years and 10 years
- **Long-term:** maturity greater than 10 years

As you can see, this is theoretical and subjective. I might speak about a 2-year bond as a short-term bond while someone else might see it as a medium-term bond. It does not matter. What is relevant for you is to have a sense of what range of maturities are considered as short, medium or long.

3.2.2 Market Price

As any other financial asset, the market price of a bond is equal to the sum of the present values of its expected cashflows.

$$Bond(t, T) = \sum_{i=1}^n C_i e^{-r_{(t, t_i)}(t_i - t)}$$

Where

- C_i $i = 1, \dots, n-1$ being the coupons paid
- C_n being the last coupon + principal
- $r_{(t, t_i)}$ being the annual interest rate for the period (t, t_i)
- $(t_i - t)$ being the number of years in the period (t, t_i)

Dirty Price

The market price of a bond may include the interest that has accrued since the last coupon date, in which case it is called the dirty price and corresponds to the fair value of a bond as it is the price effectively paid for the bond.

Clean Price

Many bond markets quote bonds as clean prices and add accrued interest on explicitly after trading.

3.2.3 Bonds' underlying risks

Generally safer and more liquid than stocks

Multiple reasons can be put forward to explain why bonds are generally considered to be a safer investment than stocks. Firstly, bonds are senior to stocks in the firm's capital structure so that bondholders receive money first in case of an event of default. Then, bonds are generally more liquid than stocks. All in all, bonds are less volatile than stocks and can be used to lower a portfolio's volatility in times of high volatility in the stock market.

But not risk-free

As we can see in the above bond price formula, bond prices are a direct function of interest rates. Since interest rates can vary greatly during the bond's life, bonds are clearly not risk-free.

Interest Rate risk

As a consequence, bonds are subject to interest rate risk. Since the interest rates are used to discount the bond's coupons, an increase in their value lowers the coupons' present value, and therefore the bond's price. The opposite applies in case interest rates go down, in which case the coupons' present value increase as well as the bond prices.

Credit risk

Moreover, bond prices depend on the issuer's creditworthiness, typically summarized by a rating given by a credit rating agency (mainly S&P, Moody's, Fitch). The higher the credit rating, the safer the bond's issuance, the less the interest rate required by the investor to bear the credit risk, the higher the bond's price.

Inflation risk

Inflation deteriorates the returns associated with bonds. This is particularly true for fixed bonds, which have a set interest rate from inception.

Inflation risk is very insidious as you cannot really see it. You still receive your coupons and principal. Everything seems intact but it really is not. Your purchasing power is suffering slowly but surely.

3.2.4 Zero-Coupon bonds (ZCB)

ZCB are debt instruments where the lender receives back a principal amount plus interest, only at maturity. No coupon is therefore paid during the bond's life.

Sold at a discount

ZCB are sold at a discount, meaning that its price is lower than 100% of the notional. This is because the interest is deducted up front.

Since no coupon is paid, a ZCB price is nothing but the present value of the par value paid at maturity.

Using continuous compounding to discount cash flows, the price of a ZCB can be expressed as:

$$B_{(t,T)} = e^{-r_{(t,T)}(T-t)}$$

3.3 Equities

Companies need cash to operate or finance new projects. They can choose to raise capital by issuing equity.

3.3.1 Dividends

Companies usually pay their shareholders dividends. Dividends can vary over time depending on the company's performance and strategy. Dividends can be expressed as discrete dividends or as a continuous equivalent dividend yield, represent by the symbol 'q'.

For the sake of consistency, I will always use continuous yields for all parameters: interest rates, dividends, borrowing rates, etc..

3.3.2 Repurchase Agreement (Repo)

If you believe a specific stock price will go down over time, you would be willing to sell it, right?

But how can you sell something if you don't own it first?

Well, you can enter into a repurchase agreement, that is a transaction in which you borrow the stock from another counterparty that actually holds the stock and you agree to give it back at a specific date in the future.

It will allow you to hold the stock and sell it directly. If your intuition about the future stock price going down happens to be correct, you would be able to buy the stock back later at a cheaper price and return it to the lending counterparty, realizing a profit.

Why would any investor lend me their stocks?

Some investors do not plan to trade in this stock for a while as they are long-term investors with a buy & hold type of strategy. Repos will benefit these investors by allowing them to earn an additional income paid by the stock borrowers. The borrowing costs are called the repo rate, represented by the symbol 'b'.

3.3.3 Liquidity

When trading stocks, an investor should also be vigilant with their liquidity, which is usually quantified by their average daily traded volumes. A stock is said to be liquid if an investor can buy and sell it without moving its price in the market.

Liquidity risk occurs when an investor wants to close his position but is unable to do it quickly enough in the market without impacting the market price. Let say you have a large long position in a stock and want to exit it by selling the stocks, you might not be able to find a buyer quickly enough so that you would have to sell at a lower price than the fair price for the transaction to be conducted. Consequently, you might not be able to make a profit from your investment.

The most popular and crudest measure of liquidity is the bid-ask spread. A narrow bid-ask spread tends to reflect a more liquid market.

3.4 Forwards and Futures

3.4.1 Introduction

A forward contract is an agreement to buy or sell a security at a certain future time for a certain price. The buyer of the forward agrees to pay the price and take delivery of the underlying asset at the pre-determined price on the agreed date.

Forwards are OTC products and are normally not traded on exchanges. However, futures with standardized features are traded on exchanges.

Note that forwards and futures are an obligation and not an option to buy/sell a security at maturity.

3.4.2 Delivery price, Forward price and Forward value

Delivery price

The price specified in a forward contract is called the delivery price. A forward contract is settled at maturity, when the holder of short position delivers the security to the holder of the long position in return for cash amount equal to the delivery price.

Forward price

The forward price is defined as the delivery price which would make the forward contract to have zero value.

Forward value

The value of a forward contract is determined by the market price of the underlying security.

At inception, the forward and delivery prices are equal and the forward contract has zero value. In other words, it initially costs nothing to enter into a forward contract.

As time passes, the underlying security price changes, causing the forward price and therefore the forward value to change.

$$\boxed{Forward_t(T) = (F_t(T) - K) e^{-r(T-t)}}$$

Where

- $Forward_t(T)$ is the value at time t of the forward contract
- $F_t(T)$ is the asset forward price
- K is the pre-determined delivery price
- $(t_i - t)$ being the number of years in the period (t, t_i)

Let's focus on $F_t(T)$ in the case of a stock as I will largely focus on Equity derivatives in these notes.

3.4.3 Forward price of a stock

The forward price of a stock is defined as the fair value of the stock at a specific point of time in the future. It can be viewed as equal to the spot price plus the cost of carrying it.

Impact of interest rates

Interest rate increases the cost of carry since the stockholder could have received the interest if he had immediately sold his shares and invested his money in a risk-free investment.

Therefore, the higher the interest rates, the higher the forward price.

Impact of dividends and repo

If a stock provides an additional income to the stockholder, this causes the cost of carry to decrease, since the stock also becomes a source of profit. Dividends and stock loans (repos) constitute a source of income when carrying a stock.

Therefore, the higher the dividend yield and the repo rate, the lower the forward price.

The forward price of a stock can be expressed mathematically as follow:

$$F_t(T) = S_0 e^{(r-q-b)T}$$

Where

- r is the interest rate (%)
- q is the stock's dividend yield (%)
- b is the stock's borrowing cost (%)

3.5 Swaps

3.5.1 Interest Rate Swaps (IRS)

Interest Rate Swaps (IRS) are over-the-counter (OTC) agreements between two counterparties to swap cashflows in the future.

Plain Vanilla Interest Rate Swaps

A plain vanilla interest rate swap is one in which two parties swap a fixed rate of interest and a floating rate. In an IRS, the notional does not actually change hands. Since all cashflows are in the same currency, payments are netted.

The swap's buyer is the party who agrees to pay the fixed rate and expects the interest rates to rise. The seller of the swap is the party who agrees to receive the fixed rate and expects the interest rates to fall.

Basis Swaps

A basis swap is an IRS where a floating rate is swapped for a different floating rate.

Value of a Swap

As usual, the swap's value is nothing but the net present value of all future cashflows, which is equal to the present value from the receiving leg minus the present value from the paying leg.

As with forwards, the terms of a swap contract are defined so that its value at inception is null.

3.5.2 Cross-Currency Swaps (CCS)

A CSS is a swap in which cash flows are based on different currencies. Unlike an interest rate swap, in a currency swap the notional changes hands both at inception and at the maturity of the swap. Since cashflows are in different currencies, interest payments are also made without netting.

In reality, market participants have different levels of access to funds in different currencies and therefore their funding costs are not always equal to LIBOR.

Funding Currency

An approach to work around this is to select one currency as the funding currency, and select one curve in this currency as the discount curve. Cashflows in the funding currency are discounted on this curve. Cashflows in any other currency are first swapped into the funding currency via a cross currency swap and then discounted.

This is something you will often do when working in structured products on a trading floor.

3.5.3 Total Return Swaps (TRS)

A TRS is a swap agreement in which a party pays fixed or floating interest and receives the total return of an asset. The total return is the sum of the capital gain/loss and any income received during the life of the swap.

The party that receives the total return obviously believes the asset will appreciate.

TRS are a good way to gain exposure to an asset without having to pay additional costs for holding it.

Equity Swaps

An equity swap is a particular type of total return swap where the underlying asset can either be an individual stock, a basket of stocks or a stock index.

Unlike stock, you do not have to pay anything up front when entering into an equity swap. Instead, you would usually deposit a margin on which you receive interest. It is therefore a good way to gain exposure to a stock without suffering additional transaction costs and local dividend taxes. It might also be a way of bypassing limitations on leverage.

3.5.4 Dividend Swaps

A dividend swap is an OTC derivative on a stock, a basket of stocks or a stock index in which two counterparties exchange cashflows based on the dividends paid by the equity underlying.

The buyer of the swap receives the dividends and pays the fixed payment. The seller of the swap obviously has the opposite position.

While dividend swaps can be used by investors to speculate on future dividends, I rather speak about them as an instrument for traders being long stocks to hedge their dividend risk.

3.6 Options

3.6.1 Introduction

Options are contracts that give their holder the right, but not the obligation, to either buy or sell an amount of some underlying asset at a pre-determined price at or before a pre-determined date in the future, the maturity date. The maturity of an option could be as short as a day or as long as a couple of years.

While the holder of an option has rights, the seller of an option has the obligation to take the opposite side of the trade if and when the owner exercises his right.

Rights do not come for free. You must therefore pay what is called the premium to buy an option.

3.6.2 Call Options

A call option gives its holder the right to buy an underlying asset, a stock for example, at a specific price per share within a specific time frame.

If you sell a call, you have the obligation to sell the stock at a specific price per share within a specific time frame if the call buyer decides to exercise his right to buy the stock at that price.

3.6.3 Put Options

A put option gives its holder the right to sell an underlying asset, a stock for example, at a specific price per share within a specific time frame.

If you sell a put, you have the obligation to buy the stock at a specific price per share within a specific time frame if the put buyer decides to exercise his right to sell the stock at that price.

Much of the time, individual calls and puts are not used as a standalone strategy. They can be combined with stock positions and other calls and puts based on the same stock to form more ‘complex’ strategies.

The next module further develops on these vanilla options.

Chapter 4

A deeper understanding of Options

This chapter aims at digging a bit more into European options and the model used to price them: the Black-Scholes model.

I have asked myself a lot of questions about the following chapters, trying to figure out where to start and how deep into mathematics should I go?

My goal is to get you to have the best understanding of options as possible. Mathematics is not always useful and necessary as you can get a great sense of options with a good intuition. I mainly rely on intuition when teaching this material.

As it might help the more quantitative profiles among the readers to dig into the mathematics behind the Black-Scholes model, I decided to share the derivation of the Black-Scholes Equation in Appendix of this chapter. For the others, there is surely enough maths in this chapter anyway.

4.1 The Black-Scholes model

Black and Scholes developed a closed-form pricing formula for European options. The market assumptions behind their model are quite strong and contained:

- Constant volatility
- Constant interest rates
- Log-normality distributed stock prices
- Constant dividend yield

The concept of risk-neutral pricing is at the heart of the Black-Scholes model.

4.1.1 Risk-Neutral Pricing

In finance, when pricing an asset, a common technique is to discount the expected cashflows. The expected cashflows are calculated using the real-world probability of each cashflow and the risk adjustment takes the form of a higher discount rate.

In the theory of risk-neutral pricing, the real-world probabilities assigned to future cash flows are irrelevant, and we must obtain what are known as risk-neutral probabilities.

The fundamental assumption behind risk-neutral valuation is to use a replicating portfolio of assets with known prices to remove any risk. By the no-arbitrage condition, we can create a pricing model by setting the temporary risk-less hedged portfolio return to the risk-free rate instead of unknown market expectations. The amounts of assets needed to hedge determine the risk-neutral probabilities.

Under the aforementioned assumptions, the Black-Scholes theory considers options to be redundant in the sense that one can replicate the payoff of a European option on a stock using the stock itself and risk-free bonds.

As such, the key feature of the Black-Scholes framework is that it is preference-free: since options can be replicated, their theoretical values do not depend upon investors' risk preferences, only upon the current stock price and its dynamics.

Therefore, an option can be valued as though the return on the underlying is riskless. This risk-neutral assumption behind the Black-Scholes model constitutes a great advantage in a trading environment. In the risk-neutral world, all cashflows can be discounted using the risk-free rate (r) whereas, in a real word, the discount rate should take into account the risk premium, which is more delicate.

A small example

Let me share with you a nice practical example that might help you to grasp the concept of risk-neutral valuation.

Suppose a first horse has 20% chance to win while a second horse has 80% chance to win.

10.000€ is bet on the first horse and 50.000€ is bet on the second.

If odds are set to 4:1, then:

- The bookmaker may gain 10.000€ if the first horse wins ($50.000€ - 4 * 10.000€$).
- The bookmaker may loose 2500€ if the second horse wins ($10.000€ - 1/4 * 50.000€$).

- The bookmaker's expected profit is 0 ($0.2 * 10.000\text{€} + 0.8 * -2500\text{€}$).

If odds are now set to 5:1, then the bookmaker won't lose or gain money no matter which horse wins. That represents the risk-neutral probabilities.

4.2 European Call Options

4.2.1 Introduction

A call option is a financial contract that gives its holder the right, but not the obligation, to buy a specified amount (nominal = N) of an underlying asset (S) at a specified price (strike price = K) within a specific time period (maturity = T).

Therefore, a call buyer profits when the underlying asset increases in price.

It is important to distinguish between payoff at maturity, premium or market price and profit.

4.2.2 Buyer's payoff at maturity

$$C_T = \text{Max}(S_T - K, 0) = (S_T - K)^+ = 1_{\{S_T > K\}}(S_T - K)$$

It is good that you get familiar with these notations as soon as possible as you will see them a lot through this material. You must know them perfectly when working in equity derivatives.

$1_{S_T > K}$ is an indicator function having a value of 1 if $S_T > K$ and 0 otherwise.

Let's plot the Call payoff with respect to the underlying spot price.

As you can see in the graph, the option's strike price is the key point which divides the payoff function in two parts. Below the strike, the payoff chart is constant and null.

Above the strike the line is upward sloping, as the call option's payoff is rising in proportion with the underlying price.

So you have got nothing to lose at maturity when buying a call option. you will therefore have to pay for this right!

How much?

Well, that's given by the call option premium and we will see now how to calculate it.

The good thing when buying a call option is that your loss is always limited to the premium paid for being long the call option. Inversely, when selling a call option, your potential loss is illimited.




Figure 4.1: Fig: 4.1 : Call Payoff at Maturity

The good thing with a long call trade is that your loss is always limited to the initial cost.

4.2.3 Market Price or Premium

The market price of the call option is called the premium. It is the price paid for the rights that the call option provides.

Option premiums are expressed as a per-share basis while most option contracts represent 100 shares of the underlying stock. Thus, a premium quoted as 1.1€ means that it will cost you 110€ as the option is generally on 100 shares.

4.2.3.1 Black-Scholes formula

The Call payoff at maturity is quite easy to understand but it does not tell you how much is the option worth at any time before expiry.

European Call Price: $C = Se^{-qt}N(d_1) - Ke^{-rt}N(d_2)$

Where

$$d_1 = \frac{\ln(\frac{S}{K}) + (r - q + \frac{\sigma^2}{2})t}{\sigma\sqrt{t}} \quad \text{and} \quad d_2 = \frac{\ln(\frac{S}{K}) + (r - q - \frac{\sigma^2}{2})t}{\sigma\sqrt{t}} = d_1 - \sigma\sqrt{t}$$

$$N(x) = \frac{1}{\sqrt{2\pi}} \int_{-\infty}^x e^{-z^2/2} dz$$

While all these equations might look scary at first. They reflect nothing but very intuitive concepts that can be understood by anyone out there.

A first observation is that the expected rate of return of the underlying S , does not enter into those equations. In fact, the relevant parameter is the risk-free rate of interest. That is due to the risk-neutral assumption behind the Black-Scholes model.

Formula for both Calls and Puts can be split into two components that have a very intuitive meaning that will surely help you understand it better. Note that we assume the dividend yield, q , to be zero in the following development.

We have seen that Call option payoff at maturity is: $C_T = \text{Max}(S_T - K, 0) = (S_T - K)^+ = 1_{\{S_T > K\}}(S_T - K)$

So, let us split this formula into two components so that: $C_T = C_T^1 + C_T^2$

and therefore $C_0 = C_0^1 + C_0^2$ by the no-arbitrage argument.

The first component is the payment of the strike price, conditional on the option finishing in the money (= conditional strike payment).

The payoff of this claim is: $C_T^1 = -K * 1_{\{S_T > K\}}$

The second component is the receipt of the stock, conditional on the option finishing in the money (= conditional receipt of the stock).

Its payoff is: $C_T^2 = S_T * 1_{\{S_T > K\}}$

The Call option can therefore be valued by valuing each of these two components separately.

The present value of the conditional strike payment will be the expected future payment, computed on the basis of the risk-adjusted probability distribution, discounted at the risk-free rate.

The expected future payoff is: $E(C_T^1) = -K * P\{S_T > K\}$ where P is the risk-adjusted probability.

The present value is: $C_0^1 = -K * e^{-rT} * P\{S_T > K\}$

It turns out that the risk-adjusted probability of the option finishing in the money is nothing but $N(d_2)$: $P\{S_T > K\} = N(d_2)$.

Therefore, the present value of the conditional strike payment is: $C_0^1 = -K e^{-rT} N(d_2)$

Don't you recognize part of the Call price formula?

The present value of the conditional receipt of the stock will also equal the expected future value of the stock, computed using the adjusted probabilities, discounted at the risk-free rate.

The expected future payoff of this component is the conditional expectation of the stock price given that option is exercised times the probability of the option being exercised: $E(C_T^2) = E[S_T | S_T > K] P\{S_T > K\}$

This expression looks slightly more complicated than for the first component. This is simply due to the fact that the strike price K is a constant (and the expected value of a constant is nothing but the constant itself as it is known with certainty) while the stock price is a random variable.

It turns out that: $E(C_T^2) = E[S_T | S_T > K] P\{S_T > K\} = e^{rT} S N(d_1)$

Therefore, the present value of the conditional receipt of the stock is:

$$C_0^2 = S N(d_1)$$

Don't you recognize part of the Call price formula?

By putting together the values of the two components of the option payoff, we get back to the Black-Scholes formula.

Why is the PV of the conditional receipt of the stock not $S N(d_2)$, corresponding to an expected FV of $e^{rT} S N(d_2)$?

Well, if it was the case the PV of the call option would be $(S - e^{-rT} K) N(d_2)$. This would be negative when the call option is OTM, which obviously cannot be the case.

Where is the flaw in the argument then?

It is that the event of exercising your option is not independent of the random magnitude of S_T . If exercise were completely random and unrelated to the stock price, then indeed it would have been the case. However, exercise is not purely random and depends on the future stock price since it happens only when S_T is high.

Therefore, $S e^{rT} N(d_2)$ underestimates the expected value.

$$E[S_T | S_T > K] P\{S_T > K\} = E[S_T | S_T > K] N(d_2) > e^{rT} S N(d_2)$$

because the correlation between S_T and the decision of exercising the call option implies that: $E[S_T | S_T > K] > e^{rT} S$

4.2.3.2 Numerical Application

In the absence of interest rates and dividends, the price of an ATM call can be approximated by the following formula: $C_T = 0.4 \sigma \sqrt{T} S_0$

4.2.3.3 Premium Graph

Plotting the Call price with respect to the spot price shows an ascending convex curve.

Ascending curve: long call positive delta long forward

Convex curve: long call positive gamma




Figure 4.2: Fig: 4.2 : Initial Call Price for different Spot levels

4.2.3.4 Moneyness

Moneyness is a term describing the relationship between the strike price of an option and the current trading price of its underlying security.

Call options are said to be:

- **In-the-money (ITM)** if Spot price $>$ Strike price.
- **At-the-money (ATM)** if Spot price $=$ Strike price.
- **Out-of-the-money (OTM)** if Spot price $<$ Strike price.

It is also time to bring up the concept of intrinsic value, which is the value any given option would have if it were exercised today. Plotting the intrinsic value yields the same graph as the payoff at maturity.

4.2.3.5 Comparison: Premium vs Payoff

As we can see, the call price (red) is always bigger than the intrinsic value (green) of the option. The difference is the time value (blue), which measures the uncertainty of the option ending in-the-money (ITM). It is almost always positive for a call and reaches its maximum at-the-money (ATM). As time to maturity decreases, the time value decreases to be equal to zero at expiry date. In other words, Call options tend to lose some of their value as time passes (*ceteris paribus*). They usually lose around two-thirds of their time value during the last third of their life.

Do not worry if you are not familiar with some terms, we will see them in more details soon.




Figure 4.3: Fig: 4.3 : Comparison: Premium vs Payoff

4.3 European Put Options

4.3.1 Introduction

A put option is a financial contract that gives its holder the right, but not the obligation, to sell a specified amount (nominal = N) of an underlying asset (= S) at a specified price (strike price = K) within a specific time period (maturity = T).

Therefore, a put buyer profits when the underlying asset decreases in price.

It is important to distinguish between payoff at maturity, premium or market price and profit.

4.3.2 Buyer's payoff at maturity

$$P_T = \max(K - S_T, 0) = (K - S_T)^+ = 1_{\{S_T < K\}}(K - S_T)$$

Let's plot the Put payoff with respect to the underlying spot price.

As you can see in the graph, the option's strike price is the key point which divides the payoff function in two parts. Below the strike, the payoff chart is in the positive territory and the payoff increases as the underlying price goes down. This relationship is linear.

Above the strike, the payoff is constant and null.

So you have got nothing to lose at maturity when buying a put option. you will therefore have to pay for this right!




Figure 4.4: Fig: 4.4 : Put Payoff at Maturity

How much?

Well, that's given by the put option premium and we will see now how to calculate it.

The good thing when buying a put option is that your loss is always limited to the premium paid for being long the put option. Inversely, when selling a put option, your potential loss can be quite substantial if the underlying stock price plunges.

4.3.3 Market Price or Premium

The market price of the put option is called the premium. It is the price paid for the rights that the put option provides.

Option premiums are expressed as a per-share basis while most option contracts represent 100 shares of the underlying stock. Thus, a premium quoted as 1.1€ means that it will cost you 110€ as the option is generally on 100 shares.

4.3.3.1 Black-Scholes formula

The Put payoff at maturity is quite easy to understand but it does not tell you how much is the option worth at any time before expiry.

European Put Price:
$$P = Ke^{-rt}N(-d_2) - Se^{-qt}N(-d_1)$$

Where


$$d_1 = \frac{\ln(\frac{S}{K}) + (r - q + \frac{\sigma^2}{2}t)}{\sigma\sqrt{t}} \quad \text{and} \quad d_2 = \frac{\ln(\frac{S}{K}) + (r - q - \frac{\sigma^2}{2}t)}{\sigma\sqrt{t}} = d_1 - \sigma\sqrt{t}$$

$$N(x) = \frac{1}{\sqrt{2\pi}} \int_{-\infty}^x e^{-z^2/2} dz$$

I am not going to develop this formula further as the idea behind it follows the same line of thought as for the Call option.

4.3.3.2 Premium Graph

Plotting the Put price with respect to the spot price shows a descending convex curve.



Descending curve: long put negative delta short forward

Convex curve: long put positive gamma

4.3.3.3 Moneyness

As for Call options, put options can be categorized as being:

- **In-the-money (ITM)** if Spot price < Strike price.
- **At-the-money (ATM)** if Spot price = Strike price.
- **Out-of-the-money (OTM)** if Spot price > Strike price.

4.3.3.4 Comparison: Premium vs Payoff

The most observant among you will have noticed that the time value is not always positive for a put!

We will further develop the concept of time value in the next section.




Figure 4.5: Fig: 4.6 : Comparison: Premium vs Payoff

4.4 Intrinsic Value and Time Value

4.4.1 Introduction

The concept of time value is very simple and intricate at the same time.

The time value can be defined as the value of the option above the intrinsic value. You should know by now that the intrinsic value is the value any given option would have if it were exercised today.

You can always express an option price as the sum of its intrinsic value and its time value.

$$\text{Option Price} = \text{Intrinsic Value} + \text{Time Value}$$

Call Intrinsic Value = $\text{Max}(0; S - K)$

Put Intrinsic Value = $\text{Max}(0; K - S)$

At first sight, one might assume that time value has to be a positive number. Well, if we take out-of-the money options, their intrinsic values are zero but their prices are typically not zero. Therefore, this differential is positive. Hence, we intuitively expect the time value to be positive.

While it is most often positive, we have seen that it could also be negative in some situations.

We will now discuss the attitude of time on call and put options.

4.4.2 Time Value of Call Options

Let us prove the positive attitude of the time on a call option by considering both OTM and ITM cases.

4.4.2.1 OTM and ATM Cases

OTM and ATM cases are trivial since premium is always positive and hence always more than the intrinsic value being worthless.

4.4.2.2 ITM Case

Suppose we own an in-the-money European call and want to lock the intrinsic value $S_0 - K$. We can do it by selling forward the stock to freeze the price risk.

If the option ends up in-the-money, we will exercise the option and our cashflow will be: $S_0 e^{(r-q)T} - K$

If the option ends up out-of-the-money, we won't exercise it and the only cash-flow will be: $S_0 e^{(r-q)T} - S_T > S_0 e^{(r-q)T} - K$

If $S_0 e^{(r-q)T} - K$ is always bigger than the intrinsic value $S_0 - K$, then the cost of this hedge (= the call price) needs to be above this final guaranteed outcome (by the no-arbitrage argument) $\Rightarrow$ $Call\ price > S_0 - K$. In such cases, the time value will always be positive.

On the contrary, if this inequality is not always verified, time value can happen to be negative.

4.4.2.3 Conclusion

The time value of a Call option is always positive except for deep ITM calls when $(r - q)$ is negative. This is the case when the dividend yield of the underlying stock is higher than the risk-free interest rate ($q > r$). Note that this includes the case of a non-dividend paying stock in a negative interest-rate world.

The always positive time value of Call options on non-dividend paying stocks in a positive rates environment explains why there is no difference in premium between American and European Calls on such underlyings.

4.4.3 Time Value of Put Options

We can wonder why this argument does not hold for put options. It looks reasonable to do the same thing.

4.4.3.1 OTM and ATM Cases

OTM and ATM cases are trivial since premium is always positive and hence always more than the intrinsic value being worthless.

4.4.3.2 ITM Case

For hedging in-the-money put options, we need to buy forward the stock to freeze the price risk and lock the intrinsic value $K - S_0$.

If option ends up ITM, we will exercise the option and our cashflow will be:

$$K - S_0 e^{(r-q)T}$$

If option ends OTM, we won't exercise and our cashflow will be: $S_T - S_0 e^{(r-q)T} > K - S_0 e^{(r-q)T}$

By the no-arbitrage argument, the cost of this hedge needs to be above this final guaranteed outcome. $\Rightarrow$ $Put\ price > K - S_0 e^{(r-q)T}$

However, since $K - S_0 e^{(r-q)T} < K - S_0$ when interest rate is bigger than the dividend yield of the underlying stock, it becomes possible that the premium is less than the intrinsic value and therefore the time value negative.

4.4.3.3 Conclusion

The time value of a Put option is always positive except for deep ITM Puts when the interest rate is bigger than the dividend yield ($r > q$). This happens more frequently as some stocks do not pay dividends. That was the case in Fig: 3.6 where I plot a 1-year put option on a non-dividend paying stock in a 2% interest-rate world. So the situations where the time value is negative can be traced back to the cashflows. Receiving money now or receiving the same amount of money in the future is not the same. And the direction of the cashflows in case of the put option are reversed!

The whole construction here is a nice example of *proxy hedging* the option.

However, the outcome of the hedging strategy is unsure. We had inequalities rather than equalities.

The reason for the residual risk has to do with the fundamental aspect of options.

We started by saying that the option was ITM and hedging accordingly. If then, afterwards, it turns out to be OTM, there arises a risk/uncertainty in the value the hedge will produce. The balance between the option and the hedge is broken at that time.

4.5 The Cost of Hedging

I will never emphasize it enough but it is very important to remember that **the price of an option should reflect the cost of hedging it!**

There are many parameters to consider when pricing an option: interest rates, volatility, repo, dividends, exchange rate, correlation, etc.

For european options, these parameters can be inserted into the Black-Scholes formula to get their prices. Such closed-form solutions that reflect the cost of hedging do not always exist for more sophisticated payoffs. Moving towards more complicated options, it is essential that you keep in mind that the cost of an option should reflect the cost of hedging the risks entailed!

Often all the risks are not hedgeable so that the product won't be fully hedged to replication. It will then be hedged so that the remaining risk exposures are tolerable according to the risk department.

Pricing an exotic product therefore comes down to understand, manage and price the underlying risks.

4.6 The Call-Put Parity

4.6.1 Introduction

Put-call parity specifies a relationship between the prices of european call and put options with the identical strike price K and expiry T . Perhaps the most important feature of put-call parity is that it must be satisfied at all times, in a model independent manner. A violation of this leads to arbitrage opportunities.

4.6.2 The Call-Put relationship

$$Call(K, T) - Put(K, T) = S_0 e^{-qT} - K e^{-rT}$$

The left side of this equality consists of a portfolio A, which is a long call position and a short position with the same strike price and maturity (K, T) on the same underlying asset S .

The right side of this equality consists of a portfolio B, which is a long position in a forward contract that gives its holder the obligation to buy the underlying asset A at a price K at maturity T . In all states, both portfolio A and B have the same payoff at maturity.

As early exercise is not possible for European options, if the values of these two portfolios are the same at the expiry of the options, then the present values of these portfolios must also be the same. If it is not the case, an investor can arbitrage and make a risk-free profit by purchasing the less expensive portfolio, selling the more expensive one and holding the long-short position to maturity.

Accordingly, we have the following price equalities:

$$Call(K, T) - Put(K, T) = Forward(K, T)$$

and

$$Forward(K,T) = S_0e^{-qT} - Ke^{-rT}$$




Figure 4.6: Fig: 4.7 : Call-Put Parity: Call - Put = Synthetic Forward

Chapter 5

The Greeks

If there is one chapter you should master, it is without doubt this one.

Some examples in this chapter are coming from the excellent book of Peter Leoni: *'The Greeks and Hedging Explained'*. A must-read if you are interested by options.

5.1 Introduction to the Greeks

5.1.1 Introduction

Whenever a bank trades a derivative product, it ends up with a position that has various sources of risk. In practice, the bank does not risk manage each product independently. Instead, it adds each trade to its existing book of options and will risk manage this book globally. Indeed, some individual risks from different exotic products may offset each other.

Where do the various risks lie?

To answer this question, traders will need to know the sensitivity of their book to the market parameters.

These sensitivities are commonly referred to as **the Greeks**.

5.1.2 Taylor Series

To understand the concept of sensitivity we must first mention the Taylor series.

Any function can be approximated by a polynomial function. The coefficients for this polynomial are determined by the derivatives at a single point (understand the current market conditions in this case).

A first order approximation means that we are approximating the option price with a straight line. The second order approximation means we are using a parabolic shape.

Since option prices are convex with respect to some of their parameters, their linear approximations always lie below the exact option prices. **The approximation generally gets better as we add more terms to it.**

This simply gives us the various orders of sensitivities.

$$\Delta V = \frac{\partial V}{\partial S} * \Delta S + \frac{\partial V}{\partial \sigma} * \Delta \sigma + \frac{\partial V}{\partial t} * \Delta t + \frac{\partial V}{\partial r} * \Delta r + \frac{1}{2} \frac{\partial^2 V}{\partial S^2} * (\Delta S)^2 + Other$$

The parabolic approximation is almost undistinguishable unless we move substantially from the initial level.

When pricing an exotic option, you will typically have to use some model. It is therefore fundamental that you understand well the model's assumptions and their implications. Wrong models will typically result in wrong hedge ratios.

5.1.3 Static and Dynamic Hedge

A hedge is an investment made with the intention to offset or at least reduce the risk of a financial asset.

5.1.3.1 Static Hedge

A static hedge is a hedge that requires no changes to its components once it is initiated. In other words, the hedge won't need to be re-balanced irrespective of how the market moves.

If a financial product has all its cashflows aligned with those of liquidly traded instruments, then a static hedge can be put into place and this static hedge is model independent. The price of the financial product is then nothing but the cost of setting up the static hedge.

The existence of a static hedge therefore provides us with both a price and a hedge!

Unfortunately, for most of the exotic products covered in this material, such a static hedge is not possible and a dynamic hedge must be implemented.

5.1.3.2 Dynamic Hedge

Unlike a static hedge, once **a dynamic hedge** is initiated, it is sensitive to movements in the market and must be readjusted to still be a hedge.

The frequency at which the hedge must be rebalanced depends on the nature of its sensitivities and its impact on the price.

The initial hedge is generally made up two parts:

- **A static part** that does not require any further adjustment
- **A dynamic part** that will need to be adjusted through the product life

5.1.4 Summary Table of the Greeks

<i>Greek</i>	<i>Symbol</i>	<i>Measures</i>	<i>Definition</i>
Delta	$\Delta = \frac{\partial V}{\partial S}$	Equity Exposure	Change in option price due to spot
Gamma	$\Gamma = \frac{\partial^2 V}{\partial S^2}$	Payout Convexity	Change in delta due to spot
Theta	$\Theta = \frac{\partial V}{\partial t}$	Time Decay	Change in option price due to time passing
Vega	$v = \frac{\partial V}{\partial \sigma}$	Volatility Exposure	Change in option price due to volatility
Rho	$\rho = \frac{\partial V}{\partial r}$	Interest Rate Exposure	Change in option price due to interest rates
Volga	$\frac{\partial^2 V}{\partial \sigma^2}$	Vol of Vol Exposure	Change in vega due to volatility
Vanna	$\frac{\partial^2 V}{\partial S \partial \sigma}$	Skew	Change in vega due to spot OR change in delta due to volatility
Charm	$\frac{\partial^2 V}{\partial S \partial t}$		Change in delta due to time passing

We will see why these greeks are relevant and how they do their job. We will see their features and behaviours when new market conditions develop.

Note that all the conclusions can easily be drawn by understanding the time value of an option. This by itself will turn out to be enough to qualitatively understand the behaviour of the Greeks! The quantitative understanding requires the expressions provided by the model.

5.2 Delta

5.2.1 Description

Delta is the first-order sensitivity of the price to a movement in the spot price, S. It gives the equity sensitivity of the option and is closely related to the probability of the option finishing in-the-money (but different...).

Delta is normally quoted in percent and gives how much of percentage of the actual stock is required to hedge the option.

5.2.2 Calls

For calls, delta lies between 0% and 100%:

- **Far OTM calls** have delta near 0% meaning there is little to no equity sensitivity.
- **Far ITM calls** have delta close to 100% meaning that they trade like a stock.
- **ATM calls** have delta around 50%.




Figure 5.1: Fig: 5.2 : Call Delta

5.2.3 Puts

For puts, delta lies between 0% and -100%:

- **Far OTM puts** have delta near 0% meaning there is no equity sensitivity.
- **Far ITM puts** have delta close to -100% meaning that they trade like a short stock position.
- **ATM puts** have delta around -50%.

As you can see, **the delta for calls and puts have the same shape**. This is a direct consequence of the call-put parity!

Both deltas are increasing with the stock value.

If there is a small movement in the spot price S , then the price of the derivative will move by delta times this small movement. For example, a delta of 0.5635




Figure 5.2: Fig: 5.3 : Put Delta

means that if the underlying asset moves by 1, then the value of the derivative will move by 0.5635×1 .

Note that **the delta of a book of options is the sum of the individual deltas.**

5.2.4 Assets to delta hedge

To delta hedge, one can use the underlying asset itself, forwards or futures on the underlying asset or another correlated asset.

5.2.4.1 Delta Hedging using Forwards/Futures

Assuming no dividends, the forward price at time t with maturity T is given by $F_t = S_t e^{r(T-t)}$ so that the delta of the futures contract is given by $e^{r(T-t)}$. To match a particular delta of Δ_S with a position in a futures contract on the underlying, you will require a position of $e^{-r(T-t)} \Delta_S$ in the futures contract.

5.2.4.2 Cross Hedging: Delta Hedging using another correlated asset

One can further exploit correlations between assets to delta hedge.

Specifically, if an option is written on an asset with price S_1 , then it is possible to use a second asset S_2 to delta hedge.

How much of this second asset is required for delta hedging is given by the chain rule:

$$\Delta = \frac{\partial V}{\partial S_1} = \frac{\partial V}{\partial S_2} \frac{\partial S_2}{\partial S_1} = \Delta_2 \frac{\partial S_2}{\partial S_1} = \Delta_2 \rho_{1,2} \frac{\sigma_2 S_2}{\sigma_1 S_1}$$

5.2.5 Delta under Black-Scholes

Under Black-Scholes assumptions, deltas for call and put options on non-dividend paying stocks are given by:

$$\text{Call Delta} = \frac{\partial C}{\partial S} = N(d_1) \quad \text{and} \quad \text{Put Delta} = \frac{\partial P}{\partial S} = N(d_1) - 1$$

where:

- $N()$ is the standard Normal Cumulative Distribution Function
- $d_1 = \frac{\ln(\frac{S}{K}) + (r + \frac{\sigma^2}{2})t}{\sigma\sqrt{t}}$

The delta of a European option is therefore sensitive to:

- the time to expiry (t)
- the volatility of the underlying (σ)
- the moneyness ($\frac{S}{K}$).

On a dividend-paying stock, the BS deltas are simply multiplied by e^{-qt} , where q is the dividend yield.

5.2.6 Delta sensitivities

5.2.6.1 Impact of Time to Maturity

At maturity, delta has a digital shape around the strike. **Once we move ourselves away from maturity, the delta becomes much smoother shaped.** The further we are from maturity, the flatter the curve looks.

The delta converges faster and faster to its intrinsic value as time to maturity goes to zero. The option picks its direction. That is true, except for ATM option, where the delta remains flatter.

Main variation from the delta is always located in the ATM range. The closer we are to maturity, the tighter the range is where the delta changes from small value to full value.

Since delta hedging is uncertain and implies transaction costs, there is a willingness to keep it to a minimum. **In practice, delta hedges are usually done on a daily basis.**




Figure 5.3: Fig: 5.4 : Impact of Time to Maturity on Delta

Since time impacts delta, one should adjust its effect on delta for weekends and holidays even though the underlying does not move during these days! The closer to maturity, the more important this adjustment.

The effect of time on delta is represented by a second-order sensitivity called **the charm**.

5.2.6.2 Impact of Volatility

The trader wants to know and understand how the delta will change when volatility changes. It will allow him to anticipate the hedge adjustments and take a position accordingly.

A higher volatility increases the delta for OTM options. The more volatility, the less OTM an OTM option really is.

A higher volatility decreases the delta for ITM options. The more volatility, the less ITM an ITM option really is.

More volatile stocks therefore have a less pronounced delta.

5.2.7 Other factors linked to delta hedge

5.2.7.1 Liquidity

Liquidity is also a concern. A trader will adjust his price if the underlying stock is illiquid and difficult to trade. For example, **if a stock is hard to short, then borrowing costs will have to be factored into the price.**




Figure 5.4: Fig: 5.5 : Impact of Volatility on Delta

5.2.7.2 Dividends

As sellers of structured products, banks are structurally long the underlying assets as their delta hedge consists of buying delta on these assets. Being long underlying means they are also long dividends. To obtain a correct option price and hedge, dividends will have to be taken into account.

However, dividends are uncertain. **So, how should dividends be factored into the price?**

It can be done in two ways:

- either in the form of a **term structure of dividend yields**
- either **priced at current levels and hedged using dividend swaps**.

On a book level, exposures to dividends can be significant and will need to be hedged.

5.2.7.3 Interest Rates

To buy delta, a trader has to borrow money. If interest rates go up, borrowing becomes more expensive making the hedging process and therefore the option price more expensive. The sensitivity of an option price to interest rate movements will be further discussed in the section on Theta.

5.2.8 Delta as a hedge ratio, not a probability

We said that delta was closely related to the probability of finishing ITM and used it sometimes this way as it is quite intuitive for option rookies. I must emphasize again that they are not the same thing. This can be intuitively shown with a small extreme binary example.

A stock is trading for 10€ and tomorrow there is a ruling:

- 10% of the time the stock goes to 100
- 90% of the time the stock goes to 0

As you can see, it is correctly priced as 10€ represents the expected value.

What would be the price of an OTM call with a strike @ 50€?

As usual, the price will be the discounted expected value of the cashflows. Discounting does not matter here as it is on a very small time step (1d) so the expected value is $0.1 \cdot 50€ = 5€$. The call is worth 5€.

What would be the delta of this OTM call?

$$\text{Delta} = \frac{\Delta C}{\Delta S}$$

In the down case, the call goes from 5€ to zero as the stock goes from 10€ to zero. $\text{Delta} = \frac{5€}{10€} = 0.5$

In the up case, the call goes from 5€ to 50€ while the stock goes from 10€ to 100€. $\text{Delta} = \frac{45€}{90€} = 0.5$

The call has a 0.5 delta.

In this binary example, the 400% OTM call has a 50% delta despite only having a 10% chance of finishing ITM.

We used this binary example to show that delta is different from the probability of finishing ITM. If you take a more balanced distribution, the two will be closer to each other though.

You could summarize things as follows: **the more positively skewed the distribution, the more the delta diverges from the probability of finishing ITM** (the further OTM the 50% delta call will be).

In this regard, you can also think of the impact of volatility on delta based on its impact on the distribution. The higher the volatility, the more positively skewed the distribution, the further OTM the 50% delta call will be.

If you are skeptical about the impact of volatility on the distribution, more volatility increases the upside of possible payoffs but decreases the probability of the stock going up and increases the probability of the stock going down.

In case you were wondering, this intuition is perfectly in line with the Black-Scholes framework previously described.

In BS, the call delta is $N(d1)$ while its probability of finishing ITM is $N(d2)$. Remember the relationship between $d2$ and $d1$ is: $d2 = d1 - \sigma\sqrt{t}$

So the higher the volatility, the more delta and probability of finishing ITM will diverge.

I hope this discussion will allow you to **avoid the common misconception that ATMF (at-the-money-forward) call and put have the same delta**. They have the same price by the call-put parity!

5.2.9 Setting up a small delta-hedging experiment

Let us try to understand how to manage a book of derivatives. Suppose we know the stock is following the Black-Scholes dynamics. We could then analyse how the delta hedge procedure works. If this theory is solid, we should be able to eliminate all risks while hedging.

5.2.9.1 Market conditions

- Interest rates are at a 2% level $\rightarrow r = 0.02$.
- Underlying stock pays no dividend and follows a Geometric Brownian Motion (GBM) with $\mu = 10\%$ and $\sigma = 20\%$.
- The initial spot price is 100€ $\rightarrow S_0 = 100$
- The call option is ATM with a maturity of 0.1 year $\rightarrow T = 0.1$

We are going to delta hedge at regular time step with $dt = 0.005$ so that there are 20 steps. We will generate sample paths of the stocks by generating standard normal deviates and use the known formula. From the time series, we can calculate back the realised volatility after the experiment $\rightarrow$ realized volatility = 17%.

5.2.9.2 First Day

The trader sells the option at a price of 2.62€, corresponding to an implied volatility of 20%. He will also use this implied volatility for hedging purpose, which prescribed a delta of 52.52% so that he buys 0.5252 stocks. He does this by borrowing $(0.5252 * 100\text{€} - 2.62\text{€}) = 49.90\text{€}$.

After the first day, the trader has the following instruments in his book:

- **A sold call option** at 2.62€.
- **A loan** with a present value of 49.90€.
- **A stock portfolio** with a value of 52.52€.

5.2.9.3 Second Day

The outstanding loan is increased because of the interest effect. In this case, the effect is about 1 cent. The stock has decreased by 10 cents. Because of the delta-position, his stock portfolio loses about 5 cents. The sold option loses more than just 5 cents.

All of this brings the trader's book at a positive value.

After the change in the market, the delta has changed to 41.83%. The trader will thus sell part of his stock portfolio at the current level. With this cash, he pays back part of the loan. The trader is in a new delta-neutral position to start the next day.

5.2.9.4 Observations

The theory of Black-Scholes is prescribing that the hedging cost leads to the option price. However, we notice a profit coming into the book after just one day of hedging! Why?

The reason is that the theory assumes that you are hedging continuously through time, rather than once a day. So the error is induced by the discrete nature of the hedging strategy.

Another way of looking at this is by analyzing the random number that was used to generate the value of the second day. This number is a draw from the standard normal distribution. This one is a small number.

One could say that the volatility of the first day, because of this small draw, is smaller than actually predicted by the distribution. It is smaller than the "typical" number in a standard normal distribution.

Which random number one has to use in order to see no profit and no loss in this first day?

$Z = -1$ or $+1$ $\rightarrow$ one standard deviation out of the mean.

5.2.9.5 Conclusion

Every day where the move of the stock is less than anticipated, based on volatility, the trader will have a gain in his book.

Every day where the move of the stock is bigger than anticipated, based on volatility, the trader will have a loss in his book.

Because of the statistical nature and the assumed independence of the moves on a daily basis, his losses and gains will average each other out, at least if the volatility that realizes corresponds to the pricing volatility at the start.

What happens if right after the price agreement has been done, the market volatility drops to 5%?

This means that the trader will see a nice profit in his book. However, this profit will die out if trader hedges the option during the lifetime, because the final P&L will depend on how the realized volatility turns out. The smart trader would try to buy back this option in the market if possible and hence lock in this margin.

As we explained in this delta-hedging procedure, there are discrete effects. Up to now, we only have a qualitative understanding of those. No worries, the introduction of gamma and theta will help us quantify them.


5.3 Gamma

5.3.1 Description

Gamma measures the change in delta due to the change in underlying price. It represents the second-order sensitivity of the option to a movement in the underlying asset's price. **The higher the gamma, the more convex the theoretical payout.**

Be careful that **gamma is not a measure of value!** A low/high gamma option does not mean a cheap/expensive option. **An option's implied volatility is a measure of its value, not its gamma.**

The price of an option as a function of the underlying price is non-linear. Gamma allows for a second-order correction to delta to account for this convexity.



This convexity in the underlying price is what gives the option value. It is important as it enables the traders to derive the profit on an option for any given stock move.

A delta-hedged portfolio is a portfolio that has been hedged by trading in the underlying asset against small movements in these assets (black line).

As a second-order effect, the gamma becomes significant when large move in the underlying occurs (convexity seen in the blue curve).

As we can see in the above figure, **a long position in option is convex and there has a positive gamma**. To delta hedge, the trader will need to sell stocks if the stock price goes up and buy stocks if the stock price goes down (sell high – buy low). **The trader can sit on the bid and offer.**

A short position in option is negative gamma. In this case, the trader will need to sell stocks if the stock price goes down and buy stocks if the stock price goes up to be delta hedged (sell low – buy high). **The trader has to cross the bid-offer spread.**

5.3.1.1 Gamma Hedging

To hedge this gamma, we need to trade other convex instruments such as other European options so that gammas cancel out. A lower gamma means that we will have a lower need for large and frequent rebalancing of delta.

To see the second-order effect in pricing, we will use models that assume some form of randomness in asset's price.

5.3.2 Gamma under Black-Scholes

Gamma for both calls and puts are the same.

This is not surprising at all. We have seen that gamma was the first derivative of delta with respect to the underlying price and that the graphics of delta for calls and puts were similar.

This can easily be seen mathematically. Since the delta of a call for a non-dividend-paying stock is given by $\Delta_C = N(d_1)$ and the delta of a put on a non-dividend-paying stock is given by $\Delta_P = N(d_1) - 1$, differentiating with respect to the spot price has the same effect!

$$\text{Put Gamma} = \text{Call Gamma} = \Gamma = \frac{\partial^2 C}{\partial S^2} = \frac{\partial \left(\frac{\partial C}{\partial S} \right)}{\partial S} = \frac{\partial N(d_1)}{\partial S} = \frac{\partial N(d_1)}{\partial d_1} \frac{\partial d_1}{\partial S} = N'(d_1) \frac{\frac{1}{S}}{\sigma \sqrt{t}} = \frac{N'(d_1)}{S \sigma \sqrt{t}}$$

Since deltas of both calls and puts are increasing with the stock value, gamma is always positive for a long option position.

When learning about delta, we have seen that **most of the change in the delta occurs around the ATM point. So the gamma should be large around this point!** As the spot price moves away from the strike, $N'(d_1)$ will become small and gamma will rapidly decrease. This can be observed in the above figure where the gamma of a call option is plotted. Notice how it resembles the bell-shaped curve of the Normal distribution.




Figure 5.5: Fig: 5.7 : Option Gamma

As we move into exotic structures, we will see that these may have quite different Gamma profiles to the European options. Their gamma can change sign!

Note that **the gamma of a portfolio of options is also the sum individual gammas of the options.**

5.3.3 Gamma sensitivities

5.3.3.1 Moneyness

Gamma tells us how much delta will move if the underlying moves. **Gamma generally has its peak value close to ATM and decreases as the option goes deeper ITM or OTM.** Options that are deeply ITM/OTM have gamma close to zero. This behavior is obvious when looking at figure 5.7.

5.3.3.2 Time to Maturity

We have seen that **delta get smoother for larger time to maturity. So we can expect a smoother gamma for a longer maturity.**

As the time to expiration draws nearer, gamma of ATM options increases while gamma of ITM/OTM options decreases.

5.3.3.3 Volatility

We have seen that **volatility attacked the delta, so it is not surprising it weakened the gamma too.** Since a higher volatility induces a less pronounced




Figure 5.6: Fig: 5.8 : Impact of Time to Maturity on Gamma

S-shape delta curve, it also induces a less pronounced/wider more stable bell-curved gamma.

A higher volatility lowers the gamma in the ATM region and increases the gamma in the ITM/OTM regions.

It is very easy to think about this effect in terms of time value of options.

For low levels of volatility, deep ITM/OTM options have little time value and can only gain time value if the underlying asset moves closer to strike.

For high levels of volatility, both ITM/OTM options have time value and the gamma near strike should not be too different from away from strike. Therefore, **for high levels of volatility, gammas tend to be more stable across all strike prices.**

We can say that volatility weakens the gamma.

5.3.4 Negative Gamma at maturity is tricky

The gamma of a European option is high when the underlying trades near the strike and there is little time left to maturity. Around this point, the trader will need to delta hedge more frequently, making the hedging process more expensive.

Intuitively, the delta of a call option on the day of expiry will change from roughly 0% when the spot is lower than the strike to roughly 100% when the spot is larger than the strike.




Figure 5.7: Fig: 5.9 : Impact of Volatility on Gamma

Therefore a small change in spot will result in a large change in delta and gamma is very high.

Any trader that has sold options can tell you how tricky it is to manage the book near maturity when the spot is around the option's strike. The gamma gets so big that the hedging needs to be done very fast. As he is the option seller, he is gamma short, with a huge gamma position. If the market keeps moving around the strike, every rebalance will cost him money and the amount of theta for is limited.

In some cases, in highly volatile markets, it is sometimes better to either fully hedge at once or fully unwind. If you are wrong, it will cost you, but rebalancing can be even more costly (slippage or bid/offer spreads).

5.3.5 Gamma P&L

This section is taken from the article on variance swaps written by the Equity Derivatives Research from JP Morgan. I will heavily rely on this excellent article when writing about variance swaps.

Options are exposed to a wide range of factors such as: performance of the underlying asset, time, volatility, interest rates, dividends, etc.

The first-order exposure to moves in the underlying can be hedged out by the familiar delta-hedging process. This leaves the exposure to volatility, paid for in time-decay as the most important sensitivity. However, this is not a pure volatility exposure as it is **path dependent**!

Suppose we hold a call option. To be delta-neutral, we can sell an amount of the underlying equivalent to the delta of the option. By frequently re-adjusting this delta-hedge, sensitivity to asset moves can be dynamically hedged out over option's lifetime.

P&L will come from the accumulated action of continuously re-balancing to keep the portfolio delta-neutral over time. **The gamma P&L is paid for in the option premium, which is marked to market as lost theta.**

How is this gamma P&L actually made ?

Essentially the gamma measures the convexity of the option. This convexity always works in favour of long options positions.

For small moves in the underlying, the replicating hedge is accurate. For larger moves in the underlying, the long options position is always outperforming the replicating hedge in both directions. This is the principle of the convexity by itself. The first derivative of a convex curve is always below the curve itself.

For a delta-hedged option, the gamma P&L will be the outperformance of the option over the replicating hedge. The more the delta changes, the more the replicating delta-hedge will underperform the long options position.

How much will it underperform?

The actual amount depends on the difference between the initial and final delta of the option, which is gamma time the change in the spot price.

Let us assume that gamma remains constant and see how the gamma P&L can be calculated.




Figure 5.8: Fig: 5.10 : Gamma P&L from delta-hedged options

0. Initially, the underlying spot price is S_0 , the option is worth p_0 , delta and gamma are respectively Δ_0 and Γ_0 .
1. The spot then increases by dS to S_1 .

2. The delta increases to $\Delta_1 = \Delta_0 + \Gamma_0 dS$
3. The average delta is therefore: $\Delta_{avg} = \Delta_0 + \frac{\Gamma_0 dS}{2}$
4. The new option premium is worth $p_1 = p_0 + \Delta_{avg} dS$
5. Replacing Δ_{avg} with 3., we have: $p_1 = p_0 + \Delta_0 dS + \frac{\Gamma_0 dS^2}{2}$
6. The P&L on the long option position is: $P\&L = p_1 - p_0 = \Delta_0 dS + \frac{\Gamma_0 dS^2}{2}$
7. The P&L on the short delta-hedged position is: $P\&L = -\Delta_0 dS$
8. The total P&L is therefore: $\text{Total P\&L} = 6. + 7. = \frac{\Gamma_0 dS^2}{2}$

So the gamma P&L, which is the difference between the performance of the option and the replicating delta hedge is $\frac{1}{2} * \Gamma * dS^2$.

The gamma P&L from a move in the underlying is proportional to:

- the gamma of the option, (Γ)
- the square of the move, (dS^2) $\rightarrow$ gamma P&L from a 2% move will be four times that of a 1% move!

5.3.6 An example

A trader is long gamma through buying a put option. If the market is right and balanced, the trader paid a fair price so that the money he can expect to make will be paid for through premium.

Assume the market suddenly realizes that the option was underpriced. Expectation now becomes that the stock will be more volatile so that the implied volatility increase and the option becomes more expensive in the market.

The trader has a choice. He can keep on delta hedging and profit from this highly volatile market or he can cash in and sell the option at a higher price.

How much money can he expect to make over the lifetime of the option if he continues hedging himself?

That amount can be shown to be identical to cashing in right now. Underlying this statement, there is the assumption that the implied volatility turns out to be the correct one.

This actually gives the trader an interesting dilemma. Because the market will keep changing his mind over time about the best implied volatility that has to be used. Option traders will take positions when they think the implied volatility is low/high. The money they will pick up by trading is directly related to the vega of these options.

5.3.7 Can volatility be captured by delta-hedging?

We have seen how the gamma P&L was made by benefiting of the volatility when delta-hedging. This exposure to volatility is not a

pure volatility exposure as it is path dependent!

Suppose a market maker buys and delta-hedges a vanilla option. Assuming that whatever volatility is realized is constant and option is delta-hedged over infinitesimally small time step, then the market maker will profit if and only if the realized volatility is larger than the implied volatility at which the option was purchased.

However, volatility is not constant! **Where and when the volatility is realized is crucial!** Furthermore the magnitude of the P&L depends not only on the difference between realized volatility and implied volatility, but where that volatility is realized in relation to the option strike.

If the volatility realizes close to maturity when the underlying trades near the strike then the gamma would be very high and the absolute value of the P&L will be larger.

For non-constant volatility, it is possible to buy and delta-hedge an option at an implied volatility smaller than the subsequent realized volatility and still losing money from delta-hedging!

5.3.7.1 Example with a 1-year SX5E option

The SX5E index is initially quoting at 3500 with an ATM implied volatility of 28.5%. During the first seven months, the index ranged between 3500 and 3800 points, realizing a volatility of 20%. After that period, the index fell rapidly to 2500 points, realizing approximately 50% volatility on the way. Over the whole year, the realized volatility was 36%.

Case 1: Long 2500-strike option @ 33.4% Suppose a trader is long a 2500-strike option at an implied volatility of 33.4% due to the skew observed in the market.

Initially, the option made a loss due to low realized volatility. However, the loss was kept small thanks to the low gamma exposure as the spot was trading far from the option strike.

As the SX5E sold off, the option gamma increased while volatility picked up and the option ended up making a large profit!

Case 2: Long 4000-strike option @ 26% Suppose a trader is long a 4000-strike option at an implied volatility of 26% due to the skew observed in the market.

Initially, the option has a higher gamma and lost more due to lower realized volatility.

As the index sold off and volatility occurred, delta-hedging failed to capitalise on it because the gamma was very low far from the strike.

Despite the fact that the option was purchased at an implied volatility 10% below the subsequent realized volatility, it ended up losing money!

Conclusion

When delta-hedging daily an option, the P&L is a daily accrual that depends on the difference between the realized volatility and the implied volatility. The magnitude of the contribution of this daily accrual is weighted by the current dollar gamma, which is unpredictable since path-dependent!

5.3.8 Dollar Gamma

Remember the formula of gamma given previously: $\Gamma = \frac{N'(d_1)}{S\sigma\sqrt{t}}$

Since $N'(d_1)$ is independent of the monetary units of the underlying, gamma will be inversely proportional to the value of the underlying (S at the denominator of the formula).

Gamma is a very useful concept but it has two drawbacks:

- It measures the **change in delta per unit of underlying**.
- It is **dependent on the absolute level of the underlying**.

Moving to the concept of the dollar-gamma is more useful because:

- It allows us to directly calculate the gamma P&L for a given percentage underlying move.
- It makes it **much easier to compare gamma exposures across different underlyings**.

Dollar delta is the cash equivalent exposure of the underlying.

Dollar gamma is the change in the dollar delta for a 1% move in the underlying.

It can easily be shown that $\$ \Gamma = \frac{\Gamma S^2}{100}$

If S changes from S to $S + \frac{S}{100}$, then Δ changes from Δ to $\Delta + \frac{\Gamma S}{100}$ and $\$ \Delta$ changes from $\$ \Delta$ to $\$ \Delta + \frac{\Gamma S^2}{100}$.

For a return R, it can also be shown that $\text{gamma P\&L} = 50 \$ \Gamma R^2$.

Remember that if S changes to $S + dS$, then the gamma P&L would be $\frac{\Gamma dS^2}{2}$.

We have just seen that $\$ \Gamma = \frac{\Gamma S^2}{100}$. By isolating Γ , it can be rewritten as:
 $\Gamma = \frac{100 \$ \Gamma}{S^2}$.

Replacing this expression into the previous one, the gamma P&L can be written

as:
$$\frac{100 \$ \Gamma}{S^2} * \frac{dS^2}{2} = 50 \$ \Gamma \left(\frac{dS}{S}\right)^2 = 50 \$ \Gamma R^2$$

For example, if we hold a position which is long \$100.000 of dollar gamma and the underlying asset moves by 3%, then the P&L will be $50 * \$100.000 * (0.03)^2 = \$4.500$.

5.4 Theta

5.4.1 Description

There is no free lunch in the world of finance. The story would be too good to be true if holding a call/put and dynamically rebalancing gave a profit due to gamma.

Positive gamma does not come for free.

We will now explain where the paradox is coming from. We left out one very important feature. **The trader who owned the option and was making free gamma money out of it had to pay the premium in the first place.** The profit he makes afterwards has to make up for this premium payment.

An option holder pays for the right of buying low and selling high by means of the theta, the time decay of an option. The holder of an option needs to earn back the daily loss of the option by taking advantage of the underlying's moves.

The seller of an option makes money on the theta and loses it by rebalancing delta by buying high and selling low.

The longer the lifetime of the option, the more time to move further ITM, the more expensive the option should be. So we know that the value of the option becomes worth less and less with every day that passes by. This process is inevitable. So little by little the option loses its value.

The change in value from one day to the next is called the theta. Clearly this number is expected to be negative for both ATM call and ATM put option.

5.4.2 Theta under Black-Scholes

Theta is defined as the rate of change of the option price respected to the passage of time: $\$ \Theta = \frac{\partial V}{\partial t} \$$.

Instead of measuring the passage of time all else being equal, **we also could simply investigate the same option with an expiry of one day earlier**. Both approaches are identical within Black-Scholes as the Brownian Motion is a stationary process. In terms of formulas, we can therefore write:

$$\Theta = \frac{\partial V}{\partial t} = \frac{\partial V}{\partial T} \frac{\partial T}{\partial t} = \frac{\partial V}{\partial T} (-1) = -\frac{\partial V}{\partial T}$$

Under Black-Scholes assumptions, thetas for call and put options on non-dividend paying stocks are given by:

$$\text{Put Theta (per year)} = \Theta = -\frac{S\sigma}{2\sqrt{T}}N'(d_1) + rKe^{-rT}N(-d_2)$$

$$\text{Call Theta (per year)} = \Theta = -\frac{S\sigma}{2\sqrt{T}}N'(d_1) - rKe^{-rT}N(d_2)$$




Figure 5.9: Fig: 5.11 : Call Theta

5.4.3 Some features of Theta

Let us turn to an example that will unfold the finesse of the Greek that makes you lose. Let us assume that the underlying stock does not pay any dividend and has a spot price of 20, interest rates are at 2.5%. A 1-year ATM put is quoting 2.5€ for an implied volatility of 35%.

Suppose there are 250 trading days in a year as we assume the option only loses value on working days.

We could write a fixed loss of 0.01€/day in our book. It would make the loss of time value completely predictable. However, there would be an inconsistency between the market value and the booked value!

By accepting a marked-to-market valuation, we have to accept that:

- We only know the theta today.
- We also know all time value will have disappeared when we reach maturity.
- But we cannot predict how the time value will behave over time!

In fact, the theta exhibits several features:

- theta loss speeds up as maturity approaches
- theta depends on moneyness
- theta is strengthened by volatility

5.4.3.1 Theta Speed of loss

The option loses little value at the start and this process speeds up as maturity approaches.

5.4.3.2 Theta depends on moneyness

The real potential of an option is found in the ATM range.

Let us take an example where we buy an ATM put initially and imagine different scenario.

Note that **at first sight the theta is negative, there are 2 situations where it turns positive.**

1. **Far ITM put options are known to have positive time decay** (see figure 4.12). Remember we have seen that the time value of ITM put options can be negative, so nothing really surprising here!
2. **For ITM call options, the theta can turn positive as well, in the case where the dividend yield is larger than interest rates so that the forward level is below the current spot level.** Figure 4.11 does not reflect this as it plots the Theta of a 1-year call option on a non-dividend-paying stock in a positive interest rate world.




Figure 5.10: Fig: 5.12 : Put Theta

Case 1 : From ATM to ITM Initially the put option is ATM but then the stock moves down and the put price rises.

Although the put value increased, the time value is less than for ATM option. Clearly, if the time value is less altogether, there is less value to lose when time passes.

Case 2 : From ATM to ITM then back to ATM Suppose the trader owns an ATM put option.

During the first week, the stock does not move so that both time value and option value decrease.

During the second week, the stock drops by 10% so that option price increases (due to intrinsic value) and time value decreases.

We know that **if the time value is smaller, the daily decrease in time value is lower** than before the actual move.

If the stock moves back up to the ATM level, the option value goes down but the time value increases. As of then, the daily decrease in time value is more substantial than before this move back up.

How all the changes in time value add up ?

This is the topic of the next section, where we see the theta in relation to the gamma.

Over the lifetime the sum of all changes in time value on a daily basis just add up to the option premium of the first day.

The reason behind this is nothing else but the fact that **an option converges to its intrinsic value**.

The creation of extra time value is a temporary effect.

5.4.3.3 Theta is strengthened by Volatility

The higher the volatility, the higher the option price and the time value, the more time value to lose, the larger the absolute value of theta.

As the vega is higher in ATM region, the effect is larger in this region so that the theta change is most noticeable when ATM. The higher the volatility, the wider this ATM region.




Figure 5.11: Fig: 5.13 : Impact of Volatility on Theta

5.4.4 Gamma and Theta are always flirting

There is no free lunch in owning an option.

You will make free gamma money and you will lose precious theta money. There is a balance to be found.

A trader that bought an option paid the intrinsic value (if any) and the time value. He knows the time value will be lost during the lifetime of the option. The only thing he does not know is how big the future daily losses/gains will be. However, the sum of all these will have to make up for the extrinsic premium paid. He will have to work for his money by actively delta hedging and locking in the gamma profits.

Every single day the trader who owns the option tries to fight the theta by:

- rebalancing his portfolio at appropriate times.
- leaving the position open if he expects the stock to keep on moving in the same direction.

By waiting, he gets to lock even more profit in one single rebalance but he is taking a risk.

5.4.4.1 Derivation of Black-Scholes equation

Within the Black-Scholes setup, we can derive an expression that exactly specifies this relation between these two greeks:

$$\Theta + \frac{1}{2}\Gamma S^2\sigma^2 = r(V - \Delta S)$$

This relation is interesting because it is telling us how all the different Greeks lead to the price. We know that the cost of any derivative is determined by the cost of hedging.

Let's first consider a portfolio of an option that we will delta hedge.

On the cash part, we have the option premium V and the loan with value equals to the value of the stock holdings in the portfolio, ΔS .

The interest accumulated over short time interval Δt on this net amount is:

$$r(V - \Delta S) \Delta t$$

That is **the deterministic cost of the construction**.

On the dynamic hedging procedure, we know we will lose the amount of theta, Θ .

The Γ term requires slightly more attention to explain. The Γ profit would be

given by:
$$\frac{\Gamma (\Delta S)^2}{2}.$$

Since the volatility is relative, we need to multiply it with the current level of the stock to get an absolute number.

In a time interval t , the typical movement would be: $\Delta S = S\sigma\sqrt{\Delta t}$.

So the combination of theta and gamma effect over t is given by:

$$\Theta\Delta t + \frac{1}{2}\Gamma S^2\sigma^2\Delta t$$

Taking all together, we have:
$$\Theta\Delta t + \frac{1}{2}\Gamma S^2\sigma^2\Delta t = r(V - \Delta S) \Delta t$$

Simplifying by Δt , it leads to the first equation above:
$$\Theta + \frac{1}{2}\Gamma S^2\sigma^2 = r(V - \Delta S)$$

This relation reveals the core of the delta-hedging procedure!

Gamma and theta constantly need to be balanced in order to make up for the premium and the cost of hedging.

Understanding this balance is key to understand how to manage a book of options!

If we rewrite the first equation by using the explicit partial derivatives, we have:

$$\frac{\partial V}{\partial t} + \frac{1}{2}S^2\sigma^2\frac{\partial^2 V}{\partial S^2} + rS\frac{\partial V}{\partial S} - rV = 0$$

This equation is a **partial differential equation (PDE)** known as the **Black-Scholes equation**. Its solution is unique when the boundary and initial conditions are set.

- The boundary condition is given by the payout profile at maturity.
- The initial conditions are set by the stock price at time t_0 .

We have often ignored in the previous discussions the fact that the time value can be negative.

Now it is time to tie up our loose ends!

The price of a put option can go below the intrinsic value, meaning the theta can be positive for these put options.

What does that mean for the theta-gamma balance that we just unveiled throughout this chapter?

This means that a trader long the option will gain from both the gamma and the theta, so there is no longer a trade-off, but money is rolling in from everywhere.

Where is the catch?

Well, the premium for these far ITM put options is quite high. In order to buy this, we need to borrow an amount of cash equal to the premium. This loan needs to be paid back as well; this cost is high and the payments are due every day in the Black-Scholes model.

Why this argument would not hold for a call option as well?

When we own a put, we also need to buy the stock to hedge ourselves so that the substantial loan we had becomes even bigger.

For a call, they would cancel each other out to some extent whereas for the put the effect gets reinforced.

5.5 Vega

5.5.1 Description

Vega is the sensitivity of the option price to a movement in the volatility of the underlying. The need to understand the vega only became important after trading options became as liquid as it is today.

Before that, if you were wrong on your volatility estimate, you would only see the extent of this after all hedging has been completed and the option expired.

Now that the option market is liquid enough and everybody agrees on σ , if its value changes over time, you can undo your transaction. This is what trading is all about after all.

For all the other Greeks, we established expressions within the Black-Scholes model. It seems ridiculous to specify the model for the vega as the volatility and the Black-Scholes model are almost one and the same.

Without Black-Scholes model there would be no σ parameter and without parameter, Black-Scholes model is missing an essential ingredient.

Note that **for most practical applications, the naked vega is normalized** so that: $v = v_{inst} * \Delta\sigma$ with $\Delta\sigma = 1\%$.

This is **the effect of a 1% point change in the volatility on the price of an option**. The practical advantage of this approach is obvious. The numbers that you get in your trading book are meaningful as a 1% point change is a typical move that you might see. Having this vega number in your head immediately indicates in terms of money-terms how your portfolio will be influenced by it.

For an option trader, **it is common practice to specify the bid/offer spread on such an option in terms of the vega**. Depending on the liquidity in the underlying, the volume of the transaction, the risk appetite and so on, he will take as a margin a multiple of the vega.

It will become clear that the parameter σ can be different for every quoted option. It gets adjusted in the market, leading to the concept of the implied volatility surface. **In practice, it means that there is not one such parameter for a certain stock, but a whole matrix $\sigma_{(i,j)}$.**

5.5.2 The Vega Matrix

Consider a collection of different options with various strikes and maturities and a stock position. We calculate the Greek position of each individual option and weight those with volumes to obtain the greek position of the book.

While the book delta, gamma and theta will be the weighted sums of individual deltas, gammas and thetas, the vega terms actually refers

to different volatility parameters, one for each different strike and maturity. They are not completely correlated, as the market decided that options can live their own life.

If we have a volatility for each strike and maturity, the volatility view in book is represented by a matrix of vega terms.

5.5.3 Vega in the Black-Scholes model

Vega is positive for both calls and puts. As for the gamma, the put-call parity also implies that the vega of call is equal to the vega of a put with the same features.

For a non-dividend paying stock, vega is given by: $\text{Put Vega} = \text{Call Vega} = v = \frac{\partial C}{\partial \sigma} = S\sqrt{T}N'(d_1)$




Figure 5.12: Fig: 5.14 : Call Vega

As for gamma, vega is greatest for ATM options and decays exponentially on both sides.

In fact, Gamma and Vega play a similar role to the options price.

Gamma is crucial to understanding how much money the hedge is making/losing. It was directly related to the realised volatility.

The Vega is in a way the forward measure for this.

Bell-Shaped Curve

The bell-shaped curve of vega makes sense intuitively.

When ATM, a change in the volatility can send the option either ITM or OTM thus the large effect on the price.

The vega sensitivity becomes really small when the option is either far ITM or far OTM because these options have already picked their direction in a way. A larger volatility won't make that much difference to the price of the option. These kinds of options are quite easy to hedge, at least to the extent that when the market stays in the same regime, the hedge adjustments are minor over the lifetime of the option.

Of course there is a turning point, where the impact becomes stronger again. As soon as volatility is above this level, one could say the vega becomes meaningful, or the option is ATM from a vega point of perspective.

The shape of the curve is very similar to the shape of the normal density function, but within the expression of vega, there is an extra S factor that gives a twitch of bias.

5.5.4 Vega sensitivities

5.5.4.1 Moneyness

Vega generally has its peak value close to ATM and decreases as the option goes deeper ITM or OTM. Options that are deeply ITM/OTM have vega close to zero. This behavior is obvious when looking at figure 4.14.

5.5.4.2 Volatility

A higher volatility results in a wider vega curve, meaning a vega curve that is more flat. This points to the fact that in highly volatile markets the ATM point is less defined. The ATM region has become bigger.

Higher volatility does not necessarily mean higher vega in the ATM point!

The vega does not depend strongly on the level of volatility used, at least not once a certain critical level of volatility is reached.

5.5.4.3 Time to Maturity

We know that options lose their time value as time passes and we approach the expiry. It is not surprising that **the volatility sensitivity or vega dies out.**

Vega/Gamma have the same qualitative behaviour: largest ATM vs ITM/OTM. However, they have **opposite dynamic behaviour:** gamma spikes towards maturity whereas vega dies out!




Figure 5.13: Fig: 5.15 : Impact of Volatility on Vega

If we want to **capture forward volatility**, then we have to **trade options with the largest vega**. If we want to trade the **spot volatility**, then we have to **trade options with largest gamma**.

When the option is away from the strike levels, vega drops very quickly with time. Only ATM options can hold their vega for a longer time.




Figure 5.14: Fig: 5.16 : Impact of Time to Maturity on Vega

For more exotic structures, the vega profile can change sign, and whether we are short or long volatility depends on the underlying's price.

The **ATM option is almost linear in volatility**.

The prices of the **ITM and OTM options are convex in volatility up to a certain level then become linear for large volatilities**.

5.6 Rho

5.6.1 Description

Rho represents the sensitivity of an option's price to a movement in interest rates.

As the delta, it is **positive for calls** and **negative for puts**.

The **prices of vanilla options are almost linear in interest rates**. In other words, it only has a first-order effect. This effect comes from:

- the impact of interest rates on the cost of the delta-hedge
- the discounting of the option price.

The second effect is generally smaller than the first one.

5.6.2 Calls

A trader sells a call option and delta hedges by buying delta shares. To buy those shares, he must borrow money at the bank. He will have to pay interest on his loan.

The higher the interest rates, the more interests to pay, the higher the cost of his hedge, the higher the call price. This is the reason why rho is positive for call options.

The discounting effect slightly offsets this delta hedge impact though.

$$\text{Call Rho} = \rho = \frac{\partial C}{\partial r} = K t e^{-rt} N(d_2)$$

5.6.3 Puts

A trader sells a put option and delta hedges by selling delta shares. By selling shares, he receives money. He can put this money in the bank and receives interest on it.

The higher the interest rates, the more interests he receives, the lower the cost of his hedge, the less expensive the put price. This is the reason why rho is negative for put options.

The discounting effect works in the same direction and reinforces this impact.

If you make more money on your delta hedge, you are going to pay more (receive less) for it.

$$\text{Put Rho} = \rho = \frac{\partial P}{\partial r} = -K t e^{-rt} N(-d_2)$$




Figure 5.15: Fig: 5.17 : Call Rho




Figure 5.16: Fig: 5.18 : Put Rho

5.6.4 Rho sensitivities

5.6.4.1 Moneyness

Rho is larger for ITM options and decreases steadily as the option changes to become OTM.

5.6.4.2 Time to Maturity

Rho increases as time to expiration increases. Long-dated options are far more sensitive to changes in interest rates than short-dated options.




Figure 5.17: Fig: 5.19 : Impact of Time to Maturity on Rho

Though rho is a primary input in the Black-Scholes model, a **change in interest rates generally has a minor overall impact on the pricing of options**. Because of this, rho is usually considered to be the least important of all the option Greeks.

5.6.5 Rho Hedging

We recall from delta-hedging procedure that we were using a cash account, where we received/were charged an interest rate r , continuously compounded, meaning every day we would receive/pay interest.

This is a source of uncertainty.

One can expect that the **interest received for cash is lower than the interest to be paid for a loan**. This typical bid/offer spread in interest rate market is the first of our problems.

What's next?

The **interest rates are not constant**. So we don't know in advance what the applicable rate interest rate is.

If we have a certain amount on the cash account, we could fix the interest rates for the lifetime of the option, but only for this fixed volume.

Because of the dynamic nature of the hedging procedure, the cash on the account will constantly change and a few days later, the conditions in the market might have changed. So, the best we can do is fix as much as we can at the start of the procedure, where we have at least an accurate view on the money in the cash account. For the future, the trader will use an estimate of the future interest rates.

The problem of uncertain interest rates in the context of equity derivatives is sometimes handled by the bank by fixing an internal system.

5.6.5.1 Internal Markets

For the sake of making the point, we will oversimplify.

Suppose you have 2 different desks in a trading floor:

- The desk that trades options on stocks.
- The desk that trades interest rates.

This will allow the interest rate risk to be transferred from one desk to the other. The management could decide that an internal system needs to be set up to accommodate for this transfer.

Therefore management could decide that the equity derivatives desk always receives an interest rate $r - \Delta r$ and pays an interest rate $r + \Delta r$. The spread Δr could be fixed internally to whatever value is reasonable. Often $\Delta r = 0$ as it is usually clear in advance if an activity will be drawing cash or is self-funding.

That just leaves the determination of r . The fixing of this is harder. Taking the overnight interest rate, which would be the closest match to what we are looking for, leaves out great opportunities.

A bank typically has a good view on the value of the cash account as it has built up some history over time. Leaving the cash on this account does not make sense. The equity desk should shift this capital to other desks where it can be invested and managed more wisely.

One could decide that, for internal purposes, the interest rate r is taken to be the 6-month swap rate, irrespective of option lifetime. One could even decide that value of r is fixed once a week and kept constant throughout the week.

That means that the internal transfer is obligated and, in some cases, the interest desk is forced to take on positions with a loss, at other times it could be a profit. If the volumes coming from the equity derivatives desk are much smaller than the typical volumes they trade, this might not be as bad as it sounds.

For the equity derivatives desk, it eliminates the burden of following up on interest rates.

If the activity is organized this way, it is clear that the equity desk could arbitrage the system. If the interest rate is fixed on Monday and valid until Friday and they are looking into a transaction on Friday, and the interest rate market has changed a lot, they could decide to keep their interest rate position open and wait to deposit the cash until next Monday. An internal check-up should be installed to prevent one-sided abuse between the internal desks.

The main conclusion to be drawn from this is that **an equity derivatives trader should be less focused on the problem of interest rate risk, and more on the value of the stock, and, on the volatility.**

In practice, desks don't have a cash account but instead they receive internal funding.

This funding has become more and more a topic of discussion as the regulator has raised the question if banks have enough capital available to run their business. **From the cost of hedging argument, it is clear that the cost of funding should be factored into the hedging cost and hence into the price of any derivative.**

Different participants in the market will have different funding rates, depending on the capitalization and credit rating they have. This leads to different internal interest rates, but also to different option prices, which they can offer to their clients.

5.7 General Practical Example

Let us go through an example that explains the concept of **vega-gamma-theta hedging**.

Suppose we have a portfolio with the following Greeks representation:

- Delta = 300.000
- Gamma = 2.500
- Theta = -500.000
- Vega = 750.000

Let us assume the following market parameters: - Stock price = 1.500 - Interest rate $r = 2.5\%$ - Dividend yield $q = 0\%$ - Volatility = 25%

Is it possible to design a portfolio that would offset all the Greeks simultaneously?

Well, we can obviously eliminate delta without making a difference to any other Greeks by selling stocks.

To hedge out the other Greeks, we need to use instruments that have non-zero Greeks. In other words, we have to use options.

Let us assume we have 3 options available:

	<i>OTM PUT</i>	<i>ATM CALL</i>	<i>OTM CALL</i>
<i>Strike</i>	1200	1500	1800
<i>Time to Maturity</i>	1	1.5	0.5
<i>Premium</i>	27.85	207.98	25.95
<i>Delta</i>	-0.13	0.61	0.19
<i>Gamma</i>	0.0006	0.0008	0.001
<i>Theta</i>	-0.09	-0.21	-0.22
<i>Vega</i>	3.2	7.05	2.90

If we construct a portfolio with the following weights:

	<i>OTM Puts</i>	<i>ATM Calls</i>	<i>OTM Calls</i>	<i>Stocks</i>
<i>Volume</i>	-542,270	-1,157,791	3,673,006	227,891
<i>Position</i>	Short	Short	Long	Long

How did we obtain these weights? Is it always possible to find those kinds of positions such that we get back what we want?

If we define the weights: w_1, w_2, w_3, w_4 respectively.

We are then trying to obtain a portfolio with portfolio Greeks equal to the ones we are after.

Mathematically speaking:

$$w_1\Delta_1 + w_2\Delta_2 + w_3\Delta_3 + w_4 = \Delta_V \quad w_1\Gamma_1 + w_2\Gamma_2 + w_3\Gamma_3 = \Gamma_V \quad w_1\Theta_1 + w_2\Theta_2 + w_3\Theta_3 = \Theta_V \quad w_1v_1 + w_2v_2 + w_3v_3 = v_V$$

From basic linear algebra, **this set of equations has a solution if its determinant is different from zero:**

$$\begin{vmatrix} \Gamma_1 & \Gamma_2 & \Gamma_3 \\ \Theta_1 & \Theta_2 & \Theta_3 \\ v_1 & v_2 & v_3 \end{vmatrix} \neq 0$$

As it turns out, if we try to use 3 options with identical maturities, this determinant gets to be very close to zero. The reason for this is of course that all options are similar derivatives on the same instrument and the relationship between the Greeks makes the difference between using one option or three options very small. In other words, these options are not independent enough to build up an arbitrary portfolio.

A good mix of strikes and maturities solves this problem.

In the above example, we were looking for the hedging portfolio. If we put on this hedge, it is typically still a dynamic hedge that will need to be adjusted as market moves. But because we minimised more Greeks, the hedge is more stable and rebalancing won't have to happen as often. It becomes much more of a semi-static hedge. **If we have more options available, we can minimise more Greeks and make it an even better hedge.**

One can prove that if we are trying to hedge an exotic derivative whose payout is only determined by the terminal value S_T , then there exists a static hedge, provided we can use all the different strikes.

In practice, one can build a portfolio with features that the trader finds desirable: long vega and short gamma for example.

This portfolio exercise is the foundation for vega, gamma and theta trading.

5.8 Second-Order Greeks

5.8.1 Volga

Vega–Gamma, or Volga, is the second-order sensitivity of the option price to a movement in the implied volatility of the underlying asset.

$$\boxed{Volga = \frac{\partial^2 V}{\partial \sigma^2} = \frac{\partial v}{\partial \sigma}}$$

When an option has such a second-order sensitivity we say it is convex in volatility, or has **Vega convexity**.

ITM and OTM European options do exhibit vega convexity but these can be captured in the skew.

Other structures exhibit a lot of vega convexity and will result in losses if we do not use a model that prices this correctly. The reason is that as volatility moves, a vega convex payoff will have a vega that now moves with the volatility and this must be firstly priced correctly and then hedged accordingly.

5.8.2 Vanna

Vanna measures the sensitivity of the option price to a movement in both the underlying asset's price and its volatility.

$$\boxed{Vanna = \frac{\partial^2 V}{\partial S \partial \sigma} = \frac{\partial v}{\partial S} = \frac{\partial \Delta}{\partial \sigma}}$$

We can thus think about vanna as:

- the sensitivity of the option's vega to a movement in the underlying's price.
- the sensitivity of an option's delta to a movement in the volatility of the underlying.

As such, vanna gives important information regarding a delta hedge by telling us by how much this delta hedge will move if volatility changes. It also tells us how much vega will change if the underlying moves. **It can thus be important for a trader who is delta or vega hedging.**

If vanna is large, then the delta hedge is very sensitive to a movement in volatility.

5.8.3 Charm

Charm, or delta decay, is the rate at which the delta of an option changes with respect to time. It refers to the second order derivative of an option's value, once to time and once to delta.

$$\text{Charm} = \frac{\partial \Delta}{\partial t} = -\frac{\partial \Delta}{\partial T} = \frac{\partial \Theta}{\partial S} = -\frac{\partial^2 V}{\partial S \partial T}$$

Charm values range from -1 to +1.

As time passes, the option loses more of its time value, OTM options see their delta approach zero and ITM options see their delta become closer to that of an equivalent position in the underlying. In other words, **each day that passes, OTM options require less delta hedging, and ITM options require more.**

- **Positive Charm:** ITM calls and OTM puts
- **Negative Charm:** ITM puts and OTM calls
- **Zero Charm:** ATM options

I personally haven't dealt much with this second-order greeks but it is particularly relevant for options traders. They must indeed pay close attention to their charm on Friday. At that moment, the market closes for more than two days and the charm's effect is magnified and impacts their options action on Monday. **Paying attention to the charm may prevent the trader from over or under hedging. Special attention is needed around a charm's expiration time, as it may become very dynamic.**

5.9 Multi-Asset Greeks

5.9.1 Cross-Gamma

The cross-Gamma is the sensitivity of a multi-asset option to a movement in two of the underlying assets. The cross Gamma involving S_i and S_j is given by:

$$\Gamma_{i,j} = \frac{\partial^2 V}{\partial S_i \partial S_j} = \frac{\partial \Delta_i}{\partial S_j} = \frac{\partial \Delta_j}{\partial S_i}$$

It is the effect of a movement in S_i on the delta sensitivity of the option to S_j .

In multi-asset options, gammas could be expressed in a matrix form, with elements on the diagonal being gammas and elements off the diagonal being cross gammas.

In multi-asset options, it is possible that the delta with respect to one asset can be affected by a movement in another underlying asset even if the first asset has not moved.

5.9.2 Cega = Correlation Delta

The correlation delta is the first-order sensitivity of the price of a multi-asset option to a move in the correlations between the underlyings.

The correlation delta with respect to assets S_i and S_j is given by:

$$Cega_{(i,j)} = \frac{\partial V}{\partial \rho_{i,j}}$$

In the contest of pricing derivatives with multi-assets, **we measure the degree of the dependence or dispersion of multi-assets through the correlation matrix.**

Correlations vary over time. A multi-asset product's sensitivity to the correlation between a pair of underlying assets can vary as the other parameters change.

While correlation is not easily tradable, there are some methods of trading correlation. **Many correlation risks are not completely hedgeable, if at all, and in many cases traders must resort to maintaining dynamic margins for the unhedged correlation risk.** Knowing the sign and magnitude of correlation sensitivity is again necessary in this case.

Some multi-asset derivatives are convex in correlation, meaning that the second-order effect on the price from a movement in the correlation is non-zero and needs to be taken into account.

5.10 Computational Methods

There are many different ways to compute the greeks, but we will focus on 3 popular methods:

1. Finite-difference approximations
2. Pathwise method
3. Likelihood Ratio method

5.10.1 Finite-difference approximations

They involve calculating the price for a given value of a parameter, then changing the parameter value slightly, by ϵ , and recalculating the price.

Let f be the payoff function and θ the parameter we are interested in.

Forward difference method

An estimate of the sensitivity will be:
$$\Delta = \frac{f(\theta + \epsilon) - f(\theta)}{\epsilon}$$

Backward difference method

An estimate of the sensitivity will be:
$$\Delta = \frac{f(\theta) - f(\theta - \epsilon)}{\epsilon}$$

Central difference method

An estimate of the sensitivity will be:
$$\Delta = \frac{f(\theta + \epsilon) - f(\theta - \epsilon)}{2\epsilon}$$

Advantages and Disadvantages

Advantages

- It is **easy to implement** and does not require too much thought.
- **Monte Carlo simulation can be used.**

Disadvantages

- It is a **biased estimator** of the sensitivity.

This is reduced by using the central difference method. However, this method is slower as there are three prices to calculate.

- There is an **issue with discontinuous payoffs**.

Let us explain further this last disadvantage using a specific discontinuous payoff such as a Digital Call.

If we want to estimate the delta of a digital call, we will get it being zero except for the few times when it will be 1. For these paths, our estimate of delta will be very big, approximately of order ϵ^{-1} .

5.10.2 Pathwise method

It gets around the problem of simulating for different values of θ by first differentiating the option's payoff and then taking the expectation under the risk-neutral measure:

$$\Delta = e^{-rT} \int \frac{\theta_T}{\theta_0} f'(\theta_T) \Phi(\theta_T, \theta_0) d\theta_T$$

Where Φ is the density of θ in the risk-neutral measure.

Advantages and Disadvantages

Advantages

- It is an **unbiased estimate**.
- It only requires simulation for one value of θ and is **usually more accurate than a finite-difference approximation**.

Disadvantages

It becomes **more complicated when payoff is discontinuous**. To get around this we can rewrite $f = g + h$, with g being continuous and h piecewise constant.

5.10.3 Likelihood ratio method

Similar to the pathwise method, but instead of differentiating the payoff we differentiate the density Φ :

$$\Delta = e^{-rT} \int f(\theta_T) \frac{\Psi}{\Phi}(\theta_T) \Phi(\theta_T) d\theta_T$$

Where Ψ is the derivative of Φ by θ .

Advantages and Disadvantages

Advantages

- Only one value of θ needs to be simulated to calculate both the price and the sensitivity.
- There is **no need to worry about discontinuities in the payoff function** as we are differentiating the density.

Disadvantages

The **density needs to be explicitly known**.

Chapter 6

All about Volatility

6.1 The many flavours of volatility

6.1.1 Some facets of volatility

There are many ways to look at volatility.

The intuitive meaning of the word is that **volatility measures the level of fluctuations for a particular price.**

The way we measure it, the unit we use, the time-scale at which we are looking all have an impact and should be specified in order to transform a single volatility number into a solid understanding of how much level of fluctuation there is in that particular stock.

By taking a more academic approach based on statistics, one can argue that **the value of the stock in one year is uncertain and assign a probability distribution to it.** It could be desirable to **use the width or standard deviation of this distribution to link to the volatility of the stock.** This point of view is exactly what has become **the market standard.**

It is clear that the statistical approach is focused on the one-year horizon, whereas a trader, who wants to delta-hedge on a daily basis, is not so interested in knowing the uncertainty accumulated over the year. What he is really interested in is to understand how the uncertainty plays a role on a much smaller scale, such that piled up over the year it leads the same distribution as the statistician has presented.

In mathematical terms, knowing the distribution at one time (or multiple times) is not enough to complete the dynamic picture. One needs to know how the distribution changes over time. Clearly, on a very short time-scale, the uncertainty is very small and the distribution function should be sharply peaked about the

current level of the stock. As the time horizon increases, the density should widen up.

One can show that **at any time t , the solution of the Black-Scholes SDE, describing a model for the movement of the stock, is a random variable $S(t)$ that behaves according to a lognormal distribution.** So **at any time t , we have a density that depends on the original parameters in the equation**, being the drift and the volatility parameter σ .

Note that **the volatility is not the standard deviation of this distribution, but it does control the wideness of the distribution.** Clearly, if we used another model for the stock price, it would lead to another family of density functions, and to other formulas for the moments.

In a way, when people use the word volatility, they also agree on the underlying mathematical model!

A higher volatility means more uncertainty about the size of an asset's fluctuations and, as such, it can be considered a measurement of uncertainty.

Volatility is dynamic and changes a great deal over time. It experiences high and low regimes, but it also has a long-term mean to which it reverts. Also, as a stock market witnesses a large decline, volatility tends to shoot up: we therefore generally see a negative correlation between such assets and their volatilities.

6.1.2 Realized Volatility

This is probably the **most common volatility measure**.

One could imagine selecting a stock and a certain time period from the past, and trying to estimate the σ parameter in the Black-Scholes model based on this data.

This requires knowledge of **Ito's formula**, which allows us to transform the Black-Scholes equation into a more suitable format. **The solution of this equation is following a lognormal distribution. So the logarithmic of the stock price follows a normal distribution.**

By applying Ito's lemma we can write down the dynamic this quantity follows:

$$d\log(S_t) = \left(\mu - \frac{\sigma^2}{2}\right) dt + \sigma dW_t$$

This equation is saying that **the change in logarithmic stock price is composed of two parts:**

The drift term, $\left(\mu - \frac{\sigma^2}{2}\right) dt$, is proportional to the period of time over which we observe this change.

The volatility term, σdW_t , which is determined by noise (Brownian motion).

If one wants to estimate the σ parameter of the stock, one can use the lognormal returns over one day and calculate the standard deviation from them. This would give the volatility over one day, which is $\sigma\sqrt{dt}$.

This follows from the property of the Brownian motion that tells us that the variance of W_t is given by the time t . That just leaves us with the normalizing effect to withdraw the value of σ .

The first obvious question is how many days there are in a year. In your dataset, there are no quotes for weekends and holidays. So there is no volatility to observe there. Therefore a common approach is to use the number of trading days in a year.

To move from the standard deviation of daily log returns to an annualized volatility therefore requires you to multiply by $\sqrt{252} = 16$. So **a volatility of 16% equates to a standard deviation of daily moves of around 1%.**

A fair challenge would be to ask if there is really no volatility during the weekend. In other words, is it fair to regard the return from Friday to Monday in the same manner as from Monday to Tuesday? Wouldn't that imply that the world stops turning in the weekend? For that matter, filtering techniques are sometimes applied to estimate the volatility parameter.

Also, it makes a difference which frequency of data you use. Typically, the longer the time period, the more normal the returns tend to be.

6.1.3 Implied Volatility

This is a **more market-related volatility concept**. This takes everything to the next step, taking into account the options market.

By observing the price of the option, one can back out the σ parameter one has to push into the formula in order to find that price. The market has adjusted for the shortcomings of the Black-Scholes model and the market-implied distribution is not lognormal anymore.

However, the beauty of the Black-Scholes formula is that you can tune your σ parameter such that you match this market price of the option.

“Implied volatility is the wrong number you put in the wrong formula to get the right price.”

Riccardo Rebonato

Vanilla options are quoted in terms of their implied volatilities since this, or a given price, amounts to the same information. The implied volatilities are in fact the market's consensus on the forward-looking volatility of the asset. This implied volatility incorporates the forward views on all market participants on the asset's volatility.

Sometimes analysts try to use implied volatility to defer conclusions about market direction. It is very tempting to think that this information is present in the option market, but the motivation for buying an OTM put does not have to be because the buyer is expecting a decrease in the stock price.

From the trader's point of view, knowing there is a huge interest in these options can bring about two thoughts:

First, he should increase his price. This is basic feature of supply and demand.

Second, if the scale of the orders becomes really big, he might become less comfortable with the risk and he might want a bigger premium for that.

This clearly means that the implied volatility does not imply anything for the future direction of the market. There are extensive studies on comparing the realized volatility to the implied volatility.

Implied volatility and realized volatility do not necessarily coincide, and although they may be close, they are typically not equal.

Using the correct implied volatility of an asset allows one to price other derivatives on the asset, in particular those that are not liquidly traded. Where the implied volatility of an asset cannot be implied from traded instruments, one may resort to using the realized volatility as a proxy for implied volatility to get an idea of what volatility would be correct to use. In contrast, the realized volatility of an asset can be used as a sanity check to ensure that the implied volatilities being used make sense.

The two are different, with implied volatility generally being higher than realized volatility, but too far a spread could imply a mistake, or if correct, an arbitrage opportunity.

6.1.4 Hedging Volatility

What volatility to use in the hedging procedure dictated by the Black-Scholes model?

It may seem a stupid question at first, because it has already been answered by the implied volatility. If you can price a particular option, meaning you have the implied volatility available, it is straightforward solution to use this value as the volatility to hedge with.

However, let assume you have a crystal ball and you know that the realized volatility is going to be 20%, but you can negotiate your client into paying an implied volatility of 25%. So you are selling the option more expensively that it is really worth.

So which of the two values do I plug into the dynamic hedging strategy to hedge this sold option?

As it turns out, it does not make that much of a difference anyway. This brings us to another strength of the Black-Scholes model. It turns out that the model is so robust that almost any value will do. Most traders choose to use the implied volatility for consistency reasons.

6.2 The Volatility Surface


6.2.1 Introduction

Vanilla options are quoted in volatility terms.

For this to work, both counterparties have first to agree on the values of the inputs to the Black-Scholes equation. The volatility one must plug into the Black-Scholes formula to get the true market price of a vanilla option is called the implied volatility.

In liquid markets, brokers will quote fairly tight two-way prices for vanilla options at several strikes and several maturities.

Plotting the associated implied volatilities in a **three-dimensional plot** results in what is called **the volatility surface**.



6.2.1.1 Strike Dependence: Smile/Skew

In Black-Scholes, a single constant volatility is used for the stochastic process followed by the spot. It means that options with different strikes would have the same volatility. In other words, the skew is flat.

In reality, it is not at all the case since volatilities for strikes that are far ITM or OTM are typically higher than the ATM volatility.

In FX markets, the implied volatility curve is usually **quite symmetrical** around ATM strike. We speak about **the smile**.

In Equity and fixed-income markets, the implied volatility curve is often far from being symmetrical, but heavily **skewed in one direction**. We speak about **the skew**. To say there is a skew means that European options with low strikes have higher implied volatilities than those with higher strikes.

Markets determine vanilla prices, which in turn determine implied volatilities.

6.2.1.2 Time Dependence: Term Structure

Implied volatility for a given contract also depends on its expiry date, T . We therefore also have an **implied volatility term structure**.

6.2.2 Trading the Term Structure of Implied Volatilities

For a given strike, implied volatilities vary depending on the maturity of the option. **In most cases, the term structure is an increasing function of maturity**. It is generally the case in calm periods where short-term volatilities are relatively low.

This curve could be decreasing if the market is volatile and short-term volatility is exceptionally high.

This term structure can also reflect the market's expectations of an anticipated near term event in terms of the volatility that such an event would imply. The term structure also reflects the mean-reversion characteristic of volatility.

One can take a view on the term structure's shape. A simple trade to provide this is the **calendar spread**, which is the difference of two call options of the same characteristics but different maturities.

6.2.3 Why a Smile/Skew?

We have seen that Black-Scholes is not the true process followed by the underlying. Therefore using Black-Scholes formula to back out the volatility given prices in the market will not give a constant number.

So what causes the volatility smile?

Supply and Demand of Vanilla Options

Every market has its own participants with their own behavior and risk profile. Therefore the volatility patterns are particular for each asset class.

FX Markets We have already said that volatility smiles in FX markets are often fairly symmetrical, ex: EURUSD. This is quite intuitive because euro investors see the market as the inverse of the way dollar investors see it.

This explains the symmetry but not why do far OTM/ITM options have higher implied volatility than ATM options?

Well, investors usually want to protect themselves from adverse moves in the FX rate. So those people for whom a drop in the FX rate would be bad buy OTM put options (low strikes) to protect themselves. Meanwhile, those for whom an increase in the FX rate would hurt buy OTM call options (high strikes) to protect themselves. Since there is greater demand from buyers than sellers, the prices are a little higher than you would otherwise expect, and as price increases with volatility, this means the volatilities at low and high strikes are higher.

Equity Markets A similar argument can be used in equity markets to show why the volatility smile is heavily skewed. **In equity markets, investors typically need to protect against decreases rather than increases in the index.**

The skew is often explained as this **concept of insurance**. It is a market where operators cover their downside risk. The market also tends to consider a **large downward move in an asset to be more probable than a large upward move**. A downward jump also increases the possibility of another such move (**volatility clustering**), again reflected by higher volatilities. Additionally, one can discuss the **leverage effect**: a leverage increase given by a decline in the firm's stock price, with debt levels unchanged, generally results in higher levels of equity volatility.

It is therefore not surprising that this skew phenomenon emerged after the 1987 market crash.

Commodity Markets Commodities often exhibit an **inverse or positive skew because the risk sits on the upside**. If commodities prices rise very sharply, this becomes a risk for industrials who process these commodities into their final products. They are often looking for protection on the upside, flipping the story around.

Fixed Income Markets Interest rates have their own behavior depending on what instrument is being considered.

The true underlying dynamics are not Black-Scholes

It expresses the fact that market participants are well aware that the returns are not Gaussian.

Both suggestions are good ways of understanding why the volatility smile is as it is. Investors' view of the market affects the way they trade in the spot as well as the vanilla hedges they put on. By no-arbitrage arguments, if it seems clear that volatilities should be higher at certain strikes for supply and demand reasons, then the true market dynamics must reflect this, and vice versa.

Think in terms of realised gamma P&L

The reason behind the skew becomes apparent when thinking in terms of realized gamma losses as a result of rebalancing the delta of the option in order to be delta hedged.

In downward spiraling market, gamma on lower strike increases, which combined with a higher realized volatility causes the option seller to rebalance his delta more frequently, resulting in higher losses for the option seller.

Option sellers want to get compensated for this and charge the option buyers a higher implied volatility for these options.

One thing that can certainly be said is that **inverted smiles are rare.**

Since the payout of a put option increases as strike increases it must be true that the value of a put option increases with strike, and the opposite is true for calls. This puts a constraint on the shape of no-arbitrage smile curve.

6.2.3.1 Misconceptions about the Skew

When markets go down, they tend to become more volatile. While this statement is true, it does not explain the skew as this realized volatility is the same regardless of any strike price!

The existence of skew is actually saying that **this increase of volatility has a bigger impact on lower strike options than on higher strike options.**

As stated earlier, it is very tempting to use the implied volatilities to predict the direction of the market. However, this is extremely difficult, if not impossible (see this nice article from Elie Ayache). It brings us back to the very basics of the Black-Scholes model. The price is set by the anticipated cost of hedging. If anything, this cost is related to the magnitude of moves or volatility rather than the direction of the market.

6.2.4 Measuring and Trading the Implied Skew

The first thing in measuring the skew is to note **its level, which is given by the ATM implied volatility.** The word skew is also used to refer to **the**

slope of the implied volatility skew. Equity markets have a negative skew since his slope is negative.

Assuming we had the set of implied volatilities as a function of strike $\sigma(K)$, then the slope is given by the first derivative, at a specific point, possibly the ATM point.

In reality, we only have implied volatilities for a discrete set of strikes. One can use some form of interpolation to obtain the function $\sigma(K)$ in order to have a parametric form, but **in practice**, and to have a standard method of measuring skew, **we take the difference between the implied volatilities of the 90% and 100% (or 110%) strike vanillas.**

If we compare the implied volatility skews of an index and that of a stock we find that **index volatilities are more skewed than those of a single stock.** The reason for this is that if stocks are all dropping during a market decline, the realized correlation between them rises, and an equity index is a weighted average of different stocks.

This is a useful property as one can use the skew of an index as a proxy for pricing skew-dependent payoffs on stocks whose implied skews are not as liquid as those of the index. Knowing that the index's implied volatilities are more skewed than those of the single stock, it is possible to take a percentage of the index skew and use this in the pricing. What percentage to use is primarily a function of whether the structure in question sets the seller short skew or long skew, and from there it is a function of how aggressive/conservative the trader wants to be on the skew position.

6.2.5 Skew through Time

The skew for any specific stock is steeper for short-term maturities than for long-term maturities. A long maturity could have a skew at a level higher than the short-term maturity, but the short-term skew will be more pronounced.

To understand this, remember that skew is mainly there because traders are afraid to lose money on downside strikes in case the market goes down and becomes more volatile.

Obviously, the larger the gamma the more imminent the problem! **Since short-term downside options have larger gammas when the stock price moves down to the lower strikes, the effect of skew is largest for those options.**

Another reason is that **for short-term maturities the trader exactly knows whether an option is a downside option or not.** For long-term options, trader cannot qualify whether it is a downside strike, as the trader does not know where the stock will be trading in X year time.

A jump in the underlying's price in the immediate future would have a large impact on the price of the put; for the short term this is more severe as the market may not have time to recover.

6.2.6 Effect of Skew on Delta Hedging

The skew curve of a particular stock can have a big impact on a trader's delta hedge against any option positions he has.

Let us take a simple example and assume a trader is long a 1-year Call 120% on a stock. The skew curve for the 1-year maturity indicates that every 10% decrease in strike translates into a 1.5% increase in implied volatility. In other words, if ATM volatility is 20%, then the 120% volatility is 17%.

Suppose instead that the trader decides to mark this 120% implied volatility at 21%, implying a smile.

Because of that, his delta is larger than it should be as he assigns too high a probability of the option expiring in-the-money. Therefore the trader sells more shares than he actually should to hedge the upside call. To make matters worse, he will continue to sell too many shares with the spot increasing. With the spot increasing, the stock is likely to realize even less, making the probability and therefore the proper delta of this option even lower.

The way a trader marks his implied volatility surface has a large impact on his delta hedge. By marking upside strikes on a higher implied volatility than ATM implied volatility, the trader sells too many shares, and especially when the spot increases, the trader not only loses money on extra shares he sold, but he also continues to sell too many shares.

It is these dynamics that force a trader to mark his implied volatility surface per maturity with a skew that has a downward sloping shape rather than a smile.

Let us speak a bit more about delta-hedging with the smile.

Since volatility is an important determinant of hedge ratios, incorrect volatility assumptions may lead to incorrect deltas. If volatility is time-varying and correlated with the price changes of the underlying asset, the delta must control not only for the direct impact of the underlying price change on the option price, but also for the indirect simultaneous change in volatility.

Intuitively, an inverse relationship between volatility changes and stock returns would suggest that the BS delta is too large.

$$\frac{dC}{dS} = \frac{\delta C}{\delta S} + \frac{\delta C}{\delta v} \frac{\delta v}{\delta S}$$

where:

- $\frac{\delta C}{\delta S}$ is BS delta assuming constant volatility

- $\frac{\delta C}{\delta v}$ is BS vega
- $\frac{\delta v}{\delta S}$ describes how the IV of the option moves as the spot price moves.

The second part of the equation is sometimes called ‘*shadow delta*’.

I usually speak about the whole delta as a *smile-adjusted delta*.

Since the vega of a long position in option is always positive, the above equation shows that **in the case of a negative correlation between stock returns and volatility changes (usual downward sloping equity skew), the delta should be smaller than the BS delta.**

Note that the last term of this equation is quite difficult to quantify but could be approximated by the slope of the volatility smile: $\frac{\delta v}{\delta K}$.

By approximating $\frac{\delta v}{\delta S}$ by $\frac{\delta v}{\delta K}$, it is assumed that as S changes by one unit, there is a parallel shift of $\frac{\delta v}{\delta K}$ units in the volatility smile.

So, **part of the smile-adjusted delta is from vega.** If one uses a lognormal model for pricing, then one needs to properly incorporate the delta due to vega in one’s delta hedging, otherwise one is not truly delta neutral.

Similarly, one needs to properly handle one’s volatility hedging and recognize that part of the vega is hedged by the delta.

If you are interested in reading further about this subject, I invite you to read the original article from Sami Vahamaa, ‘Delta hedging with the smile’.

6.2.7 The Smile Curvature

The curvature is the last parameter that is used to mark an IV surface.

For very high strikes, the implied volatility does not decrease any longer but flattens out.

At the same time, **very low strikes have an even higher implied volatility than the skew parameter indicates.** This is exactly what a trader would want. As very low strikes have very little premium, sellers want to get properly compensated, which means that the skew parameter alone will not be enough and the curvature parameter will ensure these low strikes are marked on a larger implied volatility.

6.2.7.1 Measuring and Trading the Implied Skew’s Convexity/Curvature

To quantify the **skew convexity**, one can consider the sum of the 90% and the 110% implied volatilities minus twice the 100% strike volatility and dividing by the difference in strikes squared. This makes sense as an approxima-

tion of the second derivative of a function f at the point x is indeed given by $\frac{f(x+h)+f(x-h)-2f(x)}{h^2}$.

In fact the combination of vanillas with the above strikes is known as a **butterfly spread**. If we go long a butterfly spread, we are long a 90% and a 110% strike call option, meaning that we are long the implied volatilities at these two strikes. If the implied skew becomes more convex, it means that these two implied volatilities have increased, making the butterfly spread more valuable.

The implied volatilities of a single stock generally have more curvature than those of an index. The reason for this is that downward jumps have a larger impact on single stocks than they do on an index, and the risk of a single stock crashing completely is greater than that of a whole index doing so. So, although a stock may have less negatively skewed implied volatilities than an index, the former's implied volatilities are more convex in strike than those of the index.

6.2.7.2 Smile curvature through time

Smile curvature tends to decrease as maturities increase. Smile curvature decreases as an inverse of time. This property is observed in all equity markets.

This is the expression that the **risk on short-term volatility is greater than over the long-term**. To match this observation, variance at all times must be equal. This can only be achieved with a high mean reversion process and high volatility of volatility.

6.2.7.3 Put vs Call curvature

Put curvature is higher than call curvature. This is certainly due to the fact that puts are used as a protection against default events.

6.2.8 Arbitrage Freedom of the Implied Volatility Surface

In practice we can only observe European option implied volatilities, of a fixed maturity, at a finite set of strikes: $K_1, K_2, \dots, K_m$.

It is also the case that we can only obtain these skews for a finite set of maturities: $T_1, T_2, \dots, T_n$.

For an implied volatility surface to be arbitrage free, some criteria must be met:

1. For all maturities, **all call spreads must be positive**

$$\boxed{C(K_j, T_i) - C(K_{j+1}, T_i) \geq 0}$$

2. An additional restriction on such spreads is that if we were to divide by the difference in strikes, we must have:

$$\frac{C(K_j, T_i) - C(K_{j+1}, T_i)}{K_{j+1} - K_j} \leq 1$$

3. All Calendar spreads must be positive

$$C(K_j, T_{i+1}) - C(K_j, T_i) \geq 0$$

4. All butterfly spreads must be positive

$$C(K_{j-1}, T_i) - \frac{K_{j+1} - K_{j-1}}{K_{j+1} - K_j} C(K_j, T_i) + \frac{K_j - K_{j-1}}{K_{j+1} - K_j} C(K_{j+1}, T_i) \geq 0$$

The set of European options will be arbitrage free if all these conditions are met.

We should concern ourselves that any model we use to capture skew observes these conditions. The failure of a model's calibration to meet these conditions is a solid criterion to reject such calibration. Any interpolation between the implied volatilities of two consecutive strikes in the above set must also observe these conditions to be arbitrage free.

6.2.9 Smile implied probability distribution

Without going to much into details here, we can actually make some progress without any knowledge of the true dynamic causing the implied volatility smile.

Suppose we are interested in a particular expiry T because we have a European contract only depending on the spot at T , and not on the path the spot takes between now and T .

Assuming this contract has a payout function $A(S_T)$ so that to price it, we

want to calculate: $E[A(S_T)] = \int_0^{\infty} A(S) p_T(S) dS$ where p_T is the probability density function for the spot level at time T under the risk-neutral measure.

Skipping the demonstration here, it happens that when we know the implied volatility smile at a given expiry, we can deduce the risk-neutral probability density function. Then we can easily go back to the above equation and use this density function to price any European contract.

Using vanilla prices this way to determine the probability density is known as the Breeden and Litzenberger approach.

6.2.10 Implied Volatility Dynamics

There is a **natural order of market data speed**:

- **Spot levels change faster than ATM volatility**
- **ATM volatility changes faster than volatility skew**
- **Volatilities are more volatile than dividend forecasts**

As we have seen when speaking about the effect of skew on delta hedging, **hedging performance can be improved by assuming a link between different market parameters**.

For example, when calculating a price with a new spot, or computing the delta using a spot shift, one may assume that this move is accompanied by a volatility move in the opposite direction or a change in expected dividends in the same direction.

Thus, a delta hedge also hedges part of vega if stock and volatility are correlated. If delta and vega are hedged separately one has to be careful not to double count vega exposure!

This section discusses **deterministic smile dynamics that assume that the implied volatility surface depends on spot only**. Thus there is a function, which denotes the implied volatility surface observed at time t if spot level is S_t .

There are two important special cases: **sticky strike** and **sticky delta**.

6.2.10.1 Sticky Strike

Sticky strike implies that the volatility associated with a given strike does not change when the spot moves.

The dynamics of vanilla options are thus described by the Black-Scholes model, which is the only complete model with sticky strike dynamics.

The ATM volatility for equities around the ATM strike behaves as a sticky strike movement.

6.2.10.2 Sticky Delta

Sticky delta implies that the volatility associated with an option with a given delta does not change when the spot changes.

In this case, the dynamics of vanilla options are depending on moneyness and term.

The only complete models with sticky delta dynamics are those assuming independent returns.

For currencies, behavior is sticky delta.

Volatility reacts slowly on spot moves. **In quiet markets**, volatility is quoted by strike and is updated much less frequently than spot. **The dynamics resemble sticky strike.**

When markets are volatile then implied volatilities will be updated more frequently and **dynamics may resemble sticky delta.**

Realistic models should exhibit stochastic implied volatility dynamics, in the sense that the smile dynamics may allow for sticky strike and sticky delta dynamics, as well as random changes between the two. To some extent, local stochastic volatility models capture this behavior.

6.2.11 Stylized Facts and Modelisation

Mean Reversion

Implied volatility tends to mean revert around an average level. **Model:** Orstein Uhlenbeck

Smile slope decreases as $\approx \frac{1}{\sqrt{T}}$

Volatility slope behaves as the ATM implied volatility. **Model:** stochastic volatility with two factors in order to control separately ATM volatility and the skew.

Smile curvature decreases as $\approx \frac{1}{T^\alpha}$

Volatility curve behaves as the ATM volatility with a different speed of mean reversion. **Model:** a two factor stochastic volatility enables this type of control.

Put curvature is higher than call curvature

There is dissymmetry between calls and puts. **Model:** Jump model allows the generation of put prices more expensive than call prices.

Smile dynamics

When the spot vibrates, volatility ATM approximately follows the smile (sticky strike). **Model:** a mixture of model between local volatility and stochastic volatility allows this type of behavior.

6.3 Review of Volatility Models

6.3.1 Small Historical Review

The pioneer

The long story of option pricing began in 1900 when **Louis Bachelier** developed the earliest known analytical valuation for standard options in his PhD thesis dissertation, ‘The Theory of Speculation’. Finding that stock price changes looked like a random walk process, Bachelier made the **quite revolutionary assumption that stock prices follow an arithmetic brownian motion**.

Bachelier discovered the immensity of a world in which randomness exists. After his thesis, he proposed a theory of ‘*related probabilities*’. A theory about what would, 30 years later, be called **Markov processes**. Bachelier’s work was the starting point of a major study by Kolmogorov in 1931.

While Bachelier was on the right track, **his formula had clear drawbacks**:

- It did **not take into account any discounting**.
- It **allowed for negative stock prices**.
- It **allowed for option prices superior to the prices of the underlying securities**.

However, due to the precociousness of his work, it took more than 60 years of research to propose any alternative option pricing models.

Before Black-Scholes

It is **Sprenkle** in 1961 who first extended the work of Bachelier by **switching to a geometric brownian motion (GBM) process for the stock price process**. This adaptation did not receive much attention despite ruling out negative prices by assuming the log normality of returns. The reasons often put forth are the considerable number of parameters to estimate and the lack of information about how to do so.

Three years later, **Boness** improved Sprenkle’s model by considering the time value of money. In 1965, Samuelson quickly made the consideration that an option may have a different level of risk than the underlying stock, concluding that the use of the expected rate of return as a discount rate made by Boness was wrong.

The Black-Scholes model

In 1973, Black and Scholes developed the **first completely equilibrium option pricing model**, which was going to become the **greatest breakthrough**

in the pricing of stock options. A consequence that proved more influential was the **realization that by holding stock and risk-less debt, the option position could be hedged completely in a dynamic nature.** The Black-Scholes model gave a serious impulse to the worldwide trading of options because it provided a widely suitable option pricing method.

Due to its success, more focus has been put upon the Black-Scholes model and its underlying assumptions.

Even though earlier empirical research had already started rejecting this simple hypothesis, **the Black-Scholes model relies on the assumption that stock returns have a log-Normal distribution.**

Indeed, Mandelbrot (1963) and Fama (1965) found that **stock returns exhibit excess kurtosis**, suggesting that returns have a **fat-tailed distribution.** Mandelbrot also documented what is commonly known as the **‘volatility clustering’**: *“... large changes tend to be followed by large changes and small changes tend to be followed by small changes ...”*. This stylized fact clearly violates the independence of returns assumed in the Black-Scholes model.

Furthermore, Fama (1965) and Black (1976) noticed that **large downward movements are generally more frequent than their upward counterparts.** Statistically, this means that the **stock return distribution is negatively skewed.** Black further observed the **existence of a negative correlation between stock prices and volatility**, known as the **‘leverage effect’.**

Additional studies such as Blattberg and Gonedes (1974), MacBeth and Merville (1979) have also excluded the GBM hypothesis by showing that **stock returns are heteroskedastic.** In other words, **the variance of aggregate stock returns changes over time.**

The Black-Scholes formula has been even more questioned after the Black Monday at Wall Street in 1987 since the probability of such an extreme event under the normal distribution is extremely low (less than 1.4×10^{-107}). With investors fearing a reappearance as a result of this market crash, they began putting more value on deep OTM put options. Subsequently, those options were traded at a relatively higher price than ATM puts, ATM calls and OTM calls. Their volatilities were therefore higher resulting in a **‘volatility smile’ that contradicts the Black-Scholes model under which the implied volatility surface is flat.**

Jump diffusion models

Merton (1976) and Cox and Ross (1976) were the first to **allow the stock to jump ‘up’ or ‘down’, engendering a discontinuity in the stock price process.** Using adequate parameters, Merton’s model was **able to generate a lots of volatility smiles and skews.** Particularly, choosing a negative mean

for the jump process can readily capture short-term skews. Simultaneously, the model retains the undesirable independence property. Numerous studies on jump diffusion models have been undertaken since that time.

Local volatility model

TBD...

Stochastic volatility models

The heteroskedasticity in stock returns makes it very tempting to **express the volatility as a stochastic process**.

Based on a body of work on stochastic volatility models (Scott (1987), Hull and White (1987), Stein and Stein (1991)), Heston (1993) advanced the first stochastic volatility model with a generalized solution. His model permits the capturing of essential features of stock markets, namely the **leverage effect**, the **volatility clustering** and the **tail behavior of stock returns**. However, it **cannot yield realistic implied volatilities for short maturities**.

Stochastic volatility with jumps in stock price process

Bates (1996) and **Scott** (1997) have associated a jump diffusion model with stochastic volatility.

By benefiting from the advantages of both the jumps in the stock price process and the stochastic volatility, those models seem more capable to match the market facts.

Stochastic volatility with jumps in stock price and volatility processes

Although models nesting both stochastic volatility and jumps have shown some success, Bates (2000) and Pan (2002) indicate that they are **still incapable of fully capturing the empirical features of stock index options prices**. Actually, the **significant volatility smile of index option prices cannot be described solely based on the degree of volatility of volatility**.

Several researchers have proposed to incorporate further jumps in the volatility process to amend inaccurate descriptions of significant volatility smile. Duffie, Pan and Singleton (2000) models volatility as an affine process that can jump up violently and can justify brutal and lasting market changes with upward movements in volatility. Nonetheless, their model cannot subsequently jump down as observed in the data.

Research that followed then tried to mimic volatility spikes.

I will only name Professor Zerilli as she was my teacher back in 2013 and it is the only one I remember to be honest :). In 2005, Zerilli proposed Normal jumps in an innovative log-variance process that follows an Ornstein-Uhlenbeck process. Her results revealed that mimicking volatility spikes improved the option pricing model considerably.

In the following sections, we will discuss the main models: Black-Scholes, Dupire's Local Volatility, Heston et SABR.

We can always speak about more models but choosing means eliminating!

6.3.2 Derivation of Black-Scholes PDE

We have already spoken about the Black-Scholes model but we will further derive the Black-Scholes equation in this section.

In the Black-Scholes world, the evolution of the stock price S is given by:

$$\frac{dS_t}{S_t} = \mu dt + \sigma dW_t \quad \text{for } \mu, \sigma > 0.$$

Assumptions:

- stock does not pay any dividends
- no transaction costs
- continuously compounding IR is $r > 0$ (constant).

The later assumption implies the evolution of the risk-free asset B_t is given by:

$$\frac{dB_t}{B_t} = r dt$$

It means that the risk-free bank account grows at the continuously compounding rate r and hence: $B_t = e^{rt} \rightarrow \frac{dB_t}{B_t} = r dt$.

We are interested in pricing an option which is a function of the stock price at time $T > 0$, S_T , a call option for example.

- The form of the payoff is not so important
- The fact that it is a function of the stock price at time T , and only time T , is important.

Under this condition, we can show that call price is a function of current time t and current stock price S_t only.

To price a derivative in B-S world, we must do so under a measure which does not allow arbitrage $\rightarrow$ **Risk-neutral measure**. One can show that under this

measure, the drift term of the stock price changes so that: $\frac{dS_t}{S_t} = r dt + \sigma dW_t$.

In risk-neutral world, $\frac{C(t, S_t)}{B_t}$ is a martingale and hence if we calculate its differential, we know it must have zero drift. A simple explanation of its meaning is that **we expect it to have zero growth**. Our option price is expected to grow at same rate as bank account and growth of each cancels out in the given process. This is what it means to be a martingale. We do not expect change over time so we have zero expected growth. **This translates to the discounted price having a zero drift term.**

Applying Ito's lemma to $C(t, S_t)$ gives:
$$dC(t, S_t) = \frac{\partial C}{\partial t} dt + \frac{\partial C}{\partial S_t} dS_t + \frac{1}{2} \frac{\partial^2 C}{\partial S_t^2} (dS_t)^2.$$

Under the risk-neutral dynamics of S_t and recalling that:

- $(dW_t)^2 = dt$
- $dW_t dt = dt^2 = 0$

$$dC(t, S_t) = \left(\frac{\partial C}{\partial t} + \frac{\partial C}{\partial S_t} r S_t + \frac{1}{2} \frac{\partial^2 C}{\partial S_t^2} \sigma^2 S_t^2 \right) dt + \sigma S_t \frac{\partial C}{\partial S_t} dW_t$$

Using the **Ito product rule**:

$$d\left(\frac{C(t, S_t)}{B_t}\right) = \frac{1}{B_t} \left(\frac{\partial C}{\partial t} + \frac{\partial C}{\partial S_t} r S_t + \frac{1}{2} \frac{\partial^2 C}{\partial S_t^2} \sigma^2 S_t^2 - rC \right) dt + \sigma \frac{S_t}{B_t} \frac{\partial C}{\partial S_t} dW_t$$

Since $\frac{C(t, S_t)}{B_t}$ is a martingale and hence must have zero drift:
$$\frac{\partial C}{\partial t} + \frac{\partial C}{\partial S_t} r S_t + \frac{1}{2} \frac{\partial^2 C}{\partial S_t^2} \sigma^2 S_t^2 - rC = 0$$

This is the **Black-Scholes equation**. It is a **partial differential equation (PDE)** describing the evolution of the option price as a function of the current stock price and the current time.

The equation does not change if we vary the payoff function of the derivative. However, the associated boundary conditions, which are required to solve the equation do vary!

The above implies that **two stocks with the same volatility but different drifts will have the same option prices.**

The pricing of any derivative must be done in the risk-neutral measure in order to avoid arbitrage. Under this measure, we have seen that the drift changed and was independent of the drift of the stock.

Financially, this reflects the fact that the hedging strategy ensures that the underlying drift of the stock is balanced against the drift of the option. The drifts are balanced since drift reflect the risk premium demanded by investors to account for uncertainty and that uncertainty has been hedged away.

6.3.3 Dupire's Local Volatility Model

I like the way **Colin Bennett** provides intuition about Local volatility in the appendix of his book 'Trading Volatility, Correlation, Term Structure and Skew'. I am sharing some of it with you here below.

Exotic equity derivatives usually require a more sophisticated model than the Black-Scholes model. The most popular alternative model is a **local volatility model (LocVol)**, which is the **only complete consistent volatility model**.

Complete: it allows hedging based only on the underlying.

Consistent: it does not contain a contradiction.

LocVol models try to stay close to the Black-Scholes model by introducing more flexibility into the volatility.

6.3.3.1 Some intuitions

LocVol models offer a way of capturing the implied skew without introducing additional sources of randomness; **the only source of which is the underlying asset's price that is modeled as a random variable**. In LocVol models, the **volatility is a deterministic function of the asset's level**. In the Black-Scholes model, the asset's price is modeled as a log-normal random variable, which means that the asset's log-returns are normally distributed. However, the fact that we have a skew is the market telling us that the asset's log-returns have an implied distribution that is not Normal.

LocVol is still a **one-factor model** and it **also allows for risk-neutral dynamics**, which means that, like Black-Scholes, the model is still preference free from the financial point. The LocVol model is the simplest one to account for skew and offers a consistent structure for pricing options.

How does Local Volatility work? Well, as we said, the presence of skew is the market telling us that the asset's log-returns are not normally distributed. In fact, **the market is implying some distribution**.


If we are given a set of vanilla options prices for a fixed maturity across strikes, can we find a distribution that corresponds to these prices? In other words, can we find a distribution for the asset price so that if we used such distribution to price vanilla options on this asset, it would give the same options prices as the ones seen in the market?

YES, theoretically, there is a way to find the distribution (LocVol model) which corresponds exactly to all vanilla prices taken from the skew. In fact, **LocVol extends beyond skew and can also capture term structure**. It can therefore theoretically supply us with a model that gives the exact same prices for vanillas taken from a whole implied volatility surface.

Local Volatility is instantaneous volatility of underlying Instantaneous volatility is the volatility of an underlying at any given local point, which we shall call the **local volatility**. We shall assume the local volatility is fixed and has a normal negative skew. There are many paths from spot to strike and, depending on which path is taken, they will determine how volatile the underlying is during the life of the option.

Black-Scholes volatility is average of local volatilities It is possible to calculate the LocVol surface from the Black-Scholes implied volatility surface.

This is possible as the Black-Scholes implied volatility of an option is the average of all the paths between spot and the maturity and strike of the option.



A reasonable approximation is the average of all local volatilities on a direct straight-line path between spot and strike. For a normal relatively flat skew, this is simply the average of two values: the ATM LocVol and the strike LocVol.

ATM volatility is the same for both Black-Scholes and Local Volatility For ATM implied volatilities, the LocVol at the strike is equal to ATM implied volatility. Hence the average of two identical numbers is simply equal to the ATM implied volatility. For this reason, **Black-Scholes implied is equal to LocVol ATM implied**.

Black-Scholes skew is half of LocVol skew as it is the average If the LocVol surface has a 22% implied at the 90% strike and 20% implied at the ATM strike, then the Black-Scholes implied volatility for the 90% strike is 21%.

As ATM implied volatilities are identical for both local volatility and Black-Scholes implied volatility, this means that the 90%-100% skew is 2% for LocVol

but 1% for Black-Scholes. **LocVol skew is therefore twice the Black-Scholes skew.**

6.3.3.2 Deeper look into the model

It is often used to calculate exotic option implied volatilities to ensure the prices for these exotics are consistent with the values of observed vanilla options and hence prevent arbitrage.

In the LocVol model, the only stochastic behavior introduced into the volatility function is a result of it being a function of the underlying asset price (if r_t and q_t are deterministic). So there is **still just one source of stochasticity, ensuring the completeness of the Black-Scholes model is preserved. Completeness is important, because it guarantees unique prices.** This is the stated reason to develop the local volatility model in Dupire's original paper.

From Implied to Local volatilities Since we can look in the vanilla market and find prices or equivalently implied volatilities for vanilla options at any strike and expiry, can we find the LocVol function $\sigma_{local}(S_t, t)$ so that if the spot follows

$\frac{dS}{S} = \mu dt + \sigma_{local}(S_t, t) dW$, then the fair values of vanilla options exactly match the market?

YES, and that gives us a **powerful tool to price exotic options in a model that is consistent with the vanilla market.** Finding this function $\sigma_{local}(S_t, t)$ is a process known as **calibration**. The inputs for these models are not only the current level of the asset, the curve of riskless interest rates, the size and timing of known dividends to come, but also the implied volatility skew (possibly a whole surface). Given the set of implied volatilities of vanilla options, **calibration is the process where we search for these volatilities $\sigma_{local}(S_t, t)$ so that the model matches these prices.**

So knowing the market prices of vanilla options, the LocVol function can be

derived using the following formula:
$$\sigma(K, T) = \sqrt{2 \frac{\frac{\partial C}{\partial T} + (r - q)K \frac{\partial C}{\partial K} + qC}{K^2 \frac{\partial^2 C}{\partial K^2}}}$$

The construction of the LocVol from the implied volatilities constitutes a difficult numerical problem in practice. You can have a look at Gatheral's paper for more details.

Starting from a finite set of listed option prices, a good interpolation in strike and maturities provides a continuum of option prices and we can apply the stripping formula to get the local volatilities.

Once the local volatilities are obtained, one can price exotic instruments with this calibrated local volatility model. **Properly accounting for the market**

skew can have a massive impact on the price of exotics -> example: call up-and-out.

We can complete the LocVol formulation by deriving the PDE to use for pricing. The derivation is identical to that of the Black-Scholes equation and the result

is as follows:
$$\frac{\partial C}{\partial T} = \frac{\sigma^2(K, T)}{2} K^2 \frac{\partial^2 C}{\partial K^2} - (r - q)K \frac{\partial C}{\partial K} - qC$$

The resulting LocVol surface is fully non-parametric. The LocVol model allows a full fitting of an arbitrage-free implied volatility surface.

From local volatilities to implied volatilities Given the LocVol model, **the computation of the implied volatilities is only approximate.** There exists several methods of approximation, one of which is **the most likely path.** It gives intuition as to how the implied volatility is built from local volatility.

In the Monte Carlo simulation approach, we simulate many paths and keep only the ones that finishes around the strike. We obtain a stream of trajectories that start at the initial spot and finish around the strike. We average on each date all these paths and obtain the most likely path. We can also extract the variance around this path. We obtain the implied volatility estimation from it (thanks to the most likely path and the width around it).

This is well explained in Adil Reghai's book 'Quantitative Finance: Back to Basic principles'.

6.3.3.3 Strengths and Weaknesses

Weaknesses

- **Forward skew is smaller than it should be.**
- **Volatility of volatility is small.**
- **Numerical problems in implementation.** There are computational difficulties in finding the LocVol function that will exactly fit all market prices, which is why Dupire's formula, though theoretically correct, has some practical drawbacks. In particular, **fitting all points may lead to unrealistic model dynamics.** In practice, **there may be more than one LocVol model that fits a set of vanillas**, so one must lay down a set of criteria to follow when choosing the model to use. The surface is two dimensional, one in time and one in strike, and the focus on one or both must be determined in order to correctly capture the effect of the volatility surface on certain payoffs.

Despite all of this, 90% of investment banks' production systems were still using LocVol model in day-to-day risk management.

It is important to note that **traders adjust their prices and their greeks if ever the LocVol model is not the adequate pricing model** (for Cliquet options or options on variance for example).

Strengths LocVol models offer a way of capturing the implied skew without introducing additional sources of randomness; the only source of which is the underlying asset's price that is modeled as a random variable.

LocVol is still a one-factor model and it also allows for risk-neutral dynamics, which means that, like Black-Scholes, the model is still preference free from the financial point. The LocVol model is the simplest one to account for skew and offers a consistent structure for pricing options.

The strength of the LocVol model lies in its signature to the product. Only the vega KT map can give precise sensitivities for all vanilla options (K,T). The LocVol model allows a more precise projection of the global vega. It provides a **powerful means to find the right strikes and maturities whereby to project the total vega**. It is also a **P&L explanation for the varied moves in the implied volatility surface**. This tool is used every day dozens of times to explain the impact on the book movements in the vol surface.

6.3.3.4 Which products can/cannot be priced using LocVol model?

LocVol models **can** be used to price options that have **skew dependency** yet cannot be broken down into vanillas, assuming the LocVol model is correctly calibrated to the skew (or surface) in a manner consistent with the skew sensitivity of the option.

LocVol models **cannot** be used to price payoffs that exhibit **vega convexities that are not captured in the skew** since this model has a too small volatility of volatility and therefore will tend to underestimate those payoffs.

Also, LocVol models **cannot** be used to price a derivative that has **exposure to forward skew** as the later is underestimated under this model. Although LocVol models can capture the market's consensus on the prices of vanilla options by matching the volatility surface, the evolution of future volatility implied by these models is not realistic. **Forward skews generated by LocVol models flatten out as we go forward in time, even though, in reality, forward skews have no reason to do so.** The LocVol model therefore does not provide the correct dynamics for products with sensitivities such as these and will result in wrong price and wrong subsequent hedge ratios.

6.3.4 Stochastic Volatility Model : Heston Model

The Heston Model commonly stands out among the stochastic volatility models for several reasons:

- It provides a **closed-form solution for European Call options**.
- It **allows the stock price to follow a non log-Normal probability distribution**.
- It **expresses the volatility as a mean-reverting process**
- It **fits pretty well the implied volatility surface of option prices observed in the market**.
- It **takes into account a possible correlation between the stock price and its volatility**.

The model

$$\begin{cases} dS_t = \mu S_t dt + \sqrt{v_t} S_t dZ_S \\ dv_t = \kappa(\theta_v - v_t) dt + \sigma_v \sqrt{v_t} dZ_v \\ \langle dZ_S, dZ_v \rangle = \rho_{S,v} dt \end{cases}$$

$S(0) = S_0$ where S_0 is the spot stock price.

$V(0) = V_0$ where V_0 is the spot variance.

Model's parameters

Stock price process: S The first stochastic differential equation (SDE) expresses the fact that **the stock price follows a stochastic process with a constant rate of return μ** .

$$\boxed{dS_t = \mu S_t dt + \sqrt{v_t} S_t dZ_S}$$

The presence of the square root in the diffusion coefficient ensures the non-negativity of the variance.

It is $v(t) \geq 0$ for all time with probability one.

Variance process: v It can easily be observed from historical data that volatility changes over time. While volatility is undoubtedly varying over time, it seems to fluctuate around some long-term mean level. Consequently, **the variance is represented by a square root mean-reverting process**, similar to the one used by Cox, Ingersoll and Ross (1985) for modeling the term structure of interest rates.

$$\boxed{dv_t = \kappa(\theta_v - v_t) dt + \sigma_v \sqrt{v_t} dZ_v}$$

The drift term of the SDE of the variance process indicates a **mean reversion** when $\kappa > 0$, with θ_v being the **long-term mean level of the variance**. Basically, whenever v_t is greater (lesser) than θ_v , the drift term will push the process value down (up).

κ is the **speed of the mean reversion of the variance process** and can be thought as the **degree of ‘volatility clustering’**.

σ_v is the **volatility of variance and influences the kurtosis of the distribution**. The larger σ_v , the greater the kurtosis, the fatter the tails of the distribution.

Effects of the different parameters on the implied volatility θ_v

The effect of the long-term mean reversion level is intuitive since **an increase in the long-term mean reversion level of the variance corresponds to an increase of the variance**. Consequently, an increase of θ_v will be associated with an upward translation of the volatility smile.

σ_v

Since the kurtosis of the distribution has an effect on the implied volatility, the volatility of variance σ_v indirectly affects the implied volatility. The Heston model has the pleasing attribute to mimic the volatility smile observed in the market. **The larger σ_v , the more pronounced the smile**. This is rather intuitive since it increases the probability of extreme movements and thus increases the price of OTM Calls and Puts. Note that the term structure of implied volatility is flat when the volatility is constant (as in Black-Scholes), it is when the volatility of variance is set equal to zero.

κ

Intuitively, **the speed of reversion, κ , governs the relative weights of the long-term mean level of the variance, θ_v , and the initial variance, V_0** . It is logical the impact of κ to be dependent on both the initial variance and the long-term mean reversion level. We consider the case in which V_0 and θ_v are equivalent and provide a natural interpretation.

First of all, **the impact of the mean reversion speed of the variance on option prices seems to be rather limited**.

Then, **the effect of κ appears to be different when options are deeply ITM or deeply OTM**. This is very intuitive and is closely related to the fact that **increasing the speed of the mean reversion of the variance decreases the probability of extreme movements**. As a result, when an option is deeply OTM, increasing κ decreases the probability that the option will finish ITM at maturity, therefore decreasing its price. By the same way, when an option is deeply ITM, increasing κ increases the likelihood of the option finishing ITM, therefore increasing its price.

The Feller Condition

While the variance is never negative, it can actually reach zero except if the Feller condition is satisfied: $2\kappa\theta_v > \sigma_v^2$

Intuitively, σ_v cannot be too large and $\kappa\theta_v$ cannot be too small. In practice, this condition is regularly not satisfied so that the likelihood of a zero variance is somewhat important.

Correlation between stock price and volatility: $\rho \quad \langle dZ_S, dZ_v \rangle = \rho_{S,v} dt$

Empirical studies have documented that the stock price and volatility processes are negatively correlated (Black (1976), Christie (1982), Engle and Ng (1993)). Decreasing stock prices inflate the leverage firms have. It is commonly thought that this gives rise to more uncertainty and hence volatility. Those authors also showed that both the value and the sign of the correlation have an impact on the return distribution, more specifically on its skewness. A negative correlation makes the left tail fatter than the right tail. Such negative skewness is a typical feature of the distribution of returns.

Since the skewness of the distribution has an effect on the implied volatility, the correlation */rho* indirectly affects the implied volatility. It is distinctly observed that the correlation indeed affects the shape of the implied volatility curve. In particular, a negative correlation induces a declining sloping curve. This coincides with the previous finding that a negative correlation makes the left tail fatter relatively to the right tail. Indeed, this logically contributes to a higher price for deep OTM Puts. By the Put-Call parity, it is readily found that Calls also inflate in implied volatility.

6.3.4.1 Which products can/cannot be priced using Stochastic Volatility model?

In stochastic volatility models, both the asset price and its volatility are assumed to be random processes. In allowing the volatility to be random, stochastic volatility models give rise to implied volatility skews and term structures. SV models can explain in a self-consistent manner the actual features we see in the empirical data from the market. Once such a model is specified, the skews generated by the model are a function of its parameters, and finding the parameters that fit a certain skew (or surface) is again the act of calibration.

Stochastic volatility models go beyond skew and term structure allowing for vega convexity and forward skew. Any derivative that is sensitive to vega convexity or/and forward skew should be priced using a stochastic volatility model.

A derivative exhibits vega convexity when its sensitivity to volatility is non-linear, meaning that there is a non-zero second-order price sensitivity to

a change in volatility. Vanilla options are convex in the underlying's price, but are they also convex in volatility? Well, ATM vanillas are not, but OTM vanillas do have vega convexity. However, these options are liquidly traded and their prices are obtained by using their implied volatilities in Black-Scholes. These implied volatilities give the market's consensus of the right price; therefore the cost of vega convexity of OTM vanillas is already included in the skew.

In more complex payoffs, almost all the payoffs will exhibit some form of vega convexity, although in many cases this is captured in the skew and can be correctly priced by getting the skew right (with a LocVol model). Other payoffs exhibit such convexities that are not captured in the skew and a stochastic volatility model must be used. **Since volatility is taken to be random, it must have its own volatility, known as the volatility of volatility. This parameter corresponds to the vega convexity term.** Note that the second-order sensitivity to volatility is known as **volga**.

The second feature of stochastic volatility models is that **they can generate forward skews**. If a derivative has exposure to forward skew, one must use a model that knows about forward skews in order to get a correct price. **The dynamics of stochastic volatility models are more consistent (than LocVol model) with the dynamics observed in the market.** By smile dynamics, we refer to the phenomena of how the skew moves as the underlying moves: if the underlying moves in one direction, how should the skew move?

Stochastic volatility models have also their weaknesses as they have **difficulty fitting both ends of the surface**, that is fitting the skew for both short and long maturities at the same time. One remedy for this is to add jumps to a stochastic volatility model. Jumps are able to explain the short-term skew quite well, and we recall that the reason for the existence of the steep short-term skew has to do with jumps. Adding jumps to such a model does not generally affect the long-term skews which remain relatively flatter; the long-term implied skew is not driven by jumps in the underlying.

6.3.5 Stochastic Volatility Model : SABR Model

The SABR model was introduced by Hagan and al in 2002. SABR stands for 'Stochastic Alpha Beta Rho', which are the main variables of SABR equations (α, β, ρ) . It describes a single forward F , such as LIBOR forward rate, a forward swap rate, or a forward stock price. The SABR model is widely used by practitioners in the financial industry, especially in the interest rate derivative markets.

It is easier to write down the model in terms of the forward to a fixed expiry rather than in terms of a spot level.

In terms of F , the Black-Scholes process $\boxed{\frac{dS}{S} = \mu dt + \sigma dW}$ becomes

$\frac{dF}{F} = \sigma dW$ as a simple application of Ito shows.

SABR is practically the simplest extension of Black-Scholes to a stochastic volatility model. It has SDEs:

$$\begin{cases} dF = \sigma F^\beta dW_1 \\ d\sigma = v \sigma dW_2 \\ \langle dW_1, dW_2 \rangle = \rho dt \\ \sigma(0) = \alpha \end{cases}$$

6.3.5.1 Small comparison with Heston

Like the Heston model, the SABR model has a vol-of-vol parameter v that controls the convexity of the implied volatility smile, and a correlation parameter ρ that controls the skew.

Unlike the Heston model, there is no mean-reversion. This is a drawback as it means that when the instantaneous volatility follows a path in which it becomes very large, it is likely to stay large.

6.3.5.2 Beta parameters

β is one of the key parameters and affect many fundamental characteristics of the model. It is the component that determines the shape of forward rates, leverage effect and backbone of ATM vol. It effects the distribution:

$\beta = 1$: stochastic log-normal rates

$\beta = 0$: stochastic normal rates

$\beta = 0.5$: stochastic CIR model (with 0 drift)

Ordinary, say in equity and FX markets, we choose $\beta = 1$ so that the process is approximately log-normal in the limit of small vol-of-vol.

On the other hand, interest rates traders are fond of SABR and do make use of the β parameter (usually 0.5).

If we use $\beta \neq 1$, we could rewrite the process as: $\frac{dF}{F} = \sigma F^{\beta-1} dW$ so that the instantaneous ‘log-normal volatility’ is $\sigma F^{\beta-1}$. In this way, we add some dependence between the forward level to the volatility, and therefore β also impacts the skew.

We can to some extent play off β and ρ against one another.

However, β will also impact the convexity of the implied volatility smile. To see this, you could apply Ito's lemma to calculate the process followed by the effective volatility $\sigma F^{\beta-1}$ and see an additional contribution to the vol-of-vol.

Caution is required in understanding the process when $\beta \neq 1$. In the special case of $v = 0$, for which the volatility is no longer stochastic, the process is known as constant elasticity of variance (CEV). This process is well studied and the following results are provided by Andersen and Andreasen in 1998:

For $\beta \geq 0.5 \rightarrow$ the SDE has a unique solution.

For $0 \leq \beta \leq 1 \rightarrow$ the process can reach $F_t = 0$ but never go negative.

For $\beta \geq 1 \rightarrow$ the process can never reach $F_t = 0$.

For $\beta = 0 \rightarrow$ the process is normal, and therefore F_t can go negative.

For $0 \leq \beta \leq 0.5 \rightarrow$ the SDE only has a unique solution if one adds a boundary condition at $F_t = 0$. For the process to be arbitrage free, the boundary condition must be that when F_t hits zero it stays there.

SABR is popular because Hagan and al. were able to provide an approximate solution that is valid in the limit of small time to expiry T .

6.3.5.3 λ -SABR

For equity purpose, the log-normal SABR is perfectly good (version with $\beta = 1$). Log-normal SABR has a volatility process that cannot hit zero and therefore it does not suffer from the numerical difficulties associated with Heston. Unlike the $\beta < 1$ version of SABR, log-normal SABR also has a spot that stays positive. Therefore the only drawback with the process itself is the lack of mean reversion, allowing paths in which instantaneous volatilities become very large.

This is not such a problem as we can simply add in a mean reversion term to obtain the λ -SABR model:

$$\left\{ \begin{array}{l} dF = \sigma F^\beta dW_1 \\ d\sigma = -\lambda(\sigma - \bar{\sigma}) dt + v \sigma dW_2 \\ \langle dW_1, dW_2 \rangle = \rho dt \\ \sigma(0) = \alpha \end{array} \right.$$

By including mean reversion, we lose the analytical approximation for the implied volatility smile in terms of our stochastic volatility parameters.

We have seen that the SABR or λ -SABR model itself is attractive.

It is the asymptotic implied volatility formula that causes problems. The approximation is only valid for short expiries, and furthermore, it can imply negative probability densities at strikes that are far away from ATM. These are problems for people who want to use the SABR formula as a method to define

or interpolate the implied volatility smile. If you are only interested with smile modeling, these issues will not cause you any problems.

Chapter 7

Classic Options

The Classic Options Pricer allows the user to price and analyze the greeks of three products: European, American and Digital Options (Calls and Puts).

7.1 European Options

European Options and the behavior of their Greeks have already been largely discussed in chapter 4 and chapter 5 respectively.

The Classic Options Pricer offers a perfect opportunity to put all this theory into practice.

7.2 American Options

American options can be exercised at any time during their life. Since investors have the freedom to exercise their American options at any point during the life of the contract, they are logically more valuable than European options.

The Classic Options Pricer prices American Options using the approximation method of Bjerk Sund and Stensland (1993).

7.2.1 American Calls

For American Calls, early exercise may be optimal just before the dividend payment if the dividend payment is large enough.

This can be expressed by the following condition: $\frac{D}{K} > r * (T - t)$.

Intuitively, if one exercises the American call, he pays a specific amount of money to buy the underlying shares. On the one hand, he doesn't receive interest on this cash amount; and, on the other, he would receive future dividends for holding the stocks. In other words, if the dividend yield is higher than the interest rate until maturity, it is optimal to exercise the American call. For stocks not paying dividends, it is never optimal to exercise the American call.

7.2.2 American Puts

Ultimately, it can be optimal for the holder of an American put option to choose to exercise if the interest rate that would be received on a cash deposit equal to K is higher than the dividend payments until maturity. For non-dividend-paying stocks, an American put should always be exercised when it is sufficiently deep ITM.

It is important to realize that it makes no sense to exercise an option when there is time premium remaining because you are throwing away that time premium by doing so. You would be better off selling the option than exercising it.

7.3 Digital Options

Digital options are quite straightforward. They are options that pay a fixed coupon if the underlying is below or above a predetermined level and does not give a payout at all in all other cases.

Digitals are still considered as exotic options as they cannot be perfectly replicated by a set of standard options.


7.3.1 European Digitals

We will focus on Digital Calls but the same reasoning can always be applied in the case of Digital Puts.

7.3.1.1 Payoff and Premium

European Digital Calls pay a fixed coupon C if the underlying spot price at maturity T is higher than a predetermined barrier level, K .


The Digital Call payoff can be expressed as: $C * 1_{\{S_T > H\}}$. This is why we also call them binary Calls.



Under Black–Scholes, the price of such an option is given by the following formula:

$$\text{Digital Call} = C * N(d_2) * e^{-rT}$$

So the Digital Call price is given by $N(d_2)$, which is nothing but the negative of the derivative with respect to K . It gives the probability that the spot at time T is higher than the barrier level.



This shape should be familiar to you by now. It looks like the Delta of a European Call. This is due to the fact that this Delta is given by $N(d_1)$, which is also a cumulative distribution function.

If the shape of the premium looks like the shape of the delta of European call, then the delta of a digital call will be the gamma of a vanilla call. As we will see, this means that we no longer have this bounded delta from zero to one.

As you will see, Digital Calls are used extremely often in the world of structured products as their payoffs generally contain some sort of discontinuity.


7.3.1.2 Replication of European Digital Options

The digital call can be thought of as a limit of a call spread. One can therefore make a good estimate of the price of a digital option by using option spreads.

$$\text{Digital}(K) = \lim_{\epsilon \rightarrow 0} \frac{1}{2\epsilon} (Call(K - \epsilon) - Call(K + \epsilon)) = -\frac{\partial Call(K)}{\partial K}$$

As the distance between the call option strikes and the digital strikes, ϵ , gets smaller, we need $\frac{1}{\epsilon}$ call spreads of width 2ϵ to replicate the digital. In the limit, meaning as ϵ approaches zero, the call spread replicates the digital exactly.

Note that the above expression is theoretical as, in practice, a trader will not center the call spread around the barrier. He will be more defensive and take a call spread that over-replicates the digital as shown below.



7.3.1.3 Hedging a Digital

Well, you should not be surprised if I tell you that the only real way to risk manage the digital option is with option spreads.

You can then hedge a digital call as a call spread. The gearing of the call spread used to over-replicate the digital depends on the strike width of the call spread. The wider the call spread, the lesser the gearing and the more conservative the price.

What do we mean by saying that the call spread over-replicates the digital option?

Let us have a look at Fig 6.3. here above.

Above the barrier level, the call spread has the same payoff as the digital call. Below the barrier level, the digital call has a zero payoff but the call spread has a non-zero payoff between its lower strike and its upper strike located at the barrier level. Therefore, we say that the call spread over-replicates the digital call because its payoff (and therefore its premium) is always greater or equal to the digital call's payoff.

Let us take a small practical example.

As an investor, you buy a 6-month European digital call on AB Inbev which pays 10€ if after 6 months the ABI stock trades above 50€ and pays 0 if the ABI stock trades below 50€ at maturity.

As a trader, I sell you this digital call on ABI stock. How much will I sell it to you? Well, I will replicate the digital using a geared call spread. I believe that

a 2€ wide call spread should be enough for me to risk manage this position. So I will price the digital call as if it was a 48€/50€ call spread that is 5 times geared.


You can think of different scenarii and see that this call spread over-replicates the digital call.

By doing so, I have therefore priced the digital conservatively. I could have been more aggressive by choosing a tighter 49€/50€ call spread. But remember that I have risks to manage, especially gamma and pin risk around the 50€ barrier level.

The smaller the call spread, the more aggressive the price but the more difficult the hedging.

For a digital option, Gamma can be quite large and tricky near the barrier at maturity. Think about the situation where you are just before expiry, the ABI stock trades at 50€ so that the digital would not pay you anything. If ABI stock goes up to 50.02, the digital would suddenly pay you 10€. As a trader, this would be extremely difficult to hedge. As a trader, the call spread gives me a cushion against this risk.

Using a call spread allows to smooth the Greeks. The smaller the call spread, the larger Gamma and Vega can get near the barrier. In fact, around the barrier level, they shoot up and then shoot down while changing sign.



We will analyze Call Spreads in more details in the next chapter. You will see that its gamma is smoother than that of a digital call. The larger the strike width, the more this is true.

You should have understood by now that when I sell a digital call, I actually book and trade a call spread in my risk management system. As the underlying gets closer to the barrier, you still want to be able to manage your delta hedge properly. A large Gamma means that you will have to buy/sell a large Delta of the underlying, which might be difficult in the market. It is the reason why the liquidity of the underlying is an important variable when selecting the strike width of the replicating call spread.

7.3.1.4 Width of the Call Spread and Barrier Shifts

So the width of the option spread is used as a pricing mechanism to go conservative on the price of a digital option over its model price.

It is necessary in the pricing mechanism to account for real-world difficulty in executing large deltas at the barrier that the model does not consider.

The optimal width of the call spread depends on several parameters among which the size of the digital, the size of the nominal, the underlying's liquidity, the peak delta around the barrier and the implied volatility around the barrier.

In practice, some traders rather take a constant shift of the barrier. Basically, it allows them to take an additional margin for managing the risks if the underlying was to get close to the barrier. This can be more efficient when risk managing a large book of exotic options.

When taking a barrier shift, a trader is pricing a new digital whose replicating **centered** call spread is the hedge of the actual digital.

The direction of the barrier shift obviously depends on the trader's position.

We will discuss further about barrier shifts in the chapter on barrier options.

7.3.1.5 Risk Analysis - The Greeks

We will shortly speak about the greeks of a digital call at initiation. Note that the risks and therefore the greeks are dynamics. For example, the greeks will be quite different if you get closer to maturity. As I cannot describe every scenario, the best way for you to learn this material is to use the pricer, ask yourself many questions and find your answer using the pricer. For example, what happens to the greeks if just before maturity the spot price is exactly at the barrier level? You open the pricer, you select Digital call in the option type input, you set the stock price at the strike level and you set the maturity close to 0. You will be able to calculate the greeks and see all the related graphics. You will then have to interpret them. If you have any questions, you can always drop us an email at info@protect.TU.txtddollar@protect.TU.txtddollarderivativesacademy.com

7.3.1.5.1 Delta The holder of a digital call is always long the forward price since a higher forward increases the probability of the option finishing in-the-money.

Being long the forward means being:

- Long interest rate
- Short dividends
- Short borrow costs




Figure 7.1: Fig: 7.5 :Delta of a 1-year Digital Call at initiation


I don't think I am making you a favor if I describe all the graphics with precision. The best way to develop yourself is to decipher these plots by yourself. For example, you should be able to understand why does the delta converges to zero (and not to 1 as in the case of European calls) when the stock price increases well above the barrier level. Note that the plot of the delta is simply the first derivative of the premium plot with respect to the spot price.

Since a digital call has positive delta, the trader selling it will have to buy delta of the underlying. Therefore the trader will be long dividends, short interest rates and long borrow costs of the underlying.

7.3.1.5.2 Gamma While their magnitudes are quite different, Gamma and Vega behave similarly and depend about the position of the forward price regarding the barrier. The Gamma plot can be easily deduced from the Delta plot since it is simply the first derivative with respect to the spot price. Unlike vanilla options, the gamma of digital options change sign around the barrier level. While this change is quite smooth at initiation, we have seen that it gets more spiky closer to maturity. It makes the hedging process particularly difficult for the trader as vega and gamma shoot up and down while changing sign.

Think about being the trader hedging this digital call close to maturity when the spot is around the barrier level. How does this change of sign impact your delta hedging?


This discontinuity risk (gap risk) has been discussed and is the reason why barrier shifts are applied and option spreads are used to smooth it.



7.3.1.5.3 Vega The fact that vega depends on the position of the forward price with respect to the barrier is very intuitive.

The holder of a digital call will be long volatility if the forward price is lower than the barrier level since a higher volatility will increase the probability of the spot finishing above the barrier at maturity. When the forward is lower than the barrier, you can think of the digital call as being out-of-the money. Volatility will increase the probability of the option going from OTM to ITM.

Inversely, the holder of a digital call will be short volatility if the forward price is greater than the barrier level since a higher volatility will decrease the probability of the spot finishing above the barrier at maturity. When the forward is greater than the barrier, you can think of the digital call as being in-the money. Volatility will increase the probability of the option going from ITM to OTM.



7.3.1.5.4 Theta The shape of Theta plot looks completely opposite to the shape of Vega plot. This is because time to maturity has a similar effect to a digital option price as volatility. The effect is not exactly the same as time has always a second effect that comes from the discounting impact, although this last effect is generally less important.

7.3.1.5.5 Rho When we spoke about Rho in section 5.6.1, we said that the effect of interest rate on an option's price came from two effects: the cost of delta-hedging and the discounting.




Figure 7.2: Fig: 7.8 : Theta of a 1-year Digital Call at initiation

It is therefore not very surprising to see similarities between the Delta profile and the Rho profile. Note that the discounting effect is clearly apparent in the right-side of the curve where the option is completely in-the-money and there is no delta left. On that side of the curve, Rho is negative because an increase in interest rates increases the discount factor and therefore decreases the present value of the digital call. How much the Rho is negative will then mainly depend on the time to maturity.




Figure 7.3: Fig: 7.9 : Rho of a 1-year Digital Call at initiation

7.3.1.5.6 Skew Since a trader hedges a digital option using an option spread, the skew risk is a critical consideration.

Let us assume that we just sold a digital call, we will hedge it by buying a call spread. Taking a long position in a call spread means buying a call at a lower strike and selling a call at an upper strike. The skew makes the lower strike implied volatility more expensive than the upper strike implied volatility. Since the skew makes the hedge more expensive, it makes the structure itself more expensive. Remember what we said in section 4.5, an option price is nothing else but the cost of the hedge!

Therefore the skew makes the price of digitals more expensive.

- A long position in a digital call is long the skew.
- A short position in a digital call is short the skew.

Since digital options are sensitive to skew, you must use a model that knows about skew. When pricing European digitals, then your calibration should focus on getting the skew at maturity correct. When pricing American digitals with path-dependency, you will need to use some smooth surface calibration to capture the effect of surface through time. In other words, when dealing with these path-dependent american digitals, you are not only sensitive to the volatility at maturity but to many volatilities before maturity. Your volatility hedge will then consist of several European options with different maturities. We speak about vega buckets. The book of Adil Reghai is particularly good to grasp the concept of vega buckets and vega KT.

7.3.2 American Digitals

For American digitals, the trigger condition can be activated at anytime before maturity.

There exists a fantastic approximate link between European digital options and American digital options.


I felt quite stupid while learning about it as it is actually quite intuitive :). Back in 2015, I used to try estimating the price of every exotic option before pricing them. I was trying to develop as much as possible my intuition in terms of pricing and sensitivities in every market scenario. I quickly realised that the price of American digitals were always approximately twice the price of European digitals (with the same characteristics obviously!). The price of European digitals being quite easy to estimate, the approximations for American digitals were not too bad. One day, I decided to stay a bit later on the floor and start plotting Monte Carlo simulations to compare the price sensitivity of a barrier option to the barrier level with respect to the underlying's forward. Doing so, I realised why the above relationship between American and European digitals were so consistent. This is simply the consequence of a well-known principle followed by brownian motion: the reflection principle.

7.3.2.1 The Reflection Principle

In the theory of probability for stochastic processes, the reflection principle for a Wiener process states that if the path of a Wiener process $W(t)$ reaches a value $W(s) = a$ at time $t = s$, then the subsequent path after time s has the same distribution as the reflection of the subsequent path about the value a . In other words, if $W(s) = a$ then $W(t)$ is just as likely to be above the level a as to be below the level a for $s < t$.

By assuming in our models that the log-returns of the underlying are normally distributed with zero log-drift (with mean zero), the normal distribution introduces the symmetry of the reflection principle.

Gatheral expresses it nicely in his lecture on Barrier options. In Fig 6.5 below, the dashed path has the same probability of being realized as the original solid path. We deduce that the probability of hitting the barrier B is twice the probability of ending up below the barrier at expiration. Putting this another way, the value of an American digital option is twice the value of a European digital option. Note that this relationship won't be exactly respected when the log-drift is not zero (understand when the forward level is different from the spot level).



7.3.2.2 No-Touch Options

No-touch digital options pays a coupon if the barrier has never been touched during the option life.

It seems clear that the event of never touching the barrier is complementary to the event of ever touching it. Therefore, the probability of never touching the barrier is nothing else but 1 minus the probability of ever touching the barrier. From this parity, we can easily deduce the price of a no-touch digital option knowing the price of the american digital option and vice versa.

Chapter 8

Options Strategies

TBW :)

8.1 Call Spreads

8.2 Put Spreads

8.3 Straddles

8.4 Strangles

8.5 Risk Reversals

8.6 Butterfly Spreads

8.7 (Iron) Condor Spreads

Chapter 9

Asian Options

An Asian option is a derivative with a payoff at maturity that depends on an average of the underlying on a set of predetermined observation dates. Since the payoff of Asian options is based on the average of the underlying asset prices on a set of observation dates, the uncertainty concerning the fluctuations of the underlying price at maturity decreases. Therefore it decreases the risk exposure to spot price and volatility compared to vanilla options.

Also, the higher the number of observations, the lower the volatility and the lower the option price. Because of this lower volatility, asian options are generally cheaper than vanilla options (with same characteristics).

Asian options are commonly used for currencies, interest rates, commodities and energy markets. They are useful in corporate hedging situations, for instance, a company exchanging foreign currency for domestic currency at regular intervals.

9.1 Asian-In and Asian-Out

There are typically two types of Asian options: Asian-in and Asian-out options. Be careful when using these terms as they are not used the same way by the entire financial community!

9.1.1 First case

I personally speak about asian-in options when the averaging observation dates are spread uniformly during the entire life of the option. At the contrary, asian-out options have their averaging observation dates gather during a specific period near the maturity date.

In the general case, asian-in options are less risky than asian-out options. Intuitively, the uncertainty about future spot prices is lower when averaging periodically over the option's life than when averaging periodically over a shorter period near maturity.

9.1.2 Second case

Some people speak about asian-out options when an average is computed to determine the settlement price and asian-in options when the an average is computed to determine the strike price.

According to this definition: - Asian-out options pay the difference between the average of the underlying on the predetermined observation dates and the fixed strike price. - Asian-in options pay the difference between the underlying price at expiry and the average of the underlying at these predetermined observation dates (floating strike).

9.2 Geometric Average vs Arithmetic Average

There are also several ways to compute an average. More precisely, the average part of the asian options can be either geometric or arithmetic.

Geometric asian options are easy to price since there exists analytical formulae. This is simply due to the fact the geometric average of a lognormally distributed underlying has a lognormal distribution. Therefore, you will be able to use Black-Scholes model with specific volatility and interest rate levels as shown in Exotic Options Trading from Frans de Weert.

This author emphasizes something that appears to be quite confusing for most of the readers. While σ_i and r_i are always smaller than σ_{i+1} and r_{i+1} respectively (because any deviation in period i is automatically in period $i+1$), it does not prevent the annualised volatility of σ_i to be larger than the annualized volatility of σ_{i+1} ! It is therefore important to distinguish between annualized IV and the IV associated with the actual term of the option.

This confirms the fact the volatility of an average of observations is always smaller than the volatility of the share price itself.

Arithmetic asian options are far more common and don't have closed formulas to price them since the arithmetic average of log-normal variables is not log-normal.

The pricing of arithmetic asian options would then rely on Monte Carlo simulation with appropriate smile models such as local volatility. You can also use a moment-matching technique. This consists in calculating the exact first two moments of the Asian and then matching them, making the assumption that

they are the result of a log normal. Once you have identified this particular log normal, you end up using B&S formula.

If you want to use geometric average as an estimate for the arithmetic average, keep in mind that the geometric average is a lower bound of the arithmetic average as stated by the well-known Jensen inequality.

A more economic way to understand why the IV of an asian option is smaller than the IV of the vanilla option with the same maturity is by recognising that the duration of the asian option is smaller than the duration of the vanilla option with the same maturity.

9.3 Numerical Application

In the absence of interest rates and dividends, the price of an ATM Asian call can be approximated by the following formula: $C_T = \frac{1}{\sqrt{3}} 0.4 \sigma \sqrt{T} S_0$

This proxy formula shows that the volatility of the Asian Call is lower than the volatility of the European Call: $\sigma_{Asian} = \frac{1}{\sqrt{3}} \sigma$.

9.4 Risk Analysis : The Greeks

9.4.1 Delta Hedging

Let us take the practical example from Frans de Weert's book to apprehend the delta hedging of an asian option.

You, as a trader, just sold 90.000 3-month Asian calls on Total stock with a strike price of 40€ and monthly observations (3 observations). Let's assume that each call gives its owner the right to buy one stock at maturity.

9.4.1.1 Inception

At inception, the asian call's delta is quite similar the 3-month european call, so around 50%. You will therefore delta hedge by buying 45.000 shares of Total.

9.4.1.2 First observation date

One day before the first observation date (approximately one month after inception), Total is now trading at 48€.

In this case, the delta of the asian call is larger than the 'equivalent' european call. This is because the first Asian setting will almost certainly be ITM (> 40€).

Let say the asian delta is $5/6$ (long 75.000 shares of Total) and the equivalent european delta is $3/4$ (long 67.500 shares of Total).

Since the first Asian setting is assured to be ITM ($\text{delta} = 1$), which accounts for one third of the weight in the arithmetic average (3 monthly observations), 30.000 shares of the 75.000 shares serve as a delta hedge for the first Asian setting.

You will sell these 30.000 shares at the close of the first Asian setting and will be left with a long 45.000 shares of Total as your delta hedge.

Once the first Asian setting is taken, the remaining asian call is effectively on 60.000 shares only.

9.4.1.3 Second observation date

After the first observation, Total's stock price goes down sharply. One day before the second observation date, Total is now trading at 32€. As a result, your delta hedge has decreased from a long 45.000 shares position to a long 10.000 shares position.

Since the second Asian setting is assured to be OTM ($\text{delta} = 0$), your delta hedge is purely against the third Asian Setting and nothing against the second one. For that reason you don't have to do anything on the close of the second Asian setting. You can therefore deduce that the delta for the third observation is currently $1/3$ ($10.000/30.000$).

9.4.1.4 Third observation date

At maturity, Total stock is trading at 42€ and is therefore ITM ($\text{delta} = 1$). This means that you hold a long position of 30.000 shares as delta hedge. Obviously, you have to sell these shares at the close of the last Asian setting.

9.4.1.5 To sum up

On the day of an Asian setting, the trader needs to unwind part of his delta hedge if the Asian setting is ITM and does not need to do anything if the Asian setting is OTM.

If the Asian setting is ITM, the trader needs to unwind as a share position the number of Asian options multiplied by the weight of the setting.

This practical example also shows that the duration of an Asian option is indeed shorter than the equivalent European option. Effectively, an Asian option is spread out over a set of European options with maturities equal to the Asian observation dates. Therefore the term of an Asian option can be compared to a European option with a shorter term (estimated as the weighted average of

the different Asian observation dates). That's normally the moment you should try comparing a 1-year Asian call with quarterly observations and a 6-month European call with the help of our exotic options pricer.

Since the higher the time to maturity, the higher the option price. It is clear (if it was not yet!) that the price of an Asian option should be smaller than the price of its equivalent European option.

9.4.2 Vega, Gama and Theta

We have just seen in the previous section that the duration of an Asian option is shorter than the equivalent European option. It is then easy to see that the greeks change accordingly.

It basically means that, compared to an equivalent European option, an Asian option has higher gamma and theta but a smaller vega. If you are not sure why, I invite you to (re)read about the impact of time to maturity on vega, gamma and theta in chapter 5 - The Greeks..

Chapter 10

Quanto Options

This chapter has been written based on Frans de Weert's book - Exotic Option Trading (2008).

You can price and analyze the underlying risks of quanto options using our quanto options pricer. We used it to retrieve most of the graphs from this chapter.

Quanto options are also called quantity adjusting options.

10.1 Description

A quanto option is an option denominated in a currency other than the currency in which the underlying asset is traded. A quanto product converts underlying asset prices into units of the payoff currency by applying a fixed exchange rate. It is suited for investors who want to enter into an option strategy on a foreign underlying asset but are only interested in the % return of that strategy and want this return to be paid on their own currency with no FX rate exposure. The FX rate will be fixed to the rate prevailing at the inception of the deal and the payout of the quanto strategy will simply be this FX rate times the payout of the regular strategy. Cashflows are computed from the underlying in one currency but the payoff is made in another.

The payoff of a quanto call option is given by: $\boxed{Quanto\ C_T = FX_0 * \max(S_T - K, 0)}$

Where the strike price, K , is expressed in the underlying's currency.

10.1.1 Example

A European investor is long a 1Y ATM call on Glencore (GBP) and wants to get his return in EUR. Assume that Glencore's stock price is 3£, FX rate

(investor/foreign) = 0.86 EUR/GBP and that in one year, Glencore's stock price is 3.3£.

At maturity, the payout of this quanto ATM call is 0.258€ ($0.3 * 0.86$). Since Glencore's stock has increased by 10%, the European investor also expects 10% return on his EUR investment.

As you can see, the payoff of quanto options is quite straightforward. Surprisingly, pricing them and understanding what market variables it depends on is a much harder task and will be the subject of the following section.

10.2 Additional sensitivities : correlation risk and FX volatility

In comparison to a vanilla option, a quanto option is sensitive to two new market variables:

- the correlation between the log of the underlying price and the log of the FX rate
- the FX volatility

10.2.1 Correlation between underlying's price and FX rate

A European investor long an ATM EUR quanto call is short or long the FX correlation? The answer obviously depends how we express the FX rate. Let us express it as $FX = \text{investor/foreign}$.

Let us assume the correlation is positive. In this case, if the GBP gets more valuable against the EUR, Glencore's price tends to increase. The investor would then have been better off with a vanilla call. He is therefore short this correlation.

Like many correlations, it is very hard to obtain an implied quanto correlation from market data. If you have a liquid market for quanto options, you can back out this quanto correlation as all other parameters are known. Most of the time, you will have to estimate it using the realized correlation and taking some margin.

10.2.2 FX Volatility

Note that whatever how you expressed the exchange rate, its volatility will be the same. In other words, the volatility of EUR/GBP and the GBP/EUR are the same thing.

The sensitivity to this market variable is more tricky and less intuitive to grasp.

Is the European investor long or short the FX volatility?

To shed some light on the sensitivity to the FX volatility, it is time to introduce a small and intuitive model describing the stock price difference in the quanto currency for a small time interval.

$$\frac{dF_t}{F_t} = (r_{local} - \rho\sigma_S\sigma_{FX})dt + \sigma_S dW_t$$

where

- F_t is the underlying stock quoted in the quanto currency
- ρ is the correlation between $\log(S)$ and $\log(FX)$
- r_{local} is the risk-free rate of the underlying stock's own currency. In this case, it is the GBP risk-free rate.

F_t is defined so that a vanilla option on F_t is actually a quanto option on S_t .

Many things can be learned from this equation.

First of all, it confirms our intuition about the correlation sensitivity as the holder of a quanto call (put) is always short (long) the correlation.

Secondly, it makes clear that the FX volatility sensitivity depends on the sign of the correlation. If the correlation is negative, then the holder of a quanto call (put) is long (short) FX volatility. If the correlation is positive, then the holder of a quanto call (put) is short (long) FX volatility.

Thirdly, it highlights the intuitive fact that the volatility used to price the quanto option should be the same as the IV of the underlying stock. This is obvious as the FX rate is fixed and therefore the option payout only depends on the actual stock variation.

Then, the drift part is slightly different as the delta hedge is affected by the FX rate movements. While this will be further discussed in the next section, we can already give a small intuition. If the underlying stock price doubles, it impacts positively your delta hedge financing as you will sell more stocks at a higher level and therefore receive more interest. However, if the correlation is such that a doubling in the underlying stock price results in a halving of the EUR value against the GBP, then your delta hedge financing is actually unaffected. With this in mind, it does not seem surprising to adjust the financing part of the model by $-\rho\sigma_S\sigma_{FX}$. This term is called the '*quanto adjustment*' and can be added to the typical dividend yield term. Since it has a different sign, you can always think about the quanto adjustment as having an opposite effect as dividends.

Finally, notice that the volatility of the underlying stock also appears in the quanto adjustment. It can have an opposite effect as the usual volatility effect. Generally speaking, the quanto effect will be secondary.

10.3 Hedging FX Exposure

While hedging the FX exposure is not particularly intuitive, it is actually quite simple in practice.

As we have just seen in the above example, the FX hedge is captured by the delta hedge. The notional of a quanto option is agreed in the quanto currency. Therefore the notional in the local currency (underlying's currency) changes whenever the FX changes. For example, if the quanto currency doubles with respect to the local currency, the notional of the quanto option in the local currency also doubles and therefore the trader needs to double his delta hedge even though the stock price has not moved.

This example shows that there is no need for a trader to put an FX hedge in place for a quanto option. Well, if he sold a quanto option, he would need to swap the FX on the premium received. This is in line with the small and intuitive model of the previous section that prescribes financing in the local currency.

The majority of the exotic desks are long delta, therefore long dividends and short quanto correlations.

Chapter 11

Compo Options

This chapter has been written based on Frans de Weert's book - Exotic Option Trading (2008).

Compo options are also called composite options.

11.1 Description

As a quanto option, a compo option is also an option denominated in a currency other than the currency in which the underlying asset is traded. Unlike quanto option, the holder of a compo option has exposure to the FX rate. In a compo option, the payout and the strike are fixed in the compo's currency. Trading a compo option allows the investors to protect the value in their own currency on a foreign investment.

11.1.1 Example

A European investor holds some Glencore stocks. Assume that Glencore's stock price is 3£ and the FX rate (investor/foreign) is 0.86 EUR/GBP. This means the EUR value of one Glencore share is 2.58€.

To protect his holding, he decides to buy a 1Y ATM compo put option on GLEN.

Assume that after one year, GLEN decreased to 2.7£ and the EURGBP decreased to 0.8. This means that the EUR value of one Glencore share is 2.16€ (0.8×2.7).

However, this loss is offset by the payout of the compo put option because its strike is fixed in EUR (compo put payout = 2.58€ - 2.16€).

As you can see, compo options can be used to protect the underlying value in the investor currency to both FX rate and stock price movements.

11.2 Additional sensitivities : correlation risk and FX volatility

In comparison to a vanilla option, a compo option is sensitive to two new market variables:

- the correlation between the log of the underlying price and the log of the FX rate
- the FX volatility

As we did with quanto options, let us model the stock price difference in the compo currency for a small time interval:

$$\frac{dF_t}{F_t} = r_{compo} dt + \sigma_{compo} dW_t$$

where

$$\sigma_{compo}^2 = \sigma_S^2 + 2\rho\sigma_S\sigma_{FX} + \sigma_{FX}^2$$

and r_{compo} is the RFR of the compo currency.

The dynamics of this formula is quite different from the quanto scenario. This time, the main divergence with vanilla option lies in the difference of implied volatility.

This formula makes it clear that the holder of a compo option is long correlation and long FX volatility.

11.3 Hedging FX Exposure

Whenever a trader sells a compo put option on GLEN, he will want to delta hedge himself. Since the drift part of the above formula is the same as the drift part of a normal option on a stock in the compo currency, the delta can easily be determined by setting strikes and stock price in the compo currency. However, the trader can only execute his delta on the underlying stock quoted in the foreign currency. Even if the trader delta hedges himself, he will still have an FX risk as the compo option payout and the delta hedge will not be in the same currency.

Let us illustrate this using the above example and assume that the trader only delta hedges his short position on a 1Y compo put on GLEN. Let us assume for

the argument's sake he did so on a delta of 1. That means he made 0.30£ (3£ - 2.7£) on his delta hedge per compo put, which is worth 0.24€ at maturity as we assumed the EURGBP went down to 0.8 by that time. As the EUR value of GLEN went from 2.58€ to 2.16€, he lost 0.42€ on the compo option but only made 0.12€ on his delta hedge. He ends up with a 0.18€ loss.

Only delta hedging is clearly not enough. What can the trader do to be fully hedged then?

The trader would need to buy EUR on the notional of his delta hedge to be fully hedged on his compo put position. Let us make sure that it is indeed the case.

To delta hedge a short compo put position, the trader needs to sell shares. For every share that he sold, he received 3£. He would need to sell this 3£ to get 2.58€. At maturity, the trader can buy back these 3£ for 2.4€. The total profit on this FX hedge is 0.18€ (2.58€ - 2.4€). Adding this profit to the profit of his delta hedge offsets the loss on the compo put position.

Basically if the trader sells stock as a delta hedge, he needs to sell the stock's currency and buy the compo currency in the same notional as his delta hedge. If he buys stock as a delta hedge, he needs to buy the stock's currency and sell the compo currency in the same notional as his delta hedge.

This FX hedge is dynamic and will need to be adjusted all along with his delta adjustments. In other words, he will need to have, at any time, the same notional of FX hedge as delta hedge.

We forgot about something though! he will also need a FX hedge on the compo option premium paid at inception for a long position. The reason is clear: to buy a compo option, he will first need to sell the local currency to be able to buy the compo currency. That gives him an FX position that needs to be hedged. The hedge will then be to buy the local currency and sell the compo currency.

For a short position, such an FX hedge on the premium will not have to be put in place as the model used above assumes financing in the compo currency (EUR in our example).

Chapter 12

Barrier Options

This chapter has been written using several books, namely: Frans de Weert's book - Exotic Option Trading (2008), Bouzoubaa and Osseiran's book - Exotic Options and Hybrids (2010), Encyclopedia of Quantitative Finance (2010).

You can price and analyze the underlying risks of barrier options using our barrier options pricer. We used it to retrieve most of the graphs from this chapter.

12.1 Description

Barrier options are path-dependent options, that is, their payoff is not only a function of stock level relative to option strike but also dependent upon whether or not the stock reaches certain prespecified barrier level at or before maturity.

Barrier options are very popular amongst retail investors as the barrier feature provides the investor with additional protection or leverage.

The two most common kinds of barrier options are *Knock-Out* and *Knock-In* options.

Knock-Out (KO) options are options that expire worthless when the underlying's spot crosses the prespecified barrier level.

Knock-In (KI) options are options that only come into existence if the prespecified barrier level is crossed by the underlying asset's price.

The barrier observation can be at any time during the option's life (American style) or at maturity only (European style).

Depending upon the barrier level relative to the initial underlying asset level, we can have an 'up' or a 'down' barrier.

Together, we can therefore have four types of barrier options:

- Up-and-In options (UI)
- Up-and-Out options (UO)
- Down-and-In options (DI)
- Down-and-Out options (DO)

12.2 Knock-Out Options

Knock-Out (KO) options are options that expire worthless when the underlying's spot crosses the prespecified barrier level.

In the case of KO options, an additional feature called a **rebate** can be added to the contract specifications. The rebate is a coupon paid to the holder of a KO option in case the barrier is breached.

The leverage effect of an KO option can be much more attractive than the leverage of a comparable vanilla option for an investor who believes the spot will not reach the outstrike during the investment period. He gets more profit for bearing the risk of knocking out.

12.2.1 Some sensitivities

12.2.1.1 Barrier Level

The closer the barrier level is to the initial spot, the cheaper the KO option would be. This is quite intuitive: the closer the barrier level, the higher the probability of the option expiring worthless.

12.2.1.2 Barrier Observation

For a KO option, the higher the number of barrier observations, the higher the probability of the option knocking out and the cheaper the KO option. A KO option having an annually monitored barrier would be more expensive than a similar KO option having a quarterly monitored barrier.

12.2.1.3 Volatility

A KO option is less sensitive to volatility than a vanilla option carrying the same features. Indeed, a higher volatility increases the probability of expiring ITM but also increases the probability of reaching the barrier and ending with no value. The Vega of a KO option is generally lower than the Vega of a comparable vanilla option.

12.2.1.4 Gap risk

When KOs are defined with the barrier placed in such a way that the option vanishes when it is OTM, we call these **regular** KO options. In these, it is easier for traders to hedge the associated risks.

Otherwise, KO options are classified as **reverse** and they present higher trading difficulty and risks. For example, they are subject to a *discontinuity risk* or a *gap risk*. We will discuss how does a trader handle this gap risk in more details in this chapter.

12.3 Knock-In Options

Knock-In (KI) options are options that only come into existence if the prespecified barrier level is crossed by the underlying asset's price.

The leverage effect of a KI option can be much more attractive than the leverage of a comparable vanilla option for an investor who believes the spot will reach the outstrike during the investment period. He gets more profit for bearing the risk of not knocking in.

12.3.1 Some sensitivities

12.3.1.1 Barrier Level

The nearer the barrier level to the initial spot, the more expensive the KI option would be. This is quite intuitive: the closer the barrier level, the higher the probability of the option coming into existence.

12.3.1.2 Barrier Observation

For KI options, the higher the number of barrier observations, the higher the probability of the option being activated, the more expensive the KI option. A KI option having a quarterly monitored barrier would be cheaper than a similar KI option having a monthly monitored barrier.

12.3.1.3 Volatility

Unlike a KO option, a KI option is more sensitive to volatility than a vanilla option carrying the same features. Indeed, a higher volatility can benefit the holder of the option because it increases not only the probability of maturing ITM but also the probability of reaching the barrier and being activated. The Vega of a KI option is then higher than the Vega of a comparable vanilla option.

12.3.1.4 Gap risk

When KI options are defined with the barrier placed in such a way that the options are activated when it is OTM, then we call them **regular** KI options since it is easier for traders to hedge the associated risks.

Otherwise, KI options are classified as **reverse** and they present greater trading difficulties and risks. For example, they are subject to a *discontinuity risk* or a *gap risk*. We will discuss how does a trader handle this gap risk in more details in this chapter.

12.4 In-Out Parity for barrier options

Being long a KO option and a KI option with the same features is equivalent to owning a comparable vanilla option independently from the behaviour of the spot with respect to the barrier level.

$$\boxed{\text{Knock-In (K,T,B)} + \text{Knock-Out (K,T,B)} = \text{Vanilla (K,T)}}$$

It is very easy to see that, for any given scenario of the underlying asset path before maturity, the portfolio (KI + KO) will always have the same payoff as the corresponding vanilla option. This relationship holds for both the put and call options in the absence of rebates.

12.5 Review of Payoffs

$$\text{UI Call}_T = \max[S_T - K, 0] * \mathbb{1}_{\{(\max_{t \in [0, T]} S(t)) \geq B\}}$$

$$\text{UI Put}_T = \max[K - S_T, 0] * \mathbb{1}_{\{(\max_{t \in [0, T]} S(t)) \geq B\}}$$

$$\text{UO Call}_T = \max[S_T - K, 0] * \mathbb{1}_{\{(\max_{t \in [0, T]} S(t)) < B\}}$$

$$\text{UO Put}_T = \max[K - S_T, 0] * \mathbb{1}_{\{(\max_{t \in [0, T]} S(t)) < B\}}$$

$$\text{DI Call}_T = \max[S_T - K, 0] * \mathbb{1}_{\{(\min_{t \in [0, T]} S(t)) \leq B\}}$$

$$\text{DI Put}_T = \max[K - S_T, 0] * \mathbb{1}_{\{(\min_{t \in [0, T]} S(t)) \leq B\}}$$

$$\text{DO Call}_T = \max[S_T - K, 0] * \mathbb{1}_{\{(\min_{t \in [0, T]} S(t)) > B\}}$$

$$\text{DO Put}_T = \max[K - S_T, 0] * \mathbb{1}_{\{(\min_{t \in [0, T]} S(t)) > B\}}$$

12.6 Deeper into Risk Analysis

We went into little detail about the risk analysis of barrier options. Let us go deeper into this subject as, from a risk management perspective, barrier options

are very interesting because the risks are discontinuous around the barrier. Therefore the Greeks become less predictable and often change sign around the barrier.

It is essential to understand the risks embedded in barrier options as those risks will have to be taken into account in the price. Again, the price of an option should reflect the cost of hedging it!

We will not go through the 8 types of barrier options but will use the DI put as a leading example to get across all the risks within a barrier option.

The reason we use a single leading example is because the underlying causes for barrier option risk are generic and once the drivers of barrier option risk are understood for a DI put, one will be able to derive the risks for the other types of barrier options.

The reason we choose the DI put is because many structured products use it to obtain enhanced yields or increased participation. The investor accepts the risk from selling the DI puts to generate extra funding that is used in the structure to increase the yield or participation.

12.6.1 Leading Example : Down-and-In Put (DIP)




Figure 12.1: Fig: 12.1 : Payoff of a Down-and-In Put

12.6.1.1 Delta Change over the Barrier

Traders on the sell side are usually long the DIP and have to hedge the risks associated with this position accordingly.

The trader taking a long position in the DIP is short the forward and will need to buy delta in the underlying stock at inception and adjust dynamically his delta hedge to remain delta neutral.

Assume the trader buys a 100/80% DIP on Total from an investor. The payoff of this position is represented in the above graphic. Assume the delta at inception

is 0.4 so that the trader hedges himself by buying Total shares in a ratio of 0.4 shares per option.

When Total stock price breaches the 80% barrier level, the DIP goes instantaneously from not being an option to being a 20% ITM put. As a result, the value of the DIP increases significantly over the barrier, even if the stock price barely moves down from 80.1% to 79.9% of its initial value. This means that the absolute delta of DIP becomes extremely large when the underlying price gets closer to the barrier. This absolute delta often becomes greater than 1.




Figure 12.2: Fig: 12.2 : Delta of a Down-and-In Put close to expiry

Since the absolute delta of a vanilla option can never be greater than 1, the trader will accumulate too many Total shares when the stock approaches the barrier. As soon as the stock price breaches the barrier, the trader would need to sell any excess shares. In practice, it is very likely that he will not be able to sell these excess shares exactly at the barrier level. Indeed, the share price is already going down for the barrier to be breached at the first place, and the fact that the trader needs to sell a potentially large quantity of shares will push the price further down. So the trader will probably sell these excess shares below the barrier level and incurs a loss on this sale.

Clearly, the risks of a barrier option near the barrier level can be difficult to manage. The delta of a barrier option can jump near the barrier causing hedging problems. So near the barrier, the Gamma (= the sensitivity of Delta to a movement in the underlying stock price) can be very large. To make things worst, it actually changes sign around the barrier. A long DIP position goes from being long gamma to short gamma as the share price approaches the barrier.

One method to smooth out the risks to make them manageable is to apply a barrier shift. To avoid the loss of selling the excess shares below the barrier level, the trade will give himself a cushion. To do that, he will price and risk manage a slightly different option where the barrier is shifted downwards in such a way that the trader has enough room to sell the excess shares without incurring a loss. If he believes that he needs a cushion of 2% to sell these shares over the barrier, he will price and risk manage a 100/78% DIP rather than a 100/80% DIP. It makes the option cheaper for the trader.

12.6.1.2 Magnitude of the Barrier Shift

In practice, the size of the barrier shift (2% in the above example) is not a random choice and depends on several factors:

- **The size of the transaction.** The larger the size, the more shares will have to be sold over the barrier, therefore the trader is more likely to move the stock price against him. To sum up, the larger the transaction size, the larger the barrier shift.
- **The difference between strike price and barrier level.** If this difference is large, the DIP goes from not being a put to a put that is far ITM when the underlying price breaches the barrier. To sum up, the larger the difference between strike price and barrier level, the larger the barrier shift.
- **The daily volume of the underlying share.** If the daily traded volume is low, it might be difficult for the trader to sell the excess shares over the barrier and he will likely sell them at a bad price, incurring a higher loss. To sum up, the lower the daily traded volume of the underlying asset, the larger the barrier shift.

These three first factors form a liquidity-based barrier shift and account for the discontinuity in the Delta near the barrier.

- **The volatility of the underlying stock.** The more volatile the underlying stock, the larger is the risk to the trader of the stock price approaching the barrier level. In other words, the larger the volatility, the larger the barrier shift needed to protect the trader against a larger move. This can be easily seen through our previous example. You are the trader long 100.000 100/80% DIP on Total. Total is currently trading at 81% of its initial level. Suppose you priced this DIP as a 100/78% and you are risk managing it as such. Assume the absolute delta for this 100/78% DIP is currently 2 so that you are long 200.000 of Total stocks. If Total gaps down 10%, the DIP knocks in and its absolute delta is 1 (far ITM put). It means that you need to be long a 100.000 of Total shares. Well, you were long 200.000 of these shares so you need to sell for 100.000 of them over the barrier. As the stock gapped down, you can only sell them at 71% of the initial stock price! Since you managed this DIP as a 100/78% DIP, your loss is slightly reduced but still very painful!

This kind of gap down can bring situations where the trader would actually need to buy stocks when the barrier is breached. Using the DIP as above, assume the stock is trading far away from the barrier, at 87% of its initial level for example. At this point, the absolute delta is not particularly huge and the trader has not

accumulate any excess delta. If the stock price goes through the barrier in a gap move down, the absolute delta will increase and the trader will need to buy shares (and will be able to do so at a lower price).

- **The barrier level.** Since lower stock prices tend to go hand in hand with higher volatilities and higher volatilities result in larger barrier shift, the lower the barrier level, the higher the barrier shift.
- **The time left to maturity.** The closer to maturity, the larger the absolute delta just before the barrier and therefore the larger the change in delta over the barrier. This translates into a higher risk and a higher barrier shift for shorter maturities.

Note that certain types of barrier options do not require a barrier shift. A long position in a UO put is a good example. The delta hedge of such position is always a long share position so that the trader will need to sell shares over the barrier. The trader does not need to shift the barrier as he can always let a limit order just before the barrier. If the stock gaps up, he will be pretty happy to sell them at a higher price anyway.

12.6.1.3 Types of Barrier Shift

When traders quote barrier options, they can be more or less conservative on their prices depending on the magnitude of the barrier shift but also on the way this barrier shift is applied.


Until now, we have only considered constant barrier shift but it could also be an increasing function of time. It is quite intuitive if you think about it. In the first days of an option's life, under normal levels of volatility, it is quite unlikely to see the underlying breach the barrier level. If one were to simulate paths and monitor the points in time at which the barrier was breached, it is quite obvious that the KI events occur more frequently down the line. That is the reason why traders often apply a barrier shift that is an increasing function of time.

Assume an investor is willing to sell a 100/70 KI call. 3 IBs are competing for this trade and will all take the same commissions and apply the same pricing parameters (volatility, skew, etc...). They all want to apply a shift of 2% but they have different ways of shifting the barrier.

Trader 1 is very conservative and applies a **constant barrier shift** of 2%.

Trader 2 is less conservative and decides to apply a **linear barrier shift**. At inception date, there is no barrier shift since there is no expected risk around the barrier. He believes that the maximum shift to be applied would be 2%, which is the shift value at maturity. The shift grows linearly from zero at inception to 2% at maturity.

Trader number 3 uses a **curvy barrier shift** in time which is computed from evaluating knock-in scenarios. he is the most aggressive in his barrier shift. Therefore, his bid is the highest and he wins the trade in this case.



12.6.1.4 Barrier at maturity only

While closed-form solutions exist for continuously monitored barriers, other types of barrier observation require a Monte Carlo process. This is the case when barrier options are only live at maturity or on specific days. However, barrier options that are only live at maturity can also be priced as a combination of European options.

Let us keep our example of a trader long a 1Y 100/80 DIP on Total where the option can only knock in at maturity. It is obviously worth less than the DIP that can knock in anytime before maturity and can be replicated as follows:

- The trader buys 10 times a 1Y 80% European put
- The trader sells 10 times a 1Y 78% European put
- The trader buys one 1Y 80% European put

The fact that the trader buys a 10x leverage put spread 80/78 is an overhedge and therefore a conservative way to replicate the DIP's payoff at maturity. The 2% wide put spread can be seen as a barrier shift. The tighter the put spread the more the replication converges to the actual price of the DIP with KI at maturity only. The gearing of the put spread can be calculated by dividing the size of the discontinuity (the strike/barrier differential) by the width of the put spread.

This discussion should make you think about what has been said in the chapter about digital options that can be seen as a limit of a call/put spread! This is exactly the same here as the discontinuity around the barrier corresponds somehow to a digital option.

Now that you have seen how you can replicate a barrier option with barrier observation at maturity only, you should easily see why these barrier options are sensitive to skew!

12.6.1.5 Volatility and Skew

We have already seen that a knock-in option is more sensitive to volatility (higher vega) than a vanilla option carrying the same features. Indeed, a higher volatility can benefit the holder of the option because it increases not only the probability of maturing ITM but also the probability of reaching the barrier and being activated.

The trader buying a DIP is therefore long volatility. This long vega position can be hedged, at least partially, by buying vanilla put options on the same underlying stock with strikes between the barrier and the spot.

Risk wise one can compare a KI barrier to a long option position at that barrier and a KO option to a short option position at the specific barrier. It is clear then that the owner of DIP is long the skew as his position is similar to being long a downside option (option with lower strike). In the presence of skew, the volatility around the barrier is higher than the ATM volatility, which makes the probability of crossing the barrier higher.

From a model point of view, we will need to calibrate a model to the IVs of options on the underlying asset across strikes with specific attention to the downside skew.

If the barrier is monitored continuously, we will need a model that gives a smooth calibration through all ends of the surface between short maturities and up to the option's maturity. It means we will need to calibrate to both skew and term structure. The reason is that a continuously monitored barrier option can be triggered at any time up to maturity, therefore it has vega sensitivity through the different time buckets. So European options with different maturities must be calibrated so that the model shows risk against them. You must understand that the vega sensitivity will change as the underlying moves if the underlying stock gets closer to the barrier level, then the short term vega will increase and the long term vega will decrease.

If the barrier is only monitored at maturity, then getting the skew corresponding to that maturity correct is the primary concern and we would use the exact date-fitting model.

12.6.2 Counter-example : Call Up-and-Out (CUO)

Let us highlight some differences in the risk analysis of a Up-and-Out call with respect to a Down-and-In put.

12.6.2.1 Delta Hedge

Assume a trader is short a up-and-out call. The delta at inception actually depends on how far is the barrier relatively to the spot. This clearly depends

on the volatility of the underlying asset. For example, for a 1y CUO, 20% can be seen as far from the strike if the volatility level is low and quite close if the volatility is high.

Initially, the trader usually needs to buy shares as a delta hedge. However, when the stock price gets closer to the upper barrier level, there is an inflection point where the trader will need to go short shares to be delta hedged! So delta can actually change sign before the barrier level for an up-and-out call.




Figure 12.3: Fig: 12.4 : Example of Delta of a 1Y Up-and-Out Call at inception

Once the call knocks out, the trader will need to buy back these shares as it is no longer a hedge against anything since the option does not exist anymore.

12.6.2.2 Volatility

A long position in an up-and-out call is not necessarily long vega. In fact, more often than not it is short vega, meaning that when volatility goes up, the CUO becomes less valuable. The higher IV results in a higher probability of the call knocking out and a lower chance of CUO having a payout at maturity.

The vega position also depends on how far is the barrier relatively to the spot. For example, if the barrier is very upside, then the probability of knock-out is very low and an increase in IV has a higher impact on the option part than on the fact that it increases the probability of the option knocking out.

To sum up, an up-and-out call will be long vega for very upside barriers and short vega for lower barriers. How far is the barrier is relative and depends on the barrier/strike differential and the level of volatility.

12.7 Double Barrier Options

A double barrier option is another variation that has two barriers, typically one up barrier and the other down barrier. There are typically used within structured products to achieve a specific type of payoff. These options are priced using a

Monte Carlo process. Therefore, from a trading perspective, the magnitude and direction of the barrier shift will be of great interest.

The first type of double barrier option features dependency between the trigger of the first and second barrier. In other words, the second barrier can only trigger if the first barrier has been triggered. The Monte Carlo modeling is slightly more complex as there is an additional condition but it is fact more transparent from a risk perspective than if the barriers were completely independent. Assume a trader sells an ATM call with a knock-in barrier at 90% and a knock-out barrier at 120%. Assume the call can only knock out after it has knocked in. Since the 120% barrier is dependent on the 90% barrier, one should first focus on determining the shift on the 90% one and then the shift on the 120%. The direction and magnitude of these barrier shifts have already been discussed by now.

The second type of double barrier option shows complete independence between the two barriers. When a specific barrier is breached, it triggers regardless of whether the other barrier has been triggered. The Monte Carlo modeling is easier but, from a risk perspective, it is actually less transparent. Let us take the same example of an ATM call with a 90% KI barrier and 120% KO barrier to show why it is less transparent. This option is worth less than the previous conditional one as it can already knock before it has even become a call option. This also means that the shift applied on the 120% barrier can be less if the barrier breaches before the 90% barrier has been breached. However, if the 90% has been breached first then you would require the same barrier shift as a regular CUO. As you can see, there is clearly more ambiguity regarding the magnitude of the barrier shift to apply.

You can price and analyze the risks of the second type of double barrier options using our double barrier options pricer.

Chapter 13

(Barrier) Reverse Convertibles

13.1 Reverse Convertibles

Reverse Convertibles are among the most popular yield enhancement products in Switzerland and are suited for investors who are anticipating a sideways or slightly upward trending market.


The holder of a reverse convertible gives up the potential upside exposure to the underlying asset in exchange for an enhanced coupon. The holder of the product remains exposed to the downside exposure.

The enhanced coupon of the reverse convertible is paid in any case. Because of this, the product will always outperform its underlying asset to the downside. The product will also outperform if the asset does not rise by more than the coupon. Hence, the ideal market scenario for reverse convertibles is the prospect of a sideways trending market.

13.1.1 Payoff

In its most basic form, the reverse convertible is constructed by means of a short put and a money-market placement. In most cases, the strike is placed ATM. Obviously, the lower the strike, the lower the value of the put option and the lower the guaranteed coupon.

The below graph shows the payoff of a 1y reverse convertible with a ATM Put and a guaranteed coupon of 10.5%.



The price of this reverse convertible is 99.81%, which is composed of:

- Long 1y Zero-Coupon Bond: 98.02% (assuming 2% interest rate for the sake of the example)
- Short 1y ATM put option: 8.50% (assuming 20% implied volatility and 3.5% of dividends)
- Guaranteed coupon: 10.5% paid at maturity $\rightarrow PV(10.5\%) = 10.29\%$

Let us assume that the 19 bps left (100% - 99.81%) are earned by the bank as a commission.

13.1.2 Risk Analysis

As you can see, there is no discontinuity in a reverse convertible because the coupon is not conditional. The risk analysis is therefore similar as a short put and very straightforward. The relative size of the coupon with respect to the risk-free rate of the product maturity will give you an idea of the risk of the product. The higher this difference, the more valuable the put, the riskier the product!

Since an investor buying this reverse convertible is selling a put option, he is selling volatility and should therefore invest when volatility is high or expected to fall.

The holder of a Reverse Convertible will incur a loss if the loss of the sold put option is larger than the sum of the guaranteed coupons received.

13.2 Barrier Reverse Convertibles

Barrier Reverse Convertibles (BRC) are a special variant of the classical Reverse Convertibles.

The holder of a barrier reverse convertible gives up the potential upside exposure to the underlying asset in exchange for an enhanced coupon. The holder of the

product is not exposed to the downside exposure, unless the underlying asset breaks through a predefined barrier set at the inception of the product.


The enhanced coupon of the barrier reverse convertible is paid in any case. Because of this, the product will always outperform its underlying asset to the downside. The product will also outperform if the asset does not rise by more than the coupon. Hence, the ideal market scenario for reverse convertibles is the prospect of a sideways trending market. All else remaining equal, a BRC pays a lower coupon than a reverse convertible, because of the conditional capital protection provided by the barrier.

13.2.1 Payoff

The barrier reverse convertible is constructed by means of a short Down-and-In put and a money-market placement. In most cases, the strike is placed ATM but more conservative versions placed it OTM. The barrier is usually ranging between 50% to 80% of the spot price. Obviously, the lower the barrier and the strike, the lower the value of the put option and the lower the guaranteed coupon.

As long as the underlying asset does not cross the barrier, the BRC remains capital protected. Once a barrier breach occurs, the capital protection is lost, and the BRC transforms itself in a classic reverse convertible.

The below graph shows the payoff of a 1y reverse convertible with a ATM 80% Down-and-In Put and a guaranteed coupon of 8.4%.



The price of this barrier reverse convertible is 99.76%, which is composed of:

- Long 1y Zero-Coupon Bond: 98.02% (assuming 2% interest rate for the sake of the example)
- Short 1y ATM 80% DIP: 6.50% (assuming 20% implied volatility and 3.5% of dividends)
- Guaranteed coupon: 8.4% paid at maturity $\rightarrow PV(8.4\%) = 8.23\%$

Let us assume that the 24 bps left (100% - 99.76%) are earned by the bank as a commission.

13.2.2 Risk Analysis

Since the coupon is guaranteed, it does not bring additional risk. The risk analysis is therefore similar as a short Down-and-In put and is quite straightforward if you read chapter 12.

Chapter 14

Certificates

14.1 General Description

Certificates are among the most popular participation products. There are characterized by a wide variety of risk-return profiles. This diversity of profiles can be achieved because the investor usually foregoes, with the purchase of a certificate, the dividends paid by the underlying asset during the life of the certificate. For that reason, certificates on shares that pay a high dividend often have a particularly attractive risk/reward profile.

Some advantages about investing in certificates: * Opportunity to implement sophisticated investment strategies through the purchase of a single product. * They are listed in regulated markets. * They are very liquid instruments. * Favorable taxation.

In this chapter, we will only describe the most well-known variants of certificates.

14.2 Tracker Certificates

The Tracker Certificate is a structured product that replicates (tracks) the performance of an underlying asset or often a basket of securities simply, inexpensively and with high liquidity. As such, there is no optionality involved in this product category. It is the reason one this simple product is not directly available in our exotic derivatives prices. However, you can easily price it with the certificates pricer by selecting 'Discount Certificate' and setting the Cap to 999. You can also use the Classic option pricer and price a zero-strike call option.

Excluding the issuer risk, its risk is identical to the underlying asset it tracks.

Why do they exist then? Why not simply investing in the underlying asset?

The purpose of trackers is their ability to allow investments otherwise not possible or not economically feasible.

Trackers are generally used to diversify the investor's risk exposure across a wide range of individual stocks. It provides a simple and cost-effective mean of investing in an entire stock market without having to buy each of the individual shares. With trackers, you can easily get access to exotic markets, for which you would not have access otherwise via your bank/broker.

What is the difference with a share index then?


Major share indices are priced in thousand of points. In other words, you need thousand of the denominated currency to buy a share index. That is clearly an obstacle for many investors. A tracker makes it possible to invest in the same index or any another baskets of stocks with smaller amounts of money. Exchange traded funds (ETFs) may also be a good alternative.

Tracker certificates are one of the few structured products that may have no maturity.

14.2.1 Payoff

Tracker certificates are constructed using a zero-strike call on the underlying asset. If the certificate tracks an index on which a liquid future is available, the issuer often buys the future instead of a zero-strike call. The issuer may even physically hold the underlying individual shares if the certificate is based on a basket of shares.

The below graph shows the payoff of a 1y tracker certificate on an underlying at a price of 100 with a dividend yield of 3.5% and interest rates hypothetically set to 2%.



The price of this tracker certificate is 96.56. This is simply the price of a zero-strike 1y call option but could also be split into:

- Long 1y Zero-Coupon Bond: 98.02 (assuming 2% interest rate for the sake of the example)

- Long 1y ATM Call : 7.04 (assuming 20% implied volatility and 3.5% of dividends)
- Short 1y ATM Put : 8.50 (assuming 20% implied volatility and 3.5% of dividends)

14.2.2 Risk Analysis : The Greeks

As you can see from the above payoff, this is a delta-one product. In other words, its delta is close to 1 while all the other greeks are null.

14.2.3 Bear Tracker Certificates

Bear tracker certificates represent the exact opposite of a normal tracker certificate. They gain in value when the price of the underlying instrument declines and vice versa. Thus bear tracker certificates are suitable for investors who anticipate falling prices.

14.3 Discount Certificates

Discount certificates are yield-enhancement products that allow the investors to buy an underlying instrument for less than its current market price (at a discount). As there is no free lunch, the investor is now limited to a predetermined amount (the cap). The lower the cap, the more limited the upside, the greater the discount.

Discount certificates are typically short term products with maturity ranging from one to three years.


At maturity, the investor will incur a loss only if the price of the underlying asset has decreased so far that the discount has been fully eroded. The good thing is that, if you incur a loss with a discount certificate, it will always be less than the amount you would have lost if you had invested directly in the underlying asset.

Where you place the cap really depends on your market view. If you expect a sideways market, you would place the cap closer to the current price of the underlying instrument than if you expect a modestly rising market.

14.3.1 Payoff

Discount certificates are constructed as a combination of a long zero-strike call and a short cap-strike call on the underlying asset.

The below graph show the payoff of a 2y discount certificate with a cap at 115 on an underlying at a price of 100 with a dividend yield of 3.5% and interest rates hypothetically set to 2%.



The price of this discount certificate is 88.42. This price can be split into :

- Long 2y 0 Call : 93.24
- Short 2y 115 Call : 4.81

but it could also be split into :

- Long 2y Zero-Coupon Bond: 96.08 (assuming 2% IR for the sake of the example)
- PV(guaranteed coupon of 15 in 2y) : 14.41 (assuming 2% IR for the sake of the example)
- Short 2y 115 Put : 22.07 (assuming a 20% IV and a 3.5% dividend yield)

14.3.2 Risk Analysis : The Greeks

The delta of the product is obviously positive. This positive delta position also implies that the investor is short dividends on the underlying stock. How positive is the delta depends on the position of the cap. The higher the cap, the closer to a tracker certificate is the product and the closer to 1 is the delta. In the above example, the delta is 63%. It corresponds to the delta of the combination of a long position in a 100/115 call spread (18%) and a short position in an ATM put (45%).

As the investor has more downside than upside, he will be short volatility. As the price of the underlying gets closer to the cap level, the investor becomes more sensitive to volatility. Think about the fact that the investor is short an option striking at the cap level and remember that the vega of an option is higher around its strike. On the contrary, as the price of the underlying asset decreases, the sensitivity to volatility also decreases (in absolute value).

Since the trader selling a discount certificate is buying an (OTM) option, not only he is buying volatility but he is also long gamma and short theta. He will get longer gamma as the underlying price gets closer to the cap level (the option strike) and the time to maturity gets smaller.

14.3.3 Deep-discount Certificates

Now that you understand a bit more the payoff of a discount certificate, let me introduce you to the extremely defensive version of it with a cap far below the current level of the underlying asset. This version is called the deep-discount certificates and can be seen as a substitute to a savings account since it provides much security.

Let us take a numerical example to show you the principle of the deep-discount strategy. Imagine you buy a discount certificate at 63.5 with a cap at 65 when the value of the underlying asset is at 100. Taking this into consideration, you will make a profit as long as the underlying does not drop below 63.5 (a loss of 36.5%). In addition, you will make a profit of 1.5 (2.36%) in any scenario in which the underlying does not lose more than 35 (a loss of 35%).

14.4 Bonus Certificates

Discount certificates can be a bit frustrating during bull markets since their performances lag those of the market. At the same time, investors keep on willing some safety buffer against price declines. Bonus certificates seem a good solution for the investors not willing to miss on the upside.

Bonus certificates are also typically short term products with maturity ranging from two to four years.

They guarantee a specified bonus level as long as the underlying asset has not breached an established barrier level during the term of the certificate.

In contrast to discount certificates, the payoff is not capped to a fixed amount so that the investor keeps on participating in the price gains above the bonus level.

However, if the safety barrier is breached, then the product changes into a tracker certificate and the right to a guaranteed payout of the bonus level no longer exists. If the underlying asset goes back up afterwards, the investor can still participate 1:1 in the potential gains but the guaranteed payout of the bonus level will not be reinstated after such move!


Hence, bonus certificates tend to perform well in both sideways and rising markets.

14.4.1 Payoff

Bonus certificates are constructed as a combination of a long zero-strike call and a long Down-and-Out put with striking at the bonus level on the underlying asset. The zero-strike call provides the participation, and the DOP option

generates the barrier and the bonus (determined by the DOP strike). In case of stocks, the dividends are used to buy the DOP option. Therefore, higher implied dividends generate a higher bonus level, a greater contingent protection level or a mix of both.

The below graph show the payoff of a 2y bonus certificate with a bonus level at 112.5 and a barrier level at 75 on an underlying at a price of 100 with a dividend yield of 3.5% and interest rates hypothetically set to 2%.



Note that this payoff is only valid at maturity; the mark-to-market price of the bonus certificate during its lifetime differs strongly from its final payoff. It may be stated that in general, the contingent protection and bonus will only “grip” after most of the product’s maturity has elapsed.

The price of this bonus certificate is 98.94. This price can be split into :

- Long 2y 0 Call : 93.24 (assuming a 3.5% dividend yield)
- Long 2y American DOP with barrier at 75 and strike at 112.5 : 5.70 (assuming a 20% IV and a 3.5% dividend yield)

14.4.2 Risk Analysis : The Greeks

The dynamics of the greeks are largely explained by the long DOP position. Since barrier options have already been discussed, you should be able to understand it and explain it by now. As always, the sensitivities depend on the product’s parameters.

Understanding the greeks and their dynamics is probably the most important and the most complicated aspect about derivatives. In this regard, I believe the exotic option pricer is an infinite source of information. The very large majority of interview questions related to equity derivatives asked by any financial institution can be answered using the exotic option pricer. To analyze the greeks of a bonus certificate, use the certificates pricer and select ‘Bonus Certificate’ as the type of certificate.

Remember, as soon as you see discontinuities in the payoff, think about the way the trader will smooth it on his side.


14.5 Capped Bonus Certificates

This is a particular version of the bonus certificate with limited participation in rising markets up to the cap level. As a bonus certificate, if the barrier has not been breached during the lifetime of the product, investor receives a minimum redemption equal to the bonus level. Should the barrier level be breached during the lifetime of the product, the capped bonus certificate turns into a tracker certificate with a cap.

14.5.1 Payoff

Bonus certificates are constructed as a combination of a long zero-strike call, a long Down-and-Out put with striking at the bonus level and a short cap-strike call on the underlying asset.

The below graph show the payoff of a 2y capped bonus certificate with a bonus level at 113, a barrier level at 70 and a cap level at 125 on an underlying at a price of 100 with a dividend yield of 3.5% and interest rates hypothetically set to 2%.



The price of this capped bonus certificate is 98.85. This price can be split into :

- Long 2y 0 Call : 93.24 (assuming a 3.5% dividend yield)
- Long 2y American DOP with barrier at 70 and strike at 113 : 8.24 (assuming a 21% IV and a 3.5% dividend yield)
- Short 2y 125 Call : 2.63 (assuming a 19% IV and a 3.5% dividend yield)

As you can see from the difference in implied volatilities between the down-and-out put and the OTM call, we assumed a bit of a volatility skew.

By limiting his upside to the cap level, the investor benefits from more protection and a slightly better bonus when comparing this capped bonus certificate with the bonus certificate previously discussed.

14.5.2 Risk Analysis : The Greeks

Since the product is a variant of the bonus certificate, its sensitivities to the various market parameters are very similar.

Its delta is also positive but smaller than the delta of the bonus certificate since the upside potential is capped. Selling an OTM call implies a negative delta on this position and therefore reduces the overall positive delta of the capped bonus certificate.

Its vega position is also negative and is larger in absolute value than the vega of the bonus certificate. An increase in volatility increases the probability to breach the downside barrier and lose the bonus. Since the upside potential is capped, the investor has more to lose than he has to win when volatility increases.

As for the bonus certificate, skew also has an impact on the pricing as both products include an down-and-out put option. The skew sensitivity is likely to be smaller for the capped bonus certificate as the skew will have a counterbalancing effect on the sold OTM call option (a stronger skew will tend to decrease the price of the long DOP position and decrease the price of the sold OTM call).

14.6 Airbag Certificates

Airbag certificates fully participate to the upside price performance of an underlying asset and have a certain amount of downside protection down to which the investor does not incur any losses. In other words, as long as the underlying asset does not lose more than a predefined level (the airbag level), the capital is 100% protected. Below the airbag level, the product starts to decrease in value at a leveraged pace compared to the underlying asset. The leverage is the inverse of the airbag level. There are many similarities with a bonus certificate with the one big difference that the Airbag does not knock-out and stays until maturity no matter what. For this reason, there is no discontinuity in the payoff as in the bonus certificate.


14.6.1 Payoff

Airbag certificates are constructed as a combination of a long leveraged zero-strike call, a short leveraged airbag-strike call and a long ATM call. The leverage is the inverse of the airbag level.

Let us take a specific example to make everything clear. The below graph shows the payoff of a 2y airbag certificate with an airbag at 85% of the current price of the underlying asset. This underlying asset is assumed to have a 3.5% dividend yield and interest rates are hypothetically set to 1%. In this example,

the leverage will then be of $1/0.85 \approx 1.1765$. It represents the speed at which the product's payoff will decrease below the airbag level.

As previously highlighted, the protection is bought the expected dividends of the underlying asset. The higher the dividend, the larger the protection.



The reason why you see 'Put Strike' in the legend of this graph is because you can also see the left side of the graph as a short position in a leveraged OTM put.

The price of this airbag certificate is 99.27. This price can be split into :

- Long leveraged 2y 0 Call : $93.24 * 1.1765 = 109.694$ (assuming a 3.5% dividend yield)
- Short leveraged 2y 85 Call : $16.11 * 1.1765 = 18.953$ (assuming a 21% IV, a 3.5% dividend yield and 1% IR)
- Long 2y ATM Call : 8.53 (assuming a 20% IV, a 3.5% dividend yield and 1% IR)

As you can see from the difference in implied volatilities between the OTM and the ATM calls, we assumed a bit of a volatility skew.

14.6.2 Risk Analysis : The Greeks

You can analyze the greeks using our certificates pricer and selecting 'Airbag Certificate' in the certificate type. I do not think there will be much of a surprise for this instrument. In any case, feel free to ask me any questions if you find the dynamics of a product strange for whatever the reason.

14.7 Outperformance Certificates


Outperformance certificates enable you to participate disproportionately in price advances in the underlying instrument if it trades higher than a specified threshold value. To that purpose, these certificates are equipped with a strike price and a participation factor. Depending on the product maturity, the participation factor usually lies between 120% and 200%.

While the outperformance certificate does not appear directly in the list of certificate types available in our certificates pricer, you can easily price it using the ‘Outperformance bonus certificate’ and setting the downside put barrier and the bonus level at the same level. Note that this pricer will not let you enter a downside put barrier/strike higher than the spot level.

14.7.1 Payoff

Outperformance certificates are constructed as a combination of a long zero-strike call and a long (Participation level - 1) ATM Call. The participation level represents the speed at which the product’s payoff will increase above the strike.

The below graph shows the payoff of a 2y outperformance certificate with a 165% participation level on an underlying asset at 100 with a 3.5% dividend yield and with interest rates hypothetically set to 2%.



The price of this outperformance certificate is 99.27. This price can be split into :

- Long 2y 0 Call : 93.24 (assuming a 3.5% dividend yield)
- Long Leveraged 2y ATM Call : $0.65 \times 9.29 = 6.03$ (assuming a 20% IV, a 3.5% dividend yield and 2% IR)

14.7.2 Risk Analysis : The Greeks

Risk analysis should be very straightforward with a delta larger than 1 (delta-1 + partial investment in an ATM call), a bit of vega, gamma and theta due to the partial investment in an ATM call. If you are not very comfortable with this, please refer to chapter 5 for a fresh reminder.

14.8 Outperformance Bonus Certificates


As its name indicates, it combines the strengths of both outperformance and bonus certificates by protecting the investor on the downside by a bonus level

while still offering the opportunity to participate disproportionately in upside gain in the underlying asset. Obviously, there is no free lunch and the bonus level, the barrier level and the participation level will not be as attractive in an outperformance bonus certificate as they would be in a bonus certificate and an outperformance certificate respectively.

14.8.1 Payoff

Outperformance Bonus certificates are constructed as a combination of a long zero-strike call, a long down-and-out put striking at the bonus level and a long (Participation level - 1) OTM Call striking at the bonus level.

The below graph shows the payoff of a 2y outperformance bonus certificate with a 112.5 bonus level, a 80 barrier level and a 150% participation level on an underlying asset at 100 with a 3.5% dividend yield and with interest rates hypothetically set to 2%.



The price of this outperformance bonus certificate is 98.87. This price can be split into :

- Long 2y 0 Call : 93.24 (assuming a 3.5% dividend yield)
- Long 2y American DOP with barrier at 80 and strike at 112.5: 2.93 (assuming a 21% IV and a 3.5% dividend yield)
- Long 2y Leveraged OTM call: $0.5 \times 5.4 = 2.7$

14.8.2 Risk Analysis : The Greeks

While the delta of the outperformance bonus certificate is generally larger than 1, the other greeks can take quite different values depending on the product's parameters.

For example, the above example has a slightly negative vega at inception. If you take the same structure with a barrier level at 77, a bonus level a 103 and a participation level of 145%, the vega will now be slightly positive at inception. While this result is not surprising, it is still something you need to remind yourself about as you could quickly be fooled.


14.9 Twin-Win Certificates

As the ‘twin-win’ indicates, the investor can actually win it both ways. This product has 100% participation to the upside (not capped) and 100% participation to the downside in absolute terms as long as the barrier was not breached during its lifetime. If the barrier is breached, the twin-win transforms itself in a tracker. In other words, the losses recorded down to a specified barrier level are converted into profits. For that to occur, the barrier level may never be so much as touched during the certificate’s entire term to maturity. If by expiration there has never been a breach of the barrier, you will receive cash payment of the difference between the closing price of the underlying instrument and the level where it stood on the issuance date. In any case, you can rest assured that your twin-win certificate will never perform worse than the underlying instrument.

14.9.1 Payoff

Twin-win certificates are constructed as a combination of a long zero-strike call and long two DOP, where strike is set ATM.

The below graph shows the payoff of a 2y twin-win certificate with a 70 barrier level on an underlying asset at 100 with a 3.5% dividend yield and with interest rates hypothetically set to 2%.



The price of this twin-win certificate is 98.83. This price can be split into :

- Long 2y 0 Call : 93.24 (assuming a 3.5% dividend yield)
- Long two 2y American DOP with barrier at 70 and strike at 100: $2 \times 2.795 = 5.59$ (assuming a 24% IV and a 3.5% dividend yield)

14.9.2 Risk-Analysis : The Greeks

The maturity of the product plays an important role as the twin-win delta usually amounts to a bit less than 1 at inception. Hence, the positive performance that should be reflected in the mark-to-market price of the product in a bearish market only ‘grips’ when around 70% of the time to maturity has expired. For

example, a delta of 0.98 means that if our stock price goes from 100 to 99, the twin-win price will go down by approximately 0.98. However, the payoff is larger when the stock price is at 99 rather than 100. Since there is a large time left to maturity, this decrease in price has increased the probability of breaching the downside barrier, explaining the negative effect on the twin-win price.

I advice you to play with the ‘maturity’ and ‘stock price’ parameters of the Twin-win certificates pricer to observe the interesting behaviour of the greeks when the product is getting closer to maturity.

For example, the delta of the above twin-win will not be of the same sign if you are at 105 or at 95 close to maturity. You will see that delta and gamma can get quite large in absolute value as you get closer to the barrier near to maturity. Depending on the relative distance of the stock price with respect to the barrier, the vega position can also take different signs.

Well, there are so many things that could be said about the greeks analysis of structured products. I have developped and shared the pricer exactly for this reason. It sums it all with graphical representations that will help you developing your understanding and your intuition behind the dynamics behaviour of the greeks.

Chapter 15

Multi-Asset Options

15.1 Introduction

Many payoffs that exist today are based upon the performance of multiple assets. When an option derives its value from the price of multiple assets, the relationships between these assets become important. These relationships between the underlyings are defined using correlations. Correlation gives us the strength and direction of a linear relationship between different underlyings.

We will start this chapter by a small reminder on the concept of correlation, the different types of correlations, their properties, etc.

15.2 Correlation

15.2.1 Realized/Historical Correlation

If you have never heard about correlation, please read about it first as this sections only serves as a small reminder.

The statistical correlation between two variables is measured as:

$$\rho_{X,Y} = \frac{\sigma_{X,Y}}{\sigma_X \sigma_Y}$$

This statistical correlation, also called realized correlation, will take on values between -1 and $+1$.

A **negative correlation** indicates that, historically, as one variable has moved up the other has moved down.

A **positive correlation** means that historically both variables have generally moved in the same direction.

The case of zero correlation means the two variables move in a generally random manner comparatively.

We must stress that measuring correlation as such gives us information regarding the linear relationship between two variables. A correlation of zero means that there is no linear relationship between the two variables but does not imply that these variables are independent. There can be strongly non-linearly correlated! To obtain a more thorough view of the dependence of several variables than a linear relationship between them, one can use **copulas**.

A couple of practical things can be said about correlation.

First, in practice, one computes the correlation between two financial assets using the historical series of daily log returns. It is very important that the series used to compute the correlations have matching dates. It seems quite obvious but can become problematic when the assets are quoted on different markets.

Then, correlations change dramatically through time. Therefore, the linear relationship described by the historical correlation will not necessarily hold in the future.

Finally, during a market crash, the realized correlation between financial assets tend to increase sharply and we can witness assets realizing a correlation of above 90%.

15.2.2 Correlation Matrices

A correlation matrix M_ρ is a square matrix that describes the correlation among n variables.

$$M_\rho = \begin{pmatrix} 1 & \rho_{1,2} & \dots & \rho_{1,n} \\ \rho_{2,1} & 1 & \dots & \rho_{2,n} \\ \dots & \dots & \dots & \dots \\ \rho_{n,1} & \rho_{n,2} & \dots & 1 \end{pmatrix}$$

For this matrix to have any meaning, the pairwise correlations must all be computed over the same period.

This matrix is **symmetric**. Indeed, the correlation between asset i and asset j must be the same as the correlation between asset j and asset i .

The correlation matrix is also necessarily **positive definite**. If the correlation matrix is not positive definite, then it must be modified to make it positive definite. How can you achieve this? Well, you can read about the excellent paper of Nicholas Higham to know more about this.

The good news is that when testing the multi-asset options pricer, you will be provided with a quasi-randomly correlation matrix that is already symmetric

and positive definite. You can always change it manually and the pricer will let you know if the new matrix does not respect these two constraints.

Given a basket of n assets $S_1, S_2, \dots, S_n$ with respective weights $w_1, w_2, \dots, w_n$, the realized basket correlation is defined as the weighted average of the realized correlation matrix between the components, excluding the diagonal of 1's:

$$\rho_{realized} = \frac{\sum_{1 \leq i < j \leq n} w_i w_j \rho_{i,j}}{\sum_{i < j}^n w_i w_j}$$

15.2.3 Portfolio Variance

The link between volatility and correlation is very important as it will allow us to better understand the cega (sensitivity to the correlation) profile of multi-asset options. So let us now see the implications of correlation on the variance of a portfolio.

$$\sigma_{Ptf}^2 = \sum_{1 \leq i \leq n} w_i^2 \sigma_i^2 + 2 \sum_{1 \leq i < j \leq n} w_i w_j \rho_{i,j} \sigma_i \sigma_j$$

As long as the correlation in the above formula is less than 1, holding various assets that are not perfectly correlated in a portfolio will offer a reduced risk exposure to a specific asset.

15.2.4 Implied Correlation

The market for European options on baskets of underlyings is not liquid enough to imply a correlation between the underlyings from these prices.

However, if we have both European options on an index and on each of its n underlyings, then we can use market quotes to infer an implied correlation that is a measure of the dependence between the components of the index.

$$\rho_{index \text{ implied}} = \frac{\sigma_{index}^2 - \sum_{i=1}^n w_i^2 \sigma_i^2}{2 \sum_{1 \leq i < j \leq n} w_i w_j \sigma_i \sigma_j}$$

To obtain the implied correlation over a T -day period, we must use the IVs of options with time to maturity T . In this case we make use of ATM volatilities.

The previous formula came from the formula of the variance of a portfolio. This formula assumes that all weights are constant. In the case of an index, this is usually not the case as the weights vary as the components of the index vary. This makes the above formula inexact for an index.

However it still has some implications and uses. Assume that we have a basket of stocks for which we wish to infer an implied correlation. Assume further that these stocks all belong to the same index.

The idea is to follow a simple parameterization involving a coefficient λ which relates realized and implied correlations of the index, and in turn use this coefficient and also the realized correlations between the index components to infer specific implied correlations.

Firstly, compute the realized correlation of the index, and the implied correlation using formula, then solve for λ in the equation:

$$\rho_{index\ implied} = \rho_{index\ realized} + \lambda (1 - \rho_{index\ realized})$$

If you take two stocks that belong to the same index and for which we have liquid European options on both the index and its components, you can then obtain the value of λ using the above formula and use it alongside the realized correlation of the two stocks to estimate the implied correlation between them.

Sell-side desks of multi-asset options will typically be structurally short correlation. This is due to the worst-of feature in most of the multi-asset structured products. No worries at this stage if you are not yet familiar with the structural correlation position of sell-side desks due to worst-of products. We will discuss them in further details in this chapter and you will quickly get used to it while using the multi-asset options pricer.

As with volatilities, implied correlations will usually be higher than realized correlations (positive λ).

15.2.5 Correlation Skew

Assume we have two assets for which we have liquid vanilla options at different strikes for each of them as well as for the basket. If we imply a correlation at each strike where we used the implied volatilities for the basket and the two constituents, would the implied correlations be the same? Not necessarily so. Plotting implied correlations with respect to the strikes gives a curve known as the correlation skew.

As with implied volatility, implied correlation tends to be higher for lower strike.

Many exotic products have correlation skew exposure in the sense that as the underlying assets move, the correlation sensitivity can vary significantly.

This correlation skew represents an additional risk and you must use a model that knows about it for this risk to be reflected in the product's price.

In the same way, implying correlations may also give rise to a correlation term structure. From a modeling perspective, having a correlation term structure is typically less computationally intense than a correlation skew. To go deeper into the concept of a correlation skew, and have a meaningful method to see this in a model, we will need to look at copulas.

15.2.6 Copulas

A copula is a multivariate probability distribution for which the marginal probability distribution of each variable is uniform. Copulas are used to describe the dependence between random variables.

Sklar's theorem states that any multivariate joint distribution can be written in terms of univariate marginal distribution functions and a copula which describes the dependence structure between the variables.

The importance of copulas in multi-asset derivatives is that they provide a method of expressing joint distributions between assets, allowing for the simulations of these variables, and thus the pricing of multi-asset options.

The traditional correlation does not allow us to specify different behaviours for different parts of the distribution. A copula does precisely this. Modelling with copulas can therefore provide more meaningful hedge ratios.

15.3 Basket Options

15.3.1 Description

A basket option is an option whose payoff depends on the value of a basket of assets.

At maturity, it pays off the greater of zero and the difference between the average return of the n different assets in the basket and the strike price (expressed as a return in this case):

$$\text{Basket Call}_{\text{payoff}} = \max \left[0, \sum_{i=1}^n w_i R_i - K \right]$$

The typical underlying of a basket option is a basket consisting of several stocks, that represents a certain economy sector, industry or region.

The main advantage of a basket option is that it is cheaper to use a basket option for portfolio insurance than to use the corresponding portfolio of plain vanilla options. Indeed, a basket option takes the imperfect correlation between the assets in the basket into account and moreover the transaction costs are minimized because an investor has to buy just one option instead of several ones. To avoid the problems that could potentially arise from having to deliver multiple underlyings, multi-asset options are generally cash settled.

Unlike a rainbow option in which the weighting at maturity is based on the relative performance of the various assets, the weights of each of the underlyings are known at the outset in a basket option. As such, the basket is different from

an index in that the weights in a basket stay the same, whereas in an index they can change as the components of the index move.

Since you can specify the weight of each underlying asset, our multi-asset options pricer is not limited to equi-weighted basket. However, you need to make sure that the weights sum up to 1 for the pricer to work.

Since basket options are vanilla options on a basket. The intuition behind all the greeks discussed in chapter 5 applies. We will therefore focus on a new sensitivity parameter, the Cega.

15.3.2 Cega, sensitivity to correlation

It is easily seen from the portfolio variance formula ($\sigma_{Ptf}^2 = \sum_{1 \leq i \leq n} w_i^2 \sigma_i^2 + 2 \sum_{1 \leq i < j \leq n} w_i w_j \rho_{i,j} \sigma_i \sigma_j$) that an increase in correlation implies an increase in the overall basket volatility. Since options have positive vega, the seller of a basket option is thus selling the basket volatility which, in turn, implies that he is selling the correlation between the underlyings. From this formula, you can also imply that the price sensitivity of a basket option to a movement in correlation is not linear.

Using our multi-asset options pricer, you can easily compute the sensitivity of a basket option to a specific correlation pair. To do so, you can simply bump the correlation between the two underlyings by 1% and reprice the basket option to see the difference in prices. You should normally make sure that the correlation matrix is still valid before repricing the basket option but our pricer does it for you ;).

Our pricer allows you to calculate the cega of the basket option. The cega gives you the effect of an overall move in correlations by 1% on the basket option price. It is calculated by bumping the entire matrix of correlations (understand the non diagonal elements of the correlation matrix), making sure the correlation matrix is still valid, and repricing the basket option. You can do it manually and see if you obtain the same result...

Obviously, your sensitivity to correlation depends on the assets' weight. For example, a basket option on 3 stocks weighted 96%/2%/2% will not have any sensitivity to correlation since it will behave very much like a single-stock option.

15.3.3 Pricing methods / Modelling

In the absence of volatility smile/skew, basket options can be priced using various techniques, including moment-matching or geometric conditioning. The moment-matching methodology consists of modelling the basket as a single log-normal asset so that the classical Black-Scholes formula can be applied. Basically, you find an equivalent log-normal random variable that has the same

mean and variance as the basket (which is a weighted basket of log-normal random variables).

You can find an analysis of different pricing methods for basket options in this article from Krekel, de Kock, Korn and Man.

In practice, you may want to simply use a Monte Carlo simulation of correlated log-normal random variables. In the case where the basket option has skew dependence, you will need to use a model that knows about skew.

If you have sufficient liquid underlyings for the points to which the basket option has vega exposure, then the calibration of individual loc-vol models to these skews and a simulation of these correlated variables will be enough.

Otherwise, you can handle basket skew by using an index skew as a proxy. You will then have to compare the time series of volatilities to decide the level at which to buy/sell volatility if the basket option's Vega is to be hedged with options on the index.

If you are interested about reading a more technical article on pricing basket options with skew, you can read the following article from Dong Qu.

15.4 Rainbow Options

When analyzing basket options, we only spoke about the impact of correlations on the portfolio volatility. This is because basket options are almost not sensitive to the dispersion effect. A higher dispersion does not affect the expected return of the basket. It will not increase the probability of the basket to finish ITM. This is because the weights assigned to the underlying assets do not depend on their respective performance.

For the rest of this chapter, it is quite important to understand how volatility and correlation affect dispersion. If the pairwise correlations between the underlying assets is low, the returns of these underlyings would be quite apart from each other and vice versa. Also, a higher asset volatility leads to asset returns with large deviation from its expected return. Hence, a higher volatility and a lower correlation leads to higher dispersion.

15.4.1 Description

A rainbow is an option on a basket that pays in its most common form, a non-equally weighted average of the assets of the basket according to their performance.

For instance, a rainbow call on 3 assets with weights 50%/30%/20% pays 50% of the best return of the underlying assets at maturity, 30% of the second best return and 20% of the third best.

Rainbow can take various other forms but the combining idea is to have a payoff that is depending on the asset on the assets sorted by their performance at maturity. For instance, worst-of and best-of options can be seen as a particular type of rainbow options and will be describe in the next section.

Note that when pricing rainbow options with our multi-asset options pricer, make sure to set the weights from worst to best. This might not seem very intuitive but it is the way I coded it at the time.

15.4.2 Risk Analysis

As with basket options, the sensitivities to interest rates, dividends and borrowing costs are intuitive. Sensitivity to volatility, correlation and skew are very hard to grasp. It depends on many parameters and the only way I see to determine your position regarding them is to compute the option price using different levels of volatility, skew and correlation.

For instance, the sensitivity to correlation is difficult to predict because there might be two opposite effects. It is the case in the above example. On one hand, increasing correlation would increase the overall basket volatility and will tend to push the option price higher. On the other hand, increasing correlation would decrease the forward price of the basket and will tend to push the option price lower.

Seeing our 50%/30%/20% rainbow option as $0.9 * \text{Equiweighted Basket} + 0.2 * \text{Best-Of} - 0.1 * \text{Worst-Of}$ should help you understand the second impact.

As with any multi-asset options, the rainbow option will show sensitivity to the implied volatilities of each of the underlyings. The way the vega is spread around the underlying assets depends on the probability of each asset to finish at a particular ranking at maturity (best, second best or third best).

If the weights have values far from each other, the option will behave similarly as a best-of/worst-of and the position regarding volatility and correlation is easy to grasp. Generally, the more dispersed the weights, the more expensive the rainbow option.

15.5 Worst-Of options

We only focus on Worst-Of options as everything that is said about them can be applied to Best-Of options quite easily. You can also price Worst-Of options using our multi-asset options pricer.

15.5.1 Worst-Of Put

As its name indicates, a worst-of put is a put option on the worst-performing underlying asset among a basket. It is obviously more expensive than a vanilla put option on the same basket.

$$\text{WO Put}_{\text{payoff}} = \max[0, K - \min(S_1(T), S_2(T), \dots, S_n(T))]$$

As previously mentioned, the sensitivities to interest rates, dividends and borrowing costs are intuitive.

15.5.1.1 Sensitivity to underlyings' price

A WO put option is still a put option and therefore the intuition behind the sign of the delta is straightforward.

The cross-gamma effect between the underlyings is more interesting. As an underlying asset declines and takes the role of the Worst-Of, you will expect its delta to increase in absolute value and, at the same time, the delta on the other underlyings to decrease in absolute value. That is what we call a cross-gamma effect. A movement on a underlying asset will typically have an impact on the delta of the other underlyings. The more probability an underlying asset has to finish as the Worst-Of, the more the delta will be allocated to this underlying asset. This cross-gamma sensitivity is more pronounced when the forward price of the underlying assets are close to each other (it is not clear which asset will be the worst-of). More generally, the more probability an underlying asset has to finish as the Worst-Of, the more the greeks will be allocated to this underlying asset. So the option will also show higher vega to the volatilities of the underlyings that perform the worst for example.

15.5.1.2 Sensitivity to volatility

We already know that the holder of a vanilla put option is long volatility. At the same time, higher volatility would lead to higher dispersion which again increases the price of the option. Consequently the buyer of WO put option is long volatility.

15.5.1.3 Sensitivity to correlation

A higher dispersion would lead to a higher payoff. Since lower correlation would lead to more dispersed returns, lower correlation would lead to higher payoff for worst-of put options. Therefore the buyer of a WO put option would be short correlation.

Just remind yourself that the holder of a worst-of/best-of is long dispersion and therefore short correlation and long volatility.

15.5.1.4 Sensitivity to volatility skew

Skew results in a return distribution that are negatively skewed with higher probability of downward movements leading to higher implied volatilities on the downside. Since Worst-Of put options pay on the downside, an increase in skew will raise their prices.

15.5.1.5 Sensitivity to correlation skew

Let us assume we are long a WO put option and therefore short the correlation between the underlying assets.

If the market goes down, the WO put becomes more ITM and hence we will be shorter correlation at a time when correlation tends to spike up.

If the market goes up, the correlation tends to go down but the WO put will be OTM and our sensitivity to correlation decreases.

In the first case, the skew hurts us and, in the second case, it does not benefit us. As a consequence, a WO put option is cheaper when the correlation skew is taken into account.

15.5.2 Worst-Of Call

As its name indicates, a worst-of call is a call option on the worst-performing underlying asset among a basket. It is obviously less expensive than a vanilla call option on the same basket.

$$\text{WO Call}_{\text{Payoff}} = \max[0, \min(S_1(T), S_2(T), \dots, S_n(T)) - K]$$

15.5.2.1 Sensitivities

You should be able by now to deduce most of the greeks and their behaviour for a Worst-Of call.

I will just emphasize the effects of volatility and correlation as the position in terms of dispersion is different as for a WO put option. Indeed, a WO call price decreases as dispersion goes up. Based on dispersion, you could think that the holder of a WO call is short volatility and long correlation. However, the position in volatility is not this obvious as volatility now has two opposite effects on the WO call price. On the one hand, volatility increases the expected payoff of the call (typical volatility effect that you see on a vanilla option). On the other hand, it increases dispersion which lowers the forward level of the worst performing underlying and thus decreases the expected payoff of the WO call.

Most of the time, the dispersion effect is dominant but that is not always the case (for example in a high correlation environment). It is also quite common to see a trader that is selling the WO call being short volatility on one underlying, usually the one with the highest volatility, and long volatility on the other underlyings.

Since the position in volatility is not clear, we don't know whether the WO call option holder would benefit or lose due to volatility skew. Hence skew dependence would again depend on the actual trade parameters.

As previously discussed with WO put options, you can also expect many cross-risks with WO call options.

Chapter 16

Autocallables

Autocallables are probably the most popular structures in the world of structured products. In this chapter, we will look at standard autocallables and some of their variants.

16.1 Description

The standard autocallable is a note that only pays a coupon if the underlying asset (S) is above a certain coupon barrier level (CB) and the note automatically redeems early if it breaches an autocall barrier level (AB), which can be the same or different as the coupon barrier level, at an observation date. The autocallability feature can be whether daily, monthly, yearly, or based on any schedule determined in accordance with the client. The coupon barrier level is always lower or equal to the autocall barrier level, otherwise the product does not make much sense. When the $CB = AB$, I will call the structure an Autocall and when $CB < AB$, I will call the structure a Phoenix.

An autocallable is clearly a yield enhancing strategy. The investor receives an above market yield and in turn takes the risk not to receive any coupon if the underlying asset is below the coupon barrier level on the observation date.

In this low-yield environment, it is impossible to offer attractive coupons with such a structure. Therefore, to get an enhanced coupon the investor typically also sells an option embedded in the note (usually a down-and-in put or a leveraged OTM put). We will start with analyzing the most basic autocallable structure and then add the most frequent features observed in the market.

Because of the autocall feature, this structure does not have a fixed maturity. What we call maturity is in fact the maximum duration this product can stay alive.

16.2 Payoff

At each observation date t_i with $i = 1, \dots, n$, we have :

$$\text{Coupon}(t_i) = \text{Notional} * C\% * 1_{\left\{\frac{S_{t_i}}{S_{t_0}} \geq CB\right\}} * 1_{\left\{\max_{j=1, \dots, i-1} \frac{S_{t_j}}{S_{t_0}} \leq AB\right\}}$$

and

$$\text{Redemption}(t_i) = \text{Notional} * 1_{\left\{\frac{S_{t_i}}{S_{t_0}} \geq AB\right\}} * 1_{\left\{\max_{j=1, \dots, i-1} \frac{S_{t_j}}{S_{t_0}} \leq AB\right\}}$$

Once the autocall barrier is breached on one of the observation dates, the note dies and there is no future payment afterwards.

Note that the autocall barrier level can be fixed during the life of the option as well as variable. It is not rare to see it decremental for example. This makes sense as the conditional probability of breaching a fixed autocall barrier is getting lower at each observation date. These conditional probabilities are shown in our autocallables pricer.

As you can see, the above payoff description does not contain any short position in option. This will be added and discussed later on. For now, it basically looks like a stream of conditional coupons (digitals) with an autocall feature.

16.3 Risk Analysis

As usual, the price of the product is nothing but the present value of the expected future cashflows. To calculate the expected future cashflows, we start by computing the undiscounted conditional probabilities of receiving the coupons. It will help you to develop a feeling whether the pricing makes sense or not. Once these probabilities are computed, they should be multiplied by the coupons and then discounted. These undiscounted conditional probabilities of receiving the coupons are also made available in our autocallables pricer

The first coupon is a classical European digital. All the next coupons are conditional on not autocalling at any previous observation dates. As time goes by, the probabilities of coupons being paid decrease. The value of the last path-dependent coupons can therefore be very small. As a seller, you should really be careful when offering a very large digital with a low probability.

It is no surprise that the risks associated with this basic autocallable structure is similar to those associated with digitals. That is the seller of an autocallable is short the underlying's forward (short interest rates, long dividends and long borrowing costs) and short the skew. The position in volatility depends on the relative position of the forward price of the underlying with respect to the coupon barrier level. The overall vega is split over the observation dates, we speak about vega buckets. Each of these vega sensitivities will change as the

market moves. If an autocall event is about to happen, the ST vega will increase while the other vega buckets will decrease. The vega hedge will therefore be dynamic and readjusted depending on the evolution of the underlying asset.

As the structure has a zero coupon bond embedded in it, it also has an important first order sensitivity to interest rates moves. More precisely, the structure has an overall negative exposition to interest rates.

Whenever the autocallable structure is denominated in a currency that is different from the underlying's currency, there is an exposition to the correlation between the underlying's return and the forex return. We have studied this FX exposure in chapter 10 and have seen that there is an impact on the dynamic delta hedging.

16.3.1 Equity / Interest Rate Correlation

This aspect of the pricing is often neglected but the correlation between equity and IR has an effect on the autocallable price. This cannot be seen from our pricer but will often be tackled during interviews.

Since the note is redeemed at an unknown time in the future that is dependent on the performance of the underlying stock, the autocallable structure is sensitive to the correlation between IR and the equity underlying. To get an idea about this effect, we use a hedging argument since the price of a product should reflect the cost of hedging it.

Let us start by assuming that this correlation is positive so that if the underlying increases, then interest rates also tend to increase. Since an increase in the underlying will increase the probability of early redemption after 1y, the autocallable seller will sell some of the 2y bonds to buy more of the 1y bonds. We know that the 2y ZC bond will be more impacted by the IR increase than the 1y ZC bond (as it has a longer duration). All in all, this means that the seller will net a loss on his IR rebalancing since he will be selling the bond that decreased more in value to buy the one that decreased less.

In the same way, if the underlying decreases, then interest rates also tend to decrease. Since a decrease in the underlying will decrease the probability of early redemption after 1y, the autocallable seller will sell some of the 1y bonds to buy more of the 2y bonds. We know that the 2y ZC bond will be more impacted by the IR decrease than the 1y ZC bond (as it has a longer duration). All in all, this means that the seller will net a loss on his IR rebalancing since he will be selling the bond that increased less in value to buy the one that increased more.

Let us now assume that this correlation is negative so that if the underlying increases, then interest rates tend to decrease. Since an increase in the underlying will increase the probability of early redemption after 1y, the autocallable seller will sell some of the 2y bonds to buy more of the 1y bonds.

We know that the 2y ZC bond will be more impacted by the IR decrease than the 1y ZC bond (as it has a longer duration). All in all, this means that the seller will net a profit on his IR rebalancing since he will be selling the bond that increased more in value to buy the one that increased less.

In the same way, if the underlying decreases, then interest rates tend to increase. Since a decrease in the underlying will decrease the probability of early redemption after 1y, the autocallable seller will sell some of the 1y bonds to buy more of the 2y bonds. We know that the 2y ZC bond will be more impacted by the IR increase than the 1y ZC bond (as it has a longer duration). All in all, this means that the seller will net a profit on his IR rebalancing since he will be selling the bond that decreased less in value to buy the one that decreased more.

To sum up, the autocallable price should be higher when a positive Equity/IR correlation is assumed and lower if this correlation is assumed to be negative. It makes sense when you think about the investor's position as a positive correlation means that he is likely to get his money back in a high interest rate environment, which is a pretty good scenario for him.

Knowing this, you could employ a model that includes stochastic rates in order to enter a value for this correlation and include its impact in the price. Our pricer does not go that far as using such model does not give us additional information regarding the hedging of this Equity/IR correlation. This correlation cannot be hedged in a straightforward manner. To hedge this correlation risk, one would need a liquid structure involving equity and interest rates, from which one could extract this correlation by hedging away the other parameters. Most often than not, it cannot be hedged at all. You could therefore decide on which correlation to use and add a cost accordingly without having to employ such a stochastic rate model. The magnitude of this cost will clearly be a positive function of the maturity of the structure. You can understand why by thinking about the above hedging argument. If the maturity of the autocallable would have been 3 years, then the hedging argument would have implied rebalancing 1y and 3y ZC bonds. Since the duration of the 3y ZC bond is longer than the 2y ZC bond, the impact of an IR move would have been more pronounced. So the longer the maturity of the autocallable structure, the greater the impact of the Equity/IR correlation.

16.4 Memory Feature

16.4.1 Autocall Incremental

In a standard Autocall in which the coupon and autocall barriers are at the same level, the coupon has a memory feature. We say that the coupon is incremental. It means that once the product is redeemed, the investor receives a coupon equal to the sum of all previous coupons.

At each observation date t_i with $i = 1, \dots, n$, we have :

$$\text{Redemption}(t_i) = \text{Notional} * (1 + i * C\%) \left\{ \frac{S_{t_i}}{S_{t_0}} \geq AB \right\} * 1 \left\{ \max_{j=1, \dots, i-1} \frac{S_{t_j}}{S_{t_0}} \leq AB \right\}$$

In this payoff, the digital size increases with time and, as a seller, one must be careful when offering a very large digital with a very low probability of striking.

16.4.2 Phoenix memory

In a Phoenix, the coupon barrier level is smaller than the autocall barrier level. In such a case, the memory feature is optional. In case of memory feature, we speak about Phoenix memory and the payoff becomes:

16.4.2.1 Coupons

At the first observation date t_1 , we have :

$$\text{Coupon}(t_1) = \text{Notional} * C\% * 1 \left\{ \frac{S_{t_1}}{S_{t_0}} \geq CB \right\}$$

At each observation date t_i with $i = 2, \dots, n$, we have :

$$\text{Coupon}(t_i) = \left(\text{Notional} * \left[(i * C\%) - \sum_{j=1}^{i-1} \text{Coupons} \right] \right) * 1 \left\{ \frac{S_{t_i}}{S_{t_0}} \geq CB \right\} * 1 \left\{ \max_{j=1, \dots, i-1} \frac{S_{t_j}}{S_{t_0}} \leq AB \right\}$$

16.4.2.2 Redemption

At each observation date t_i with $i = 1, \dots, n$, we have:

$$\text{Redemption}(t_i) = \text{Notional} * 1 \left\{ \frac{S_{t_i}}{S_{t_0}} \geq AB \right\} * 1 \left\{ \max_{j=1, \dots, i-1} \frac{S_{t_j}}{S_{t_0}} \leq AB \right\}$$

16.5 Put Feature

In this low-yield environment, it is impossible to offer attractive coupons with such a structure. Therefore, to get an enhanced coupon the investor typically also sells a put option embedded in the note: a down-and-in put or a leveraged OTM put whose maturity is the maturity of the autocallable structure. This means that the investor's capital is no longer protected.

To price this structured product, one should you deduct the price of the put option from the previously explained price.

Adding a put feature does not change the seller's position with respect to the forward price of the underlying asset: he is still short the forward. However, his overall position in volatility and skew are not straightforward as there are potentially offsetting effects from the digitals (coupons) and the put.

While a trader selling an autocallable is always long volatility with respect to the put, his vega with respect to the autocallable digitals depends on the relative position of the forward price with respect to the coupon barrier level. Most of the time, it results in a long position in volatility. We have seen that it was relevant to analyze the specific vega sensitivities across time until maturity. In this way, we should rather evaluate the vega for each time period by progressively shifting parts of the volatility term structure. This process makes it possible to clearly conclude on the local exposition to volatility across varying maturities. Moreover, not only Vega exposure can differ across volatility maturities, but it can also be ambiguous for a single volatility maturity. Anyway, despite these local variations across the term structure, the investor is globally selling volatility: the put has the biggest impact on the Vega.

As for the skew position, the seller of an autocallable is short skew with respect to the digitals but long skew with respect to the long position in the downside put. Most of the time, it results in a short skew position.

I invite you to review the risk analysis of the down-and-in put option as all the risks associated with it are present in most of the autocallable structures.

16.6 Smoothing the autocall barrier

It is important for you to understand that each barrier (DIP barrier, coupon barrier, autocall barrier) has some sort of digital associated with it. Therefore, everything that has been said in chapter 12 about the importance of smoothing discontinuities applies. There are many discontinuities linked to the autocallable's payoff: around the DIP barrier, around the coupon barrier, around the autocall barrier. These discontinuities have crucial consequences in terms of risk management and hedging as they imply potentially unstable and explosive Greeks.

The autocall barrier at each observation date represents an up-and-out barrier that can be approximated by a call spread with spread between $AB-\epsilon$ and AB . The thing is that it is often impossible to determine ex-ante the direction of the discontinuity in the payoff. Because of this, you would not know now in what direction you should shift the autocall barrier.

If **continuation value** < **exit value** $\rightarrow$ *epsilon* > 0

If **continuation value** > **exit value** $\rightarrow$ *epsilon* < 0





Figure 16.1: Fig: 16.1 : Barrier shift when continuation value < exit value



In theory, the exit cashflow received in a knock-out event is known at the barrier's observation date t_i . At t_i , the exit value is simply this exit cashflow.

On the contrary, the continuation cashflows received in the absence of a knock-out event depend on the underlyings' trajectory after the barrier's observation date t_i . At t_i , the continuation value is the expected value of these continuation cashflows. Evaluate this continuation value implies evaluating the expected value of future cashflows at each observation date t_i for each simulated trajectory so that $|S_{t_i} - AB| < \epsilon$. The available methods such as Monte Carlo or American Monte Carlo are not much robust.

To apply the optimal smoothing (most conservative price), it is not necessary to know the continuation value. It is enough to know the sign of the discontinuity **Exit value - Continuation value**.

By definition, the sign $\delta_{optimal} \in \{-1, 1\}$ of the shift that guarantees an optimal smoothing verifies : $\delta_{optimal} = \underset{\delta_{optimal} \in \{-1, 1\}}{argmax} \text{ Price with a shift of } \delta * |\epsilon|$.

With N the number of barrier observations, 2^N pre-evaluations enable to define the directions of the optimal smoothing $\delta_{optimal} \in \{-1, 1\}^N$.

When N is large, it is not feasible to test 2^N combinations in a single evaluation.

There are several solutions to speed up the process.

You could divide the problem into M blocks and $2^{N/M}$ combinations per block are tested: * The directions of the smoothing of the last N/M autocall barriers are calibrated during the first calculation by setting to 0 the shift on all the previous barriers. * Then, you move to the next N/M previous barriers by taking into account the previously calculated values from the first calculation

and setting to 0 the shift on all previous barriers. * You repeat this operation M times.

You could also get the directions of the smoothing independently from each other. It requires $2 * N$ evaluations but the solution will not be systematically optimal: * You set all the other shifts to 0. * You test the two combinations: $\{-1,1\}$ and keep the one for which the price is the most conservative. * You repeat this operation N times, meaning that it requires $2*N$ evaluations.

In practice, $\delta(i)$ depends greatly from the smoothing of the next autocall barrier, $\delta(i+1)$. It is therefore desirable to take it into account when evaluating $\delta(i)$ by modifying the previous algorithm as follows: * You set all the other shifts to 0 except the one from $i+1$. * You test the two combinations: $\{-1,1\}$ and keep the one for which the price is the most conservative.

Note that this algorithm neglects the fact that the continuation value at observation date i depends on the smoothing on every observation dates after i (not only the next one). This solution is therefore suboptimal but faster as it requires $2^1 + 2^2 * (N - 1)$ combinations instead of 2^N .

16.7 Multi-Asset: Worst-Of


The payoff is exactly the same except that we observe the performance on the worst performing share of the basket.

In terms of risk analysis, it is quite similar except that there is an additional dispersion dimension. The holder of a WO autocall is long correlation as it increases the probability of receiving the coupons and being autocalled and decreases the probability of being into the Down-and-In Put.

16.8 Other variants

16.8.1 Autocall Twin-Win

It is an autocall structure with a down-and-in put, with the potential of capturing the absolute performance of the underlying at maturity. The name Twin-Wins comes from the fact that this note enables the holder to get a participation in both the upside and the downside movements of the underlying asset if no knock-in event occurred.



As a seller, you must be careful to shift and smooth correctly around the barrier level as the discontinuity is twice the size of the discontinuity in a standard autocall.

16.8.2 Autocall with one star feature

In this variant, even if the worst-performing underlying closes below the Down-and-In Put barrier level, if at least one underlying closes at or above the one star barrier (usually its initial level), the product is redeemed at 100%.

While this one star feature is quite appreciated by the investors, it is not so much by the bank's risk department. The reason behind this aversion is that the cega (sensitivity to correlation) can change sign because of this feature. Can you intuitively explain the impact of the star effect on the sensitivity to correlation?

Chapter 17

Variance Swaps

I was quite surprised to frequently receive questions about variance swaps during zoom meetings with future practitioners eager to step into the industry. So let's add a chapter about these products by summarizing the fantastic article written by Allen, Einchcomb and Granger on this subject.

17.1 Introduction

Variance swaps are instruments which offer investors straightforward and direct exposure to the volatility of an underlying asset without the path-dependency issues associated with delta-hedged options. As their names indicate, they are swap contracts where the parties agree to exchange a pre-agreed variance level for the actual amount of variance realised over a period.

Buying a variance swap is like being long volatility at the strike level; if the market delivers more than implied by the strike of the option, you are in profit, and if the market delivers less, you are in loss. However variance swaps are convex in volatility: a long position profits more from an increase in volatility than it loses from a corresponding decrease. For this reason variance swaps normally trade above ATM volatility.

The directness of the exposure to volatility and the relative ease of replication through a static portfolio of options make variance swaps attractive instruments for investors and market-makers alike.

17.2 Liquidity

Bid/offer spreads have come in significantly over recent years. For indices, they are typically in the region of 0.5 vegas. For single-stocks, they are in the range

of 1-2 vegas in Europe and 2-2.5 in US and Japan. Spreads are naturally higher in emerging markets although these too are becoming more liquid.

The most liquid variance swap maturities are generally from 3 months to around 2 years. The maturities generally coincide with the quarterly options expiry dates, meaning that they can be efficiently hedged with exchange-traded options of the same maturity. The VIX, VSTOXX and VDAX indices represent the theoretical prices of 1-month variance swaps on the SP500, SX5E and DAX indices respectively, and are calculated by the exchanges from listed option prices, interpolating to get 1-month maturity. These volatility indices are widely used as benchmark measures of equity market risk, even though they are only short-dated measures and are not directly tradable.

17.3 Uses of Variance Swaps

Increasingly, investors have come to view volatility itself as an asset class, one that can diversify investment returns or hedge unwelcome investment scenarios.

Some common uses of variance swaps :

Exploiting a volatility view: Ideal for taking a direct view on the volatility of an underlying without the path-dependency issues of a delta-hedged option.

Specific hedging purposes: They can be used for macro-hedging and for hedging specific volatility exposures, such as that resulting from structured products.

Rolling short variance: Short variance swaps can be used to capture the observed equity index volatility risk premium. Rolling short index variance is an attractive systematic volatility strategy from a risk-return perspective. With returns from short volatility trades somewhat uncorrelated with the underlying, these types of strategy work well as overlay strategies aimed at boosting alpha and diversifying returns.

Diversification: Volatility can be thought of as an asset class in its own right, and as such can act to diversify returns within a portfolio. Returns of rolling short volatility index has many similarities to returns of bond index (regular periods of positive P&L and punctual large losses). Short variance swap sometimes replace bonds within an efficiently allocated portfolio as relatively low correlation with equity market.

Index variance spreads: They can be used to trade the spread of volatilities between two indices. Such trades can be thought of as 'volatility-beta' trades aiming to profit from a spread of volatilities widening as volatility increases. It attempts to get & long volatility exposure, but using a correlated index with lower beta to mitigate the carry. Be careful as some indices have local 'regime-changes' as their composition or behaviour of the constituents changes. A typical example is the Nasdaq before and after the dot-com bubble.

Relative value single-stock volatility: Use volatility pairs, or cross-sectional regression volatility models to find rich/cheap single-stock volatilities.

Variance dispersion and correlation trading: Trading variance swaps on an index against variance swaps on its constituents provides exposure to equity correlation. Selling variance on an index and buying variance on its constituents has been a profitable strategy. This can be partly explained by the demand for protection at the index level. This is also another way to trade implied correlation.

Forward variance and volatility spikes: Long forward volatility can avoid potentially negative carry at the cost of slide down the term structure, and can be a useful way of positioning for volatility spikes.

The holder of a long variance swap position will have to pay the spread of implied volatility over realised volatility everyday he holds the swap. Since this spread is usually positive, he has a negative carry. A forward starting instrument will have no exposure to carry until the forward starting date is reached.

On the contrary, both spot and forward starting variance swaps have exposure to slide. For upward sloping variance term structure, a long variance swap position loses out through time as it slides down the term structure.

Trading the variance term structure: Variance swaps can be used to trade the shape of the variance term structure. Using variance swaps, you could try to take advantage of a flattening of the volatility term structure or take advantage of a volatility term structure that appears overly convex for example.

Skew and convexity trades: Variance swaps are long skew and convexity. Trading variance against (delta-hedged) vanilla options provides interesting exposures to skew and/or convexity.

Cross asset class trades: Equity Volatility and credit spreads are correlated, both being measures of corporate risk. Variance swaps are useful instruments in debt/equity trades, either at the index or single name level.

17.4 Mechanics

Realized Volatility

We have already broached this subject in chapter 6 but, to my knowledge, variance swaps use a different volatility measure than the annualised standard deviation of daily log returns over a fixed period of time.

It is the root mean-squared measure that is used to define the payout of the variance swap. This RMS volatility measure is like a standard deviation but assuming a zero mean.

$$\sigma^2 = \frac{252}{T} \sum_{i=1}^T \left[\ln \left(\frac{S_i}{S_{i-1}} \right) \right]^2$$

where

- S_i is the underlying's price on day i
- T is the number of days

Why do we use variance and not volatility?

While volatility is more intuitive as it is measured in the same units as the underlying, variance is in some sense more fundamental as:

- Variance is additive, whereas volatility is not.
- Delta hedging of options seeks to capture realised variance, although exposure is complicated by path dependency issues.

Variance swaps pay-out is purely based on realised variance.

17.5 The Contract

The strike of a variance swap is set at trade inception so that the swap initially has zero value according to the prevailing market conditions.

Obviously, the buyer of a variance swap is long volatility. He pays the strike and receives the realised variance at expiry.

By convention, variance swap strikes are quoted in terms of volatility, not variance.

17.5.1 P&L for a long variance swap

The P&L of a variance swap is non-linear (convex) with volatility, although of course it is linear in terms of variance.

$$\text{P\&L} = N_{VAR} * (\sigma^2 - K^2)$$

where:

- K is the variance swap strike
- σ^2 is realised variance
- N_{VAR} is the variance notional

17.5.2 Vega Notional / Variance Notional

The notional of a variance swap can be expressed either as a variance notional or a vega notional.

As shown in the equation above, the variance notional represents the P&L per point difference between the strike squared (implied variance) and the subsequent realised variance.

Since most market participants are used to thinking in terms of volatility, trade size is typically expressed in terms of vega notional. The vega notional represents the average P&L for a 1% change in volatility.

$$\text{Vega Notional} = \text{Variance Notional} * 2K$$

$$\text{P\&L} = N_{\text{Vega}} * \left(\frac{\sigma^2 - K^2}{2K} \right)$$

The variance swap payout, expressed in vega notional, is locally linear around the strike. In other words, when volatility remains close to the variance swap strike, the variance swap payout is similar to the payout of a linear volatility swap.

For a vega notional of €100k, a gain of €500k is expressed as a profit of 5 vegas (5 times the vega notional).

The daily P&L pattern of a short variance position profits modestly most of the time, but loses heavily on large moves.

17.6 Convexity

As previously mentioned, variance swap payoffs are linear with variance, but convex with volatility.

The vega notional represents only the **average** P&L for a 1% change in volatility.

As it is convex in volatility, a long variance swap position will always profit more from an increase in volatility than it will lose for a corresponding decrease in volatility. This convexity is the reason that variance swaps strikes trade above ATM volatility.

With no surprise, the convexity premium should depend on the expected variability of the realised volatility. The higher the variability of volatility, the more beneficial the convexity.

For a long position, the maximum vega loss of $K/2$ happens when realised volatility.

For a short position, the loss is potentially unlimited unless the variance swap is capped. With the standard cap of 2.5%, the maximum vega loss would be 2.625K.

Example of variance swap convexity:

Let us take a 20-day short variance swap position with a strike at 16.5 and a vega notional of €100k. Let us assume that over the 20-day period, the realised volatility was 14%, 2.5% lower than the level of implied volatility sold. You would expect a P&L for this short position of €250k. In reality, the actual P&L will be slightly less than this due to the convexity.

17.7 Mark-to-Market

Marking to market of variance swaps is easy since variance is additive.

$$V_T = V_t + V_{t,T}$$

At an intermediate point in the lifetime of a variance swap t , the expected variance at maturity V_T is simply the time-weighted sum of the variance realised over the time elapsed V_t and the implied variance over the remaining time to maturity $V_{t,T}$.

To compute the mark-to-market of a variance swap, you need:

- The realised variance since the start of the swap
- The implied variance from now until expiry
- A discount factor between now and expiry

$$\text{Variance P\&L at time } t \text{ per unit of variance notional} = \frac{t}{T}(\sigma_{0,t}^2 - K_{0,T}^2) + \frac{T-t}{T}(K_{t,T}^2 - K_{0,T}^2)$$

You can see from this formula that the total remaining exposure to variance decreases linearly with time.

17.8 Forward Variance

The additive feature of variance can be used to calculate the fair strike of a forward-starting variance swap.

Suppose we know the strike for a short-maturity variance swap expiring at time t (K_t) and the strike for a longer maturity variance swap expiring at time T (K_T), then we can easily calculate the fair strike of the forward-starting variance swap ($K_{t,T}$) with a combination of:

- Long $\frac{T}{T-t}$ variance notional of spot variance maturity T
- Short $\frac{t}{T-t}$ variance notional of spot variance maturity t, but with payment delayed until maturity T.

$$K_{t,T}^2 = \frac{T}{T-t} K_T^2 - \frac{t}{T-t} K_t^2$$

Two observations can be made from this formula:

- More variance is needed on the longer leg (tends to be less liquid) than the shorter leg.
- The total notional of the two legs will be greater than the notional of the forward.

Forward-starting variance swap can be useful for taking a direct view on the future value of implied variance or on the future shape of the variance term structure curve.

17.9 Contract Specifications

Variance swaps termsheets contain some interesting conditions that will be highlighted in this section.

Margins and collateral

Variance swaps are OTC products that are usually margined with an initial amount to be posted as collateral (ex: 3*vega). Further margin calls will be made during the life of the trade as necessary.

Disrupted days

Disrupted days are not considered as observation dates on which the realised variance is calculated. This can be impactful and work in both direction as shown by the below examples:

Example of a disrupted day decreasing the realised volatility: a 5% loss is followed by a 6% gain over two consecutive trading days. If the first day is declared as disrupted, this will result in a single combined return of 0.7%.

Example of a disrupted day increasing the realised volatility: a 5% loss is followed by a 6% loss over two consecutive trading days. If the first day is declared as disrupted, this will result in a single combined return of -10.7%.

Index reconstitution risk

Variance swaps on indices pay out on the returns on the index and not on the weighted returns of the basket of current constituents. As a consequence, they

are exposed to reconstitution risk as the index may end up with exposure to a different set of stocks with different volatilities!

Dividends Adjustments

Variance swaps on single names are typically adjusted for dividends, meaning that returns on ex-dividend dates are calculated after adjusting for the dividend.

Variance swaps on indices are typically not adjusted for dividends as they are more spread out and therefore have small impact compared to the average daily move.

Caps

Variance swaps are usually sold with caps with a standard cap being 2.5 times the strike.

Size

Typical sizes are €100k-200k vega notional for indices and €50k vega notional for single-stocks.

17.10 Observations from historical prices

- Variance swap prices tend to follow **high and low regimes** in a similar manner to realised variance.
- **Longer maturity variance swaps levels vary less and react less to spikes.** This is not very surprising as, for example, a sudden unexpected event is likely to increase dramatically short-term volatility but is less likely to cause the same level of high volatility over the next few years.
- **Single-stock variance swaps trade at higher levels than index variance swaps** (diversification effect).
- **Variance Swaps tend to trade above levels of realised volatility** for two reasons:
 1. **Volatility risk premium** due to risk aversion and hedging programmes
→ ATM implied volatility > realised volatility.
 2. **Convexity premium** since variance swaps are convex in volatility → variance swap strike > ATM implied volatility.

The theoretical price of variance swaps is calculated from prices of a replicating portfolio of options. With skew and skew convexity, the average volatilities will usually be above ATM implied volatility, making the variance swap more expensive. The price of a variance swap can be thought of as a function of ATM volatility level and the slope of the skew. In practice, contribution of skew

component means that variance swap strikes tend to trade at similar level to OTM puts. Skew and convexity become more important factors at longer dates as the probability of reaching more OTM strikes increases.

17.11 Variance Swaps and Option Implied Volatilities

Not surprisingly, variance swap strikes are well correlated with Black-Scholes implied volatilities derived from option prices since both can be interpreted as market estimates of future volatility.

These measures are similar but difference since:

- ATM IV reflects the market estimate of future volatility realised around the current level.
- Variance swap strike represents the market estimate of variance, independent of future market level.

17.12 Pricing Rules of Thumb

Theoretically, it is necessary to have prices available for the entire strip of options to calculate the true price of a VS.

However, under some assumptions about the skew, we can reasonably approximate variance swaps prices.

Under Flat Skew

Since under flat skew all strikes trade at an identical implied volatility, the variance swap level will be the constant implied volatility level.

Under Linear Skew

If skew is assumed to be linear, at least for strikes relatively close to the money, then Derman's approximation can be used.

Derman's approximation assumes a linear put skew and a flat call skew. It is a function of three variables, namely:

- ATM (forward) volatility
- Slope of the skew
- Maturity of the swap

$$K_{0,T} = \sigma_{ATMF} * \sqrt{1 + 3T * skew^2}$$

Where skew is generally taken to be the slope of 90/100 skew.

In practice, this approximation works best for:

- Short-dated variance swap: as maturity increases, OTM strikes have greater effect on the variance swap price and the contribution of skew is therefore more important. The inability of this approximation to account for skew convexity can make it less accurate.
- Index variance swap: single stocks tend to show more convexity, even at shorter dates.

Most often than not, this approximation tends to underestimate the variance swap price.

Under Log-Linear Skew

In reality, volatility skew is not linear across all option strikes and more accurate approximations can be used.

Assuming that the skew curve is log-linear of the form:

$$\sigma_K = \sigma_{ATMF} - \beta \ln\left(\frac{K}{F}\right)$$

Where

- F is the forward price
- K is the strike
- β is the slope of the skew

Then

$$K_{0,T} = \sqrt{\sigma_{ATMF}^2 + \beta \sigma_{ATMF}^3} = \frac{\beta^2}{4} (12\sigma_{ATMF}^2 T + 5\sigma_{ATMF}^4 T^2)$$

Gatheral's formula

Gatheral expresses the VS strike as an integral of the IVs across the entire range of strikes.

This formula characterises the skew curve in terms of the BS d2 parameter, which measures the 'moneyness' of the associated OTM option.

This leads to potentially powerful methods of variance swap approximation: by fitting a quadratic, or higher order polynomial to the skew surface parameterised in terms of d2, it is then possible to directly calculate a theoretical variance swap price from this parameterisation.

Example:

If skew curve is quadratic in variable $z = d_2 : \sigma^2(z) = \sigma^2 + \alpha z + \beta z^2$ then theoretical variance swap strike is $K_{0,T} = \sigma_0^2 T + \beta T$

- in ‘d2-space’ the linear component of the skew, α , has no effect on the variance swap price.
- base level of volatility σ_0 affects the swap strike.
- the convexity β affects the swap strike.

17.13 Drivers of Variance Swaps levels

Variance swap levels and options implied volatilities move in closest step with each other. They are driven mostly by the same factors although variance swaps are also driven by the shape of the skew curve.

Historical realised volatility

It is one of the most important drivers of variance swap levels at least at short maturities since the correlation between the 1-month VS and the realised volatility over the previous month is about 0.91. As maturity increases, this correlation between variance swap levels and realised volatility naturally decreases.

Risk-aversion

Risk-aversion expresses itself in the demand for protection in the form of put options. The more risk-aversion, the higher the skew and the skew convexity and the more expensive the variance swap prices.

Market level

Volatility itself is directional, at least over short time frames, with volatility tending to increase if market sells off.

Additional info

Variance swaps flows can have a feedback effect into the market as they account for an important part of the traded vega demand. Their hedging by market makers can therefore influence implied volatilities and skew.

17.14 Are variance swaps good predictor of future volatility?

In theory, variance swap strike represents the market prediction of realised variance over the term of the swap.

In practice, this is complicated by supply/demand factors, especially at longer maturities where driven by structured products flows. Also, at shorter maturities, risk-aversion can bias the variance swap strike to be above the fair expectation of future realised variance.

Anyway, at short dates, implied variance tends to do a relatively good job of forecasting future realised variance. Backtesting shows that it does a better job than previous realised variance.

Given historical data, a simple model of future realised volatility as 0.9 times the variance swap strike provides the best fit with the data.

17.15 Is Variance Swap Convexity fairly priced?

How can we quantify the volatility of volatility priced in the variance swap price?

Well, we can calculate the standard error of the future volatility as estimated by the variance swap strike. The average error is giving information about the historical variability of future realised variance around the estimate provided by 0.9 times the variance swap strike. Backtesting shows that this standard error is about 6%. In other words, the standard deviation of implied variance minus future realised variance is 6%.

We can also calculate an implied variability of volatility by considering spread between short-dated ATM implied volatilities and the variance swaps strikes. This spread for the SX5E suggests an average implied volatility of volatility of approximately 7%. This is fairly close to the 6% average error of the variance swap in predicting volatility.

This 6% vol of vol for SX5R is a long-term average. Like volatility, it changes over time with a min of 2% and a maximum of 10% (historically).

Also, there is a strong correlation between volatility and the spread between variance swap levels and ATM volatility. It makes sense if we believe that volatility of volatility is correlated with volatility.

17.16 Variance Term Structure

The shape and potential movement of the implied variance curve is important in determining the most promising parts of the curve to buy or sell variance.

Variance Swap term structures:

- are usually upwards sloping

- have tendency to flatten following increases in volatility
- generally steeper than ATM vol curves as there is an increasing effect of skew at longer maturities.
- can be thought of representing the mean-reverting nature of volatility.
- Short end is most sensitive to prevailing levels of realised volatility.
- Long end is more anchored to some LT estimate of average volatility.
- Long end is also driven by structured product flows and tends to be more susceptible to supply/demand dynamics. Other important point is the effect of the mark-to-market P&L.
- A short-date variance swap is principally exposed to realised variance $\rightarrow$ gamma.
- A long-dated variance swap takes on significant exposure to changes in implied variance before expiry $\rightarrow$ vega.

An investor has just bought a 5y variance swap. His principal exposure is the 5y implied variance, which is driven by factors not necessarily correlated with current realised variance. His P&L can be quite unpredictable.

Typical movements of term structures can be explained in part by the “root-time” rule. For “normal” move in volatility, the change in implied variance at a given point on the curve will be proportional to $\frac{1}{\sqrt{T}}$.

- Flattens normal upward-sloping term structure as volatility increases.
- Steepens normal upward-sloping term structure as volatility decreases.
- Changes will be most visible at shorter end of the curve.

Example: if 1y variance increases by 1% $\rightarrow$ 3m variance increases by 2% and 4y variances increases by 0.5%.

17.17 From Options to Variance Swaps

To create a portfolio of options with constant exposure to variance, you need to ensure that $\$ \Gamma$ is constant for moves both in S and t .

The cost of this portfolio represents the price of exposure to realised variance.

When considering how to construct such an exposure, there are **three possible approaches** to take:

1. Use a single vanilla option, but buy/sell additional amounts of it to keep $\$ \Gamma$ constant over time.

Advantages: use only a single option strike

Disadvantages:

- option needs to be dynamically traded.
- position could end up with enormous amounts of option as gamma decreases.

2. Re-strike the option to maintain a constant gamma.

Start with an ATM option and on each re-hedging step, sell/hold option and buy new ATM to achieve constant gamma.

Better than the first approach but still requires dynamic trading of options.

3. Construct a portfolio of options so that dollar Gamma ($\$ \Gamma$) remains constant over both moves in S and t .

Advantages:

- no dynamic trading of options although dynamic delta-hedging.
- to some extent independent of the volatility process driving the underlying.

Disadvantage: it requires a strip of options across a continuum of strikes.

In theory, the third approach is used.

In practice, it is more like a combination of the 2nd and 3rd approach.

The third approach is to some extent independent of the volatility process driving the underlying.

The first two approaches would require continually calculating the gamma over the course of the trade, and this gamma will be highly (and dangerously) dependent on the assumed IV (and volatility process) at that time.

So what kind of portfolio is needed to achieve a constant dollar gamma across strikes?

Let us have a look at a few graphics to answer this question.

Fig. 16.1:

- peak dollar gamma increases linearly with strike.
- contribution of low-strike options is small compared to high-strike options.
- we need to increase weights of low-strike options and decrease weights of high-strike options.




Figure 17.1: Fig: 16.1 : Peak dollar gamma of an option increases linearly with underlying




Fig 16.2:

- Naively, it may be thought that weighting by $\frac{1}{K}$ will achieve constant dollar gamma.
- It has the property that each option in the portfolio has an equal peak dollar gamma.
- However, dollar gammas of higher strike options spread out more.




Fig. 16.3:

- summing $\frac{1}{K}$ - weighted options across all strikes leads to dollar gamma exposure increasing linearly with S.
- weighting each option by $\frac{1}{K^2}$ will achieve constant dollar gamma.

Why a portfolio of options weighted by $1/K^2$ gives a constant exposure to volatility?

Answering this question requires a bit of math but we will stick to the basics as always.

Delta-neutral P&L = $\frac{1}{2}\Gamma S^2((\frac{\Delta S}{S})^2 - \sigma^2\Delta t) \rightarrow$ only ΓS^2 prevents the direct exposure to variance.

We have to find a portfolio with gamma proportional to $\frac{1}{S^2}$ to have a constant exposure. Mathematically speaking, we have to find a portfolio whose second derivative with respect to the underlying is proportional to $\frac{1}{S^2}$.

The negative natural log of S represents such a payoff: $\frac{\delta^2(-\ln S)}{\delta S^2} = \frac{1}{S^2}$

Unfortunately, log contracts are not traded in the market.

But we can replicate such a contract with vanilla options using an infinite sum of calls and puts across the continuum of strikes, each weighted by the inverse square of strike.

Integrating the value of this portfolio at expiry demonstrates that the non-linear part of the payoff is a negative log contract.

$$K_{VAR}^2 = \frac{e^{rT}}{T} \left[\int_0^{F_0} \frac{P_0(K)}{K^2} dK + \int_{F_0}^{\infty} \frac{C_0(K)}{K^2} dK \right]$$

The strike-squared, K_{VAR}^2 , can be thought of as the (future-valued) inverse strike-squared weighted sum of the time values of the options portfolio.

To summarise:

- A portfolio of calls and puts, weighted as 1/strike-squared, has constant dollar-gamma;
- Delta-hedging this portfolio provides constant exposure to the difference between implied and realised variance regardless of where the volatility is delivered;
- Hence the P&L from delta-hedging this portfolio is proportional to difference between realised and implied variance.

A variance swap can therefore be created by replicating a log contract with options which are then delta-hedged.

17.18 Sensitivity to Skew and Convexity

Skew is commonly thought of as an important component of variance swap prices, with put skews seen as having a much greater impact on prices than call skews.

Although, in practice, this is a useful framework for thinking about how variance swap prices behave, it is not theoretically correct as shown with the below example.

Case 1:

- 3-month ATM IV @ 20%
- Linear put skew of 5% (per 10 volatility points) with put IVs capped @ 35% and all OTM call volatilities flat at the level of the ATM IV.

The theoretical 3-month variance strike can then be calculated to be 23.05.

Case 2: Situation where the skew is a mirror image of Case 1.

- 3-month ATM IV @ 20%
- All OTM put volatilities are flat @ 20%, but call IVs increase linearly by 5% per 10 points as they become more OTM, capped @ 35%.

In this case the theoretical variance swap price is virtually identical at 23.15.

So the exposure to the skew curve is symmetrical.

Why does this symmetry exist?

Since the $\frac{1}{K^2}$ replicating portfolio has a much higher weighting of puts, it would be natural to assume that their associated implied volatilities should have a proportionally greater effect on the variance swap strike.

Also, in practice, variance swaps can effectively be priced and hedged with ATM volatility plus a contribution from the skew.

But, as we have seen, the exposure to the skew curve is symmetrical. That is, the contribution to the variance swap price of the volatility of an OTM call is exactly the same as from an OTM put with same (risk-neutral) probability, $N(d2)$, of ending ITM. The OTM puts have a greater weighting in the replicating portfolio, simply because their dollar gamma is lower and so they must be scaled up in order to provide constant dollar gamma across the range of strikes.

So what exactly determines the contribution of volatilities across the skew curve to the variance swap price?

Clearly a very OTM option should make a relatively small contribution to the variance swap price. It then seems sensible that a variance swap represents a

kind of weighted average of volatilities across the skew curve, with the closer-to-the-money volatilities higher weighted. In fact, this is exactly the case, with the average being taken over the variances rather than the volatilities, and the weighting function simply being the risk-neutral probability density function, $N(d_2)$, that the corresponding OTM option ends up ITM.

If we define the variable z to be the standard Black-Scholes parameter d_2 : $d_2 = \frac{\ln(\frac{S}{K}) + (r - \frac{\sigma^2}{2})T}{\sigma\sqrt{T}}$ then it can be shown that $K_{VAR}^2 = \int_{-\infty}^{\infty} N'(z) \sigma_{BS}^2 dz$

The $N(d_2)$ term is **the probability density function** for the underlying at expiry, T .

That is, **the cumulative distribution $N(d_2)$** gives the (risk-neutral) probability that the underlying will be trading above z at time T . Thus, the parameter z , simply represents the ‘moneyness’ of the corresponding OTM option. This means that the variance swap price is a weighted sum of squared option implied volatilities weighted by the probability that the (OTM) option will end ITM.

Furthermore, if the skew curve is quadratic in the variable z (the moneyness of the option) of the form: $\sigma_{BS}^2(z) = \sigma_0^2 + \alpha z + \beta z^2$ then substituting and integrating gives $K_{VAR} = \sigma_0^2 T + \beta T$, i.e. in ‘ d_2 -space’ the variance swap price is not affected by the linear component of the skew α , but only on the base level of volatility σ_0 , and the convexity parameter β .

This explains why very different (linear) skews, but similar convexities give almost identical variance swap strikes. The reason that there is any difference at all in variance swap strikes of our example, is because in d_2 -space, the convexity in case 2 is slightly greater, since the maximum volatility is at a strike of 130, which is slightly less OTM than the 70 strike where the maximum volatility is achieved in case 1.

This also helps to explain why the convexity has a greater effect on longer maturity variance.

As maturity increases, the probability of far OTM (in relative strike terms) options ending ITM increases. Therefore, the relative weight of e.g. the 160-strike call will increase with maturity.

At shorter maturities perhaps only the 80-120 portion of the skew surface significantly affects the variance swap price, and this part of the skew is relatively linear (at least for indices).

At 5 years, strikes out to say 50-180 become relevant to the pricing and the convexity of these strikes can be much higher.

In practice, skew is most often thought about in terms of relative strikes, rather than moneyness – but the point remains: **the contribution of a point on the skew curve to the variance swap price depends on the (risk-neutral) probability $N(d_2)$ of the associated OTM option ending ITM.**

Therefore, ATM volatility will provide the greatest contribution to variance swap prices – particularly for short maturities. Both high put skews, and high call skews (where OTM calls have higher volatilities than ATM) will increase variance swap prices.

In fact, there is some kind of **feedback effect** here because the contribution of each volatility is determined by the probability of an option being ITM, but, for example high put skews will increase the probability of OTM puts ending ITM so both the volatility and the weighting factor will increase.

To sum up:

- Given a flat skew, variance will price at the same level as ATM volatility.
- Positive convexity will always act to increase variance swap strikes.
- With a negatively convex skew – OTM volatilities are (on average) less than ATM volatility – it is theoretically possible that variance could price below ATM volatility.

17.19 Greeks

In this section we consider the Greeks for variance swaps, giving information about the sensitivity of variance swaps to various market variables. We work directly by differentiating the mark-to-market value of the variance swap contract.

The value of a variance swap, per unit vega-notional, at time t is given by:

$$P_t = \frac{1}{2K_0} [\sigma_{Expected,t}^2 - K_0^2]$$

Note that intra-day there is a term representing the square of the move which will act to give the variance swap delta on an intra-day basis:

$$\sigma_{Expected,t}^2 = \frac{t-1}{T} \sigma_{0,t-1}^2 + \frac{252}{T} [\ln(\frac{S_t}{S_{t-1}})]^2 + \frac{T-t}{T} K_{t,T}^2$$

where the three terms are written in chronological order:

- The first term is the realised variance accrued from inception to day $t-1$.
- The second term is the realized variance based on the current daily return at the valuation time on day t .
- The last term is the implied variance expressed as the strike of a variance swap on day t expiring at T squared.

The Greeks of the variance swap can then be calculated by differentiating P_t .

17.19.1 Gamma

Gamma comes only from the exposure to realised volatility on each day:

$$\Gamma = \frac{\partial^2 P}{\partial S_t^2} = \frac{252}{K_0 T} \left[\frac{1}{S_t^2} \right]$$

Since the dollar gamma is achieved by scaling the gamma by the spot squared, this gives a constant dollar gamma as expected.

17.19.2 Theta

$$\theta = \frac{\partial P}{\partial t} = \frac{-K_{t,T}^2}{2TK_0}$$

In particular, if the variance strike does not change, theta remains constant. Note, that since T is measured in days this value represents a daily theta.

You simply have to multiply by 252 to annualize it:

$$\theta = \frac{\partial P}{\partial t} = 252 \frac{-K_{t,T}^2}{2TK_0}$$

which can then be shown to satisfy the formula $\theta = -\frac{1}{2}\Gamma S^2 \sigma^2$ taking σ^2 as the implied variance from the variance swap strike.

17.19.3 Vega

We can calculate exposure to volatility in terms of sensitivity to changes in:

1. ATM volatility
2. Variance strike
3. Implied variance (strike squared)
4. To compute the vega in terms of sensitivity to ATM volatility, we must make some assumptions about how the variance swap strike relates to ATM volatility.

If we assume the Derman approximation: $K_{t,T}^2 = \sigma_{ATMF}^2 * [1 + 3(T-t)skew^2]$, then:

$$\frac{\partial P}{\partial \sigma_{ATMF}} = \frac{1}{2K_0} \frac{T-t}{T} \frac{\partial \left(\sigma_{ATMF}^2 [1 + 3(T-t)skew^2] \right)}{\partial \sigma_{ATMF}} = \frac{\partial \sigma_{ATMF}}{K_0} \frac{T-t}{T} [1 + 3(T-t)skew^2]$$

This approximation also allows the calculation of sensitivity to the skew or the skew squared:

$$\frac{\partial P}{\partial skew} = \frac{3}{K_0} \frac{(T-t)^2}{T} \sigma_{ATMF}^2 skew$$

Increasing skew will increase the value of the variance swap, and do so by more, if there is more time remaining until expiry. 2. / 3. Computing sensitivities to the strike (or strike squared) is more straightforward and needs no such assumptions about the skew surface:

$$\frac{\partial P}{\partial K_{t,T}} = \frac{T-t}{T} \frac{K_{t,T}}{K_0}$$

$$\frac{\partial P}{\partial K_{t,T}^2} = \frac{1}{2K_0} \frac{T-t}{T}$$

These all tell us that the exposure to implied variance (or volatility) decreases with time as the accrued realised volatility is locked in to the P&L.

17.19.4 Delta

Firstly, assuming that the variance strike K has no sensitivity to the underlying, the variance swap can be seen to take on delta only intra-day:

$$\Delta = \frac{\partial P}{\partial S_t} = \frac{T-t}{T} \frac{252}{K_0} \left[2 \ln \left(\frac{S_t}{S_{t-1}} \right) \frac{1}{S_t} \right]$$

This represents the replication of the log contract which will have to be done at the end of the day to capture that day's realised variance.

If the variance strike itself has a dependency on the underlying (implied variance is directional) then the variance swap acquires other sources of delta, in addition to the intra-day delta.

For example, using Derman's approximation, calculating delta (only on the close to avoid the intra-day delta discussed above) gives:

$$\Delta = \frac{\partial P}{\partial S_t} = \frac{T-t}{T} \frac{1}{2K_0} \frac{\partial K_{t,T}^2}{\partial S_t} = \frac{T-t}{T} \frac{1}{2K_0} \frac{\partial \left(\sigma_{ATMF}^2 (1 + 3(T-t)skew^2) \right)}{\partial S_t} = \frac{T-t}{T} \frac{1}{2K_0} 2\sigma_{ATMF} \frac{\partial \sigma_{ATMF}}{\partial S_t} (1 + 3(T-t)skew^2)$$

Thus with a 'normally shaped' negatively sloping skew, the delta of the variance swap will be negative, at least if the skew curve is sticky with strike.

This fits with the intuition that IV will tend to go down as the underlying rallies – which is exactly what a (negative) linear skew represents.

17.19.5 Convexity

As explained before, it is really convexity and not skew which acts to increase the variance swap strike above ATM volatility. Clearly then, rising convexity will increase variance swap strikes as the ‘average’ implied volatility used in Gatheral’s formula goes up.

17.20 Setting up a replicating portfolio

We previously saw that a variance swap could be statically replicated by a portfolio of OTM options, weighted according to $\frac{1}{K^2}$.

How do we weight this portfolio to achieve a specific variance-notional or vega-notional exposure?

If we wish to replicate a variance swap for €100K of vega notional, how many of the calls and puts do we need in the replicating portfolio?

Let Π be forward of the price of the portfolio of OTM options: $P = e^{rT} \left[\int_0^{F_0} \frac{P_0(K)}{K^2} dK + \int_{F_0}^{\infty} \frac{C_0(K)}{K^2} dK \right]$

This portfolio associated with the dynamic futures positions will pay $\frac{T}{2} \sigma_R^2$ at expiry.

Scaling this portfolio by $2/T$ will therefore produce a payout of RV squared but in % terms, so we have to multiply by 100^2 to be consistent with the standard quoted variance swap strikes.

The cost of this portfolio represents the fair cost of future realised variance.

$$K_{VAR}^2 = \frac{2 \cdot 100^2 \cdot \Pi}{T}$$

To be long N_{VAR} of variance notional, you must buy $\frac{2 \cdot 100^2 \cdot N_{VAR}}{T}$ lots of portfolio P, which is equivalent to buying $\frac{2 \cdot 100^2 \cdot N_{VAR}}{T K^2}$ of each OTM K-strike option.

The total cost of this portfolio will then be $N_{VAR} * K^2$. The value of this portfolio and its associated delta hedging at expiry will be $N_{VAR} * \sigma^2$ giving the desired overall P&L of $N_{VAR} * (\sigma^2 - K^2)$.

At inception if the cut-off between puts and calls is taken to be the current forward price of the underlying, no net delta- hedge is needed.

However, during the lifetime of the variance swap, at the end of each day’s close the notional value of the delta hedge is adjusted so that it equals:

$$\frac{2 \cdot 100^2 \cdot N_{VAR}}{T} \left(\frac{F_0 - F_t}{F_0} \right)$$

where F_0 and F_t were respectively the original forward price at time 0 and the current forward time.

To achieve a specific vega notional, $N_{VAR} = \frac{N_{Vega}}{2K_{VAR}}$, we must trade $\frac{100^2 * N_{Vega}}{T * K_{VAR}}$ lots of the portfolio P.

Since $K_{VAR} = 100 * \sqrt{\frac{2\Pi}{T}}$, the required amount of the portfolio is $\frac{100 * N_{Vega}}{\sqrt{2\Pi T}}$ or equivalently $\frac{100 * N_{Vega}}{K^2 \sqrt{2\Pi T}}$ of each OTM option.

In practice, it is obviously not possible to trade a continuum of option strikes, so we will briefly see how to approximately construct a variance swap from a tradable strip of OTM options. For simplicity, we assume the strikes are equally spaced.

To get a variance notional exposure N_{VAR} , we will need $\frac{2 * 100^2 * \Delta_K * N_{VAR}}{TK^2}$ of each option.

Be careful with the contract size of the option when using that formula. If one option represents 10 underlying shares, then you will need to divide by 10 the amount of options you would otherwise require.

Also note that this kind of approximation will tend to slightly overvalue the true theoretical variance swap strike due to convexity issues.

17.21 Replicating and Hedging in Practice

In theory, a variance swap can be statically hedged with a portfolio of OTM European options, weighted according to the inverse squares of their strikes. Assuming option prices are available across the entire range of strikes, it makes it easy to value it.

In practice, traded strikes are not continuous, although for major liquid indices they are closely spaced. A more serious limitation is the lack of liquidity in OTM strikes, especially for puts, as these provide a relatively large component of the variance swap price in the presence of steep put skews.

When replicating the variance swap:

- The long futures position is used to create a pay-out which is equivalent to a long log contract plus realised variance.
- The long options/short forward position is used to create a short log contract and pay the fixed strike.

Supposing the market falls significantly, the delta-hedge will be long the log contract, while the options should counteract this by being short the log contract. However, if not enough downside puts were used, the options portfolio will not fully reflect the short log-contract exposure needed and hence the overall hedge will lose money.

This lack of liquidity at the wings has led to the development of conditional variance swaps which can remove exposure to volatility once the underlying moves into areas where vanilla options are illiquid. **In practice**, market-makers will not attempt to hedge with the entire strip of options but typically will use only two or three – including one close to the money and one or more OTM, but liquid, puts.

Alternatively they could approach the replicating by hedging the vega with an OTM put whose implied volatility coincides with the variance swap strike – close to the money for 1-month maturity, 95% for 3-6 months, 90% for 1-year, 85% for 2-3 years etc. In this case they would also look to buy back the wings/convexity separately.

One problem with this kind of approach is that the partial hedge is no longer static, and must be dynamically managed. For example if the market sells off towards the strike, the market maker will have to trade further OTM puts to ensure that their exposure to volatility, in the form of dollar-gamma, remains constant.

This makes the actual variance swap replication more akin to a combination of both alternatives. Here the constant dollar gamma would be maintained by a combination of holding a portfolio which has roughly constant dollar gamma if the underlying does not move too much, and re-hedging by trading more options if the underlying does move significantly.

17.22 Effects of Variance Swaps Hedging

Market-makers who trade variance swaps may hedge their positions by replicating the opposite variance swap position through the replicating options portfolio. This replicating portfolio then needs to be delta hedged.

The effects of delta-hedging this portfolio are different to that from normal delta-hedged options for two principal reasons:

First, the actions of delta-hedging the options could potentially act to the disadvantage of the counterparty's position.


If the market-maker has sold options to a counterparty, his delta hedging will have the effect of increasing volatility in the underlying: magnifying both up-moves and down-moves.

Similarly, if the market-maker has bought options from a counterparty, his delta hedging will have the effect of suppressing volatility in the underlying, potentially to the advantage of the counterparty who is short the option.

The situation with variance swaps is different.

Suppose that the market is such that market participants have generally sold index variance swaps to market-makers. Note that no exchange of options has

taken place here – the parties have just taken opposite sides in a contract for difference. Suppose that the volatility sellers do not hedge their variance swaps (they have sold the variance swaps specifically for the direct volatility exposure they offer). But assume that market-makers hedge their short volatility exposure. A market-maker who is long the variance swap can offset the risk by shorting the replicating portfolio of options, and delta-hedging. They will therefore be short gamma in the options market.



As described, a delta-hedger who is short options will act to increase volatility in the underlying – buying as it rallies, and selling as it sells-off. However, the action of these market makers hedging their short options will not necessarily act to increase volatility in the underlying, as the counterparties they have sold options to may be counteracting this effect by themselves hedging their long volatility positions.

Second, since variance swap contracts typically measure close-close realised volatility, the options must be delta-hedged on the close only to capture this.

The important difference between the two groups of hedgers is that the variance swap market-makers who are short options, must hedge only on the close to capture the close-close realised variance specified in the variance swap contract. In contrast, the hedgers who are long the options, will generally be free to choose when to delta-hedge, as they attempt to capture the true volatility of the underlying process.

Therefore, the overall effect of hedging these variance swaps need not have the effect of increasing overall market volatility, although it may if the long options positions are not being hedged. However, the important point is that the hedging of long variance swap positions may act to increase close-to-close volatility, with option hedges on the close having the potential effect of magnifying daily moves.

17.23 Why not Volatility Swaps?

As pointed out, the exposure of delta-hedged options to volatility, after accounting for the gamma, is actually an exposure to the difference between implied and realised volatility squared. In this sense, a variance swap mirrors a kind of ideal delta-hedged option whose gamma remains constant. Furthermore, variance swaps are relatively easy to replicate. Once the replicating portfolio of options has been put in place, only delta-hedging is required. No further buying or selling of options is necessary.

All this explains why variance swaps are attractive instruments to trade, but still does not explain why volatility swaps are not also frequently traded.

The main theoretical difficulty with volatility swaps is that they cannot be statically replicated through options. A replicating portfolio must dynamically trade options, and make relatively strong assumptions about the underlying volatility process – in particular about the volatility of volatility. This makes any replication process model dependent, and therefore much more prone to errors than the theoretically robust variance swap replication.

In fact the convention of quoting variance swap notionals in vega, rather than variance, amounts can be seen as an attempt to treat variance swaps like volatility swaps. The vega notional represents the average P&L for a 1 vega change in volatility, but with the convexity meaning that longs will profit by more if volatility increases, and lose by less if volatility decreases. Thus for small changes in volatility, where the effect of the variance swap convexity is relatively limited, variance swaps (measured in vega notional) locally approximate volatility swaps.

Volatility swaps can then be thought of as variance swaps without the convexity. The discount of the volatility swaps to variance swaps should therefore reflect the value of this convexity, which in turn is determined by the volatility of volatility. Seen in this light, the ability to calculate a fair price for volatility swaps, requires not only a price for the volatility (variance), but also a price for the volatility of volatility – i.e. the means of valuing options on volatility.

Since it is variance which arises naturally from delta-hedging options, in terms of volatility products, variance should be thought of as the true underlying. In this framework, volatility swaps are naturally thought of as derivatives of variance – paying the square-root of the variance swap contract.

In fact we could dynamically trade a long variance swap (buying more as volatility decreases and selling as volatility increases) to hedge out this convexity bias.

Similar to delta-hedging an option the P&L made from the resulting buy low – sell high strategy (for variance) will lead to a P&L based on the volatility of volatility: the larger this volatility-of-volatility, the bigger the discount of the dynamically replicated volatility swap to the variance swap.

However, besides issues arising from the transaction costs of dynamically trading variance swaps, estimating this volatility-of-volatility contribution from the dynamic hedging of the variance swap is problematic and model-dependent, making the volatility swap contracts difficult to price.

17.24 Third Generation Products

Those products include gamma swaps, corridor swaps and conditional variance swaps.

17.24.1 Gamma Swaps

Gamma swaps are very similar to variance swaps but maintain a dollar gamma linear with spot, rather than a constant dollar gamma, leading to lower a skew exposure and a slightly lower strike compared to standard variance swaps.

17.24.2 Corridor Variance Swaps

The original corridor variance swaps only accrue variance within a pre-specified range, meaning a long would make a maximal loss when the underlying fails to trade within the range.

17.24.3 Conditional Variance Swaps

Conditional variance swaps have been the most popular of these products, allowing investors to take exposure to realised variance contingent upon the underlying instrument trading within a pre-specified range. Outside this range no P&L accrues. Conditional variance swaps can be useful for hedging complex volatility exposures, taking a view on the volatility levels encapsulated in the skew or buying/selling variance at more attractive levels given a view on the underlying.

As with variance swaps, the sign of the P&L for a conditional variance swap is controlled by the difference between the volatility realised (in the range) over the lifetime of the swap and the pre-agreed fixed strike level. In the case of the conditional, P&L is accrued only when the underlying is within a pre-specified range. In addition, the magnitude of the final P&L is scaled by the proportion of time the underlying has spent in the pre-specified range. If the underlying trades entirely outside the range, the P&L will be zero, so the maximal loss for a long will be when the underlying trades within the range but with very low volatility.

Whilst investors are free to specify the range associated with a conditional variance swap, the two principal types are up- and down- conditional variance swaps (up-variance and down-variance). Up-variance accrues realised volatility only when the underlying is above a pre-specified level (i.e. no upper barrier), while down-variance is accrued only when the underlying is below the specified barrier (i.e. no lower barrier).

In the presence of a positive put-skew, down-variance will normally price above up-variance for close to ATM barrier levels.

Conditional variance swaps can be useful for expressing views on volatility contingent on market level.

For example investors seeking crash protection may purchase conditional down-variance, which only becomes activated in the event of a market sell-off. That is, if the market stays above the down-barrier, P&L is zero.

Conversely investors who believe volatility will be realised on the upside can buy conditional up-variance swaps, which are usually cheaper than a standard variance swap for which a significant amount of the premium is used to fund the downside skew exposure.

Conditional variance swaps have typically been traded on index underlyings. More recently, single-stock conditionals have gained liquidity, allowing investors to trade the volatility of a stock, contingent on stock price.