KruegerAlgorithms

Foundations · market mechanics from the ground up

Futures, options and who makes the price

Chapters 1 to 10 of 10

One connected explanation from the ground up: where futures and options come from and what they are for, who is in the market, how a price forms in the order book, what exactly a future and an option are, how option dealers offset their risk and what open interest and GEX say about it. In plain language, with every central formula derived step by step and worked through with numbers.

Field	Value
Series	Futures, options and who makes the price
Scope	Chapters 1 to 10 of 10
Type	Explanation with worked examples, not a measurement
As of	11 November 2026
Author	Timon Krüger
Online	kruegeralgorithms.com/en/research/futures-options-market-structure-guide

An explanation of market mechanics with worked examples. Not investment advice and not a trading recommendation. Trading futures and options involves the risk of loss up to and beyond the capital employed.

Contents

1 Why futures and options exist: history, participants and the common thread **5**

1. The basic problem: a price that changes before delivery 5
2. From the rice bill to the future 5
3. From grain to the index 6
4. The option: the right without the obligation 6
5. From the floor to the computer: who quotes the prices? 7
6. Who is in the market today 8
7. The common thread 8

2 The market as an auction: order book, market makers and market takers **11**

1. An exchange is an auction that never ends 11
2. Reading the order book 11
3. Two kinds of orders 12
4. Makers and takers 12
5. What a market maker earns from 13
6. What a market maker loses on 13
7. Derivation: how wide does the spread have to be? 14
8. From the floor to the computer 15

3 Liquidity: depth, price impact and when the resting orders disappear **16**

1. What liquidity is: three measures 16
2. A large order runs through the book 16
3. The square-root rule of thumb 17
4. When liquidity disappears 18
5. Auctions: one price for all 19
6. VWAP and TWAP: the yardstick of execution 19

4 Futures: what they are, who uses them for what and why their price is tied to the index **22**

1. What a future is 22
2. Who uses futures and for what 23
3. Margin, leverage and daily settlement 23
4. The fair price: derivation from arbitrage 24
5. Basis and expiry 25
6. The roll 26
7. How index CFDs are tied to it 27

5 Options: call, put, payoff, time value and put-call parity **29**

1. What an option is 29
2. The payoff at expiry 29
3. The four basic positions 30
4. Hedging with puts instead of futures 31

5. Intrinsic value and time value 31
6. Put-call parity: derivation 32
7. Synthetic positions 34

6 The option price: from a coin toss to the Black-Scholes formula 36
1. The obvious calculation, and why it is wrong 36
2. The replication: a portfolio that pays the same 36
3. Why it has to be this price 37
4. The general formula 37
5. Two steps: the replication has to be adjusted 38
6. Many steps 39
7. The Black-Scholes formula 40
8. Implied volatility 41
9. The skew: insurance costs extra 41
10. The volatility premium: implied lies above realised 42

7 The Greeks: delta, gamma, theta, vega, vanna and charm 44
1. What the sensitivities are for 44
2. Delta: how closely the option follows the index 45
3. Gamma: how fast the delta changes 45
4. Theta: what a day costs 46
5. Gamma and theta belong together 47
6. Vega: what a volatility point is worth 47
7. Vanna: how volatility shifts the delta 48
8. Charm: how time shifts the delta 48
9. All Greeks at a glance 49

8 Hedging in the underlying: why option dealers with long gamma trade against the move and with short gamma trade with it 51
1. What a dealer wants: the spread, not the direction 51
2. The delta hedge, step by step 52
3. The price moves: the hedge no longer fits 52
4. What the move brings in: ½ · gamma · move² 52
5. What waiting costs: theta 53
6. Five days worked through 54
7. The other side: short gamma 55
8. From the single book to the market 56

9 Dealer positioning: open interest, GEX, gamma flip and walls 58
1. The question 58
2. Open interest and volume 58
3. Who has which side? 59
4. GEX derived 59

5. Worked example: a small option chain 59
6. The profile and the gamma flip 60
7. What walls are supposed to mean 61
8. 0DTE: the gamma that is missing from open interest 61
9. Expiry and pinning 62
10. Vanna and charm flows 62
11. The assumptions everything depends on 63

10 What can be measured: the mechanics against our own tests 65

1. Three stages that have to be kept apart 65
2. How we measure 65
3. The overview 66
4. What is confirmed 66
5. What remains mechanics 66
6. What failed as a signal 67
7. The pattern behind it 68
8. What can be done with it 68

Chapter 1

Why futures and options exist: history, participants and the common thread

In one sentence: Futures and options exist because some want to get rid of a price risk and others take it on for payment, and almost everything that follows in this series is the question of how those who take it on get rid of that risk again themselves.

What this chapter is: the entry into the series. It explains the origin and purpose of the products, introduces the participants and shows how the following chapters connect. The history is told in shortened form, the dates come from the chronicles of the exchanges. **Not a trading recommendation.**

1. The basic problem: a price that changes before delivery

A farmer sows wheat in spring. He can only sell in autumn. What the wheat will cost then, he does not know. If the price falls, the harvest may not cover his costs.

A miller has the opposite problem. He has promised bread deliveries at fixed prices and has to buy wheat in autumn. If the price rises, he loses money.

Both carry the same risk with the opposite sign. The obvious solution: they agree today on a price for delivery in autumn. The farmer knows what he will get, the miller what he will pay. Such a contract is called a **forward transaction**: concluded today, fulfilled later.

Three questions follow from this, and they drive the whole history of these markets:

- **Who is the other side when no suitable miller happens to be there?** Someone has to be willing to take on the risk without needing wheat himself.
- **What if one side does not pay?** If the price falls sharply, the miller has an incentive to break the contract.
- **How do you get out of the contract again?** A contract between two persons is hard to sell on.

The future is the answer to all three.

2. From the rice bill to the future

Osaka, 1730. In Japan, taxes were paid in rice. Warehouses in Osaka issued bills on stored rice, later also on rice from future harvests. These bills were traded. In 1730 the government officially permitted trading at the Dojima rice exchange. What was traded there was not only rice for money but also rice for future delivery, in uniform quantities and with a clearing office. Dojima is therefore regarded as the first futures exchange.

Chicago, 1848 and 1865. Chicago was the hub for the grain of the American Midwest. After the harvest prices fell because everyone delivered at once, in spring they rose. In 1848 merchants founded the Chicago Board of Trade. At first they traded delivery contracts on grain arriving later. In 1865 the exchange turned these into uniform contracts with fixed rules: the same quantity, the same quality, fixed delivery months, and a security deposit that both sides had to post. That is the future in its present form.

Three innovations are contained in it:

- **Standardisation.** Every contract is the same. That is why it can be sold on to anyone.
- **Security deposit (margin).** Both sides post money. Gains and losses are settled continuously. Nobody can let a large loss build up and then disappear.
- **Clearing house.** A central body steps between buyer and seller. Everyone has only it as a counterparty. If a participant fails, it does not hit his trading partner.

This answers the three questions from section 1. The other side no longer has to be a miller but can be anyone who wants to trade at that price. The clearing house carries the default risk. And you get out by selling the same contract again.

3. From grain to the index

For more than a hundred years futures were written on things that can be delivered: wheat, corn, pork bellies, metals. Then, within ten years, two steps created today's market.

1972: futures on currencies. Until 1971 the exchange rates of the major currencies were fixed. When that system ended, they floated freely, and firms with foreign business suddenly had the farmer's problem. The Chicago Mercantile Exchange (CME) introduced futures on seven currencies in May 1972. For the first time a future was written on something purely financial.

1982: futures on a stock index. An index cannot be delivered. It is a number computed from 500 share prices. The solution was **cash settlement**: nothing is delivered at the end, only the difference is settled in money. In April 1982 the CME launched the future on the S&P 500. Anyone who wanted to hedge a share portfolio no longer had to sell hundreds of shares. A single trade was enough.

1992 and 1997: trading becomes electronic and smaller. Until then traders stood in a ring on the exchange floor and shouted prices at each other. In 1992 the CME put its electronic trading system Globex into operation. In 1997 the E-mini on the S&P 500 followed: one fifth the size of the original contract and tradable only on screen. It became one of the most heavily traded futures in the world. Since 2019 there has also been the Micro E-mini at one tenth of that size. In Germany, trading in the DAX future began in 1990 at the Deutsche Terminbörse, which became Eurex in 1998.

4. The option: the right without the obligation

A future binds both sides. The farmer gets his fixed price, but he also gives up the gain if wheat becomes more expensive. Some do not want that. They only want to be protected against the bad case and keep the good one. That is insurance, and insurance on a price is called an **option**: the right to buy or sell at a fixed price, without the obligation to do so. For this the buyer pays a premium up front.

Such contracts had existed for centuries, but only individually negotiated, without an exchange and without a recognised method for determining the fair price. In 1973 both changed in the same year:

- **The first options exchange.** In April 1973 the Chicago Board Options Exchange (Cboe) opened in Chicago. On the first day there were call options on 16 shares. As with the future, the contracts were uniform, and a clearing house stood between the parties. Put options were added in 1977.
- **The pricing formula.** Also in 1973, Fischer Black and Myron Scholes published their formula for the price of an option, and Robert Merton extended it. Its core is not a look into the future but a construction manual: whoever has sold an option can offset its risk by trading continuously in the underlying asset, and the cost

of this offsetting is the fair price. This manual is the working basis of option dealers to this day. Chapters 6 and 8 derive it.

After that came the same development as with the future:

- **1983:** options on indices, among them the S&P 500, with cash settlement.
- **1987:** On 19 October the Dow Jones fell by 22.6 % in one day. Since then, options that protect against sharp losses have permanently cost more than the 1973 formula gives. Chapter 6 explains this surcharge.
- **1993:** The VIX is introduced, an index that derives the expected fluctuation of the S&P 500 from option prices.
- **2005 to 2022:** Maturities become shorter. In 2005 options with weekly expiry arrived. Since 2022 an option on the S&P 500 has expired on every trading day. Options that expire on the day they are traded are called 0DTE (zero days to expiration).

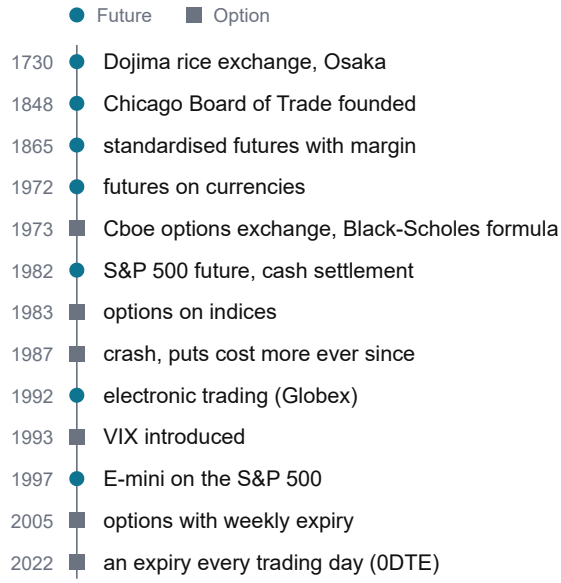


Fig. 1: The stations from sections 2 to 4 in chronological order. Dots stand for the future, squares for the option.

5. From the floor to the computer: who quotes the prices?

For a farmer to be able to sell at any time, someone has to want to buy at any time. On the old exchanges people took on that role: traders on the floor who continuously quoted buying and selling prices for their own account and lived on the small spread in between. This role is called **market maker**, literally: the one who makes the market.

With electronic trading the role moved from people to computers. Today specialised trading firms quote the prices, with programs that adjust their quotes many thousands of times per second. Well-known names are Citadel Securities, Virtu, Jane Street, Optiver, Susquehanna and IMC, plus the trading desks of large banks. They have no opinion on where the market is going. Their business is the spread, many times a day.

Whoever wants to buy or sell immediately and accepts the market maker's price for it is the **market taker**. Chapter 2 shows what this interplay looks like in the order book, what a market maker earns from and what he loses on.

6. Who is in the market today

Six roles are enough to understand the market. One firm can hold several of them.

Role	Who that is	What he wants
Hedger	funds, insurers, pension funds, companies	to get rid of a risk he already has
Investor and speculator	hedge funds, asset managers, private investors	to take a risk because he has a view
Market maker (dealer)	trading firms, banks	to earn the spread, with no view on direction
Arbitrageur	trading firms, banks	to earn from price differences in the same asset
Exchange	CME, Eurex, Cboe	to provide rules, order book and settlement, for a fee
Clearing house	CME Clearing, Eurex Clearing, OCC	to step between all parties and collect collateral

The first two groups bring the orders that move the market. They want something: protection or profit. Market makers and arbitrageurs want nothing of the kind. They react. The market maker takes the other side and tries to offset the risk he has taken on straight away. The arbitrageur makes sure that future, index and option fit together.

This difference is the key to the rest of the series. Whoever wants something trades on his view, and nobody knows that in advance. Whoever only reacts trades by rules that follow from his position. Those rules can be derived.

7. The common thread

One example connects all chapters. A pension fund holds shares worth 35 million USD and wants to protect itself against falling prices for one month.

1. **It buys put options on the S&P 500.** With these it may sell at a fixed price, however far the market falls. What it pays for that and why is shown in chapters 5 and 6.
2. **A market maker sells it the puts.** He now has the risk the pension fund has got rid of: if the market falls, he loses. How strongly his risk reacts to price changes is measured by the Greeks in chapter 7.
3. **He offsets the risk by selling futures on the S&P 500.** If the market falls further, he has to sell even more. This mechanism is the subject of chapter 8.
4. **The future is tied to the index.** If it deviates too far, arbitrageurs step in and trade the future against the shares. Why this works is shown in chapter 4.
5. **Each of these orders lands in an order book,** in which other market makers quote prices. How far an order moves the price depends on how much is resting there at that moment. That is what chapters 2 and 3 are about.
6. **From the positions of all dealers together** comes an estimate of whether their offsetting trades tend to dampen or amplify moves. That is the gamma exposure (GEX) in chapter 9.
7. **Whether this shows up in real prices** is checked in chapter 10 against our own measurements.

A pension fund's wish for protection thus turns, over four stations, into selling pressure in the index. Nobody in this chain has expressed a view on the market, except the pension fund itself.

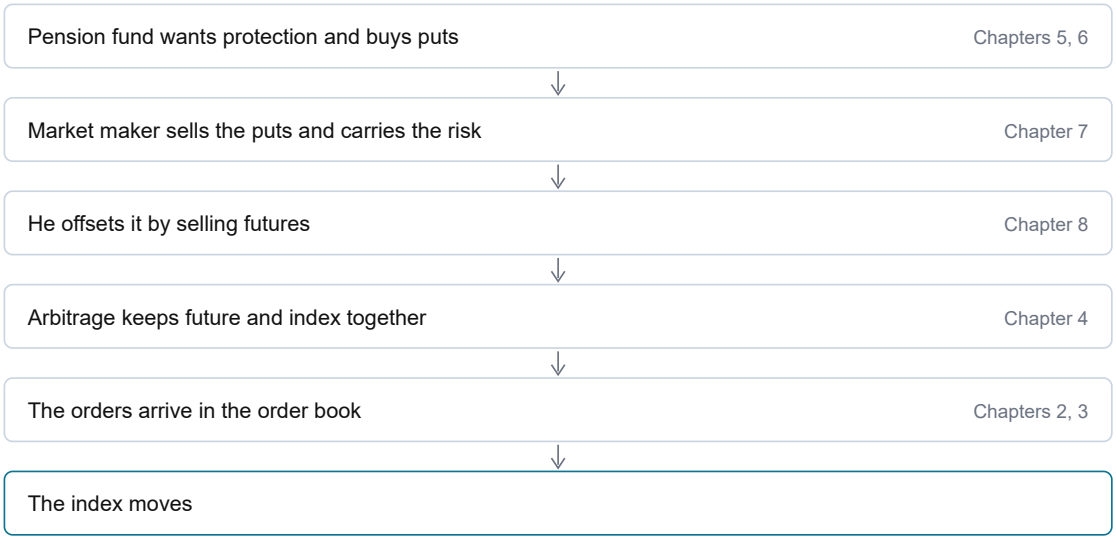


Fig. 2: The chain from the example. Each station is a chapter of this series.

Common misconception

Common misconception: "Futures and options are betting slips with no connection to the real economy." Both came about as hedges: the future for farmers and millers, the option as insurance on a price. This use carries the market to this day. The other half of the sentence is right: whoever trades without a risk to hedge is playing a zero-sum game. What one wins, another loses, and after fees the sum is negative.

What you should understand now

- A forward transaction fixes today the price for later. The future is its standardised form, with margin and a clearing house.
- Since 1982 there have been futures on stock indices. Nothing is delivered, settlement is in money.
- An option is a right without an obligation, that is, insurance on a price. Exchange-traded and with a pricing formula it has existed since 1973.
- The market has participants who want something (hedgers, investors) and participants who react (market makers, arbitrageurs). The behaviour of the second group can be derived.
- The chain of hedging, option, market maker, future and order book connects all chapters of this series.

Limits of this account

- **Shortened history.** Forward transactions existed long before 1730, for instance in the grain trade of antiquity and on the Amsterdam exchange of the 17th century. Only the stations that lead to today's market are named.
- **Simplified roles.** A bank is often hedger, market maker and investor at once. The roles describe activities, not firms.
- **Not covered:** interest-rate and commodity derivatives, swaps and over-the-counter trading. The series stays with futures and options on stock indices.

Continue with chapter 2: how a price forms in the order book and what market makers and market takers do there.

Chapter 2

The market as an auction: order book, market makers and market takers

In one sentence: The price forms where an order that wants to be executed immediately meets an order that has been waiting, and whoever waits demands a spread for it that is just large enough to cover his losses against better-informed traders.

What this chapter is: an explanation with a worked example, not a measurement. The order book in the example is made up to show the mechanics. The contract data of the E-mini future come from the CME. **Not a trading recommendation.**

Chapter 1 introduced the participants. This one is about the place where they meet.

1. An exchange is an auction that never ends

At an art auction there is one seller and many bidders. On an exchange there are many buyers and many sellers at the same time, and trading goes on without a break. The principle is the same: whoever wants to buy says what he will pay at most. Whoever wants to sell says what he wants at least. If two orders fit together, a trade happens.

The exchange collects all orders that are still waiting for a counterpart in a list. This list is called the **order book**. Every participant can see it.

2. Reading the order book

The example is the E-mini future on the S&P 500. Its price moves in steps of 0.25 points. One point is worth 50 USD, so one step is 12.50 USD per contract. Suppose the order book looks like this:

Side	Price	Resting contracts
Sell	7,001.00	310
Sell	7,000.75	260
Sell	7,000.50	180
Sell (ask)	7,000.25	120
Buy (bid)	7,000.00	110
Buy	6,999.75	170
Buy	6,999.50	240
Buy	6,999.25	300

Four terms are enough:

- **Bid:** the highest price at which someone wants to buy. Here 7,000.00.
- **Ask:** the lowest price at which someone wants to sell. Here 7,000.25.

- **Spread:** the distance between the two. Here 0.25 points, that is 12.50 USD per contract.
- **Depth:** how many contracts are resting at each price level.

There is no single "price" at this moment. There is a price for whoever wants to buy immediately (7,000.25) and one for whoever wants to sell immediately (7,000.00). What appears on the chart as the price is the price of the last trade.

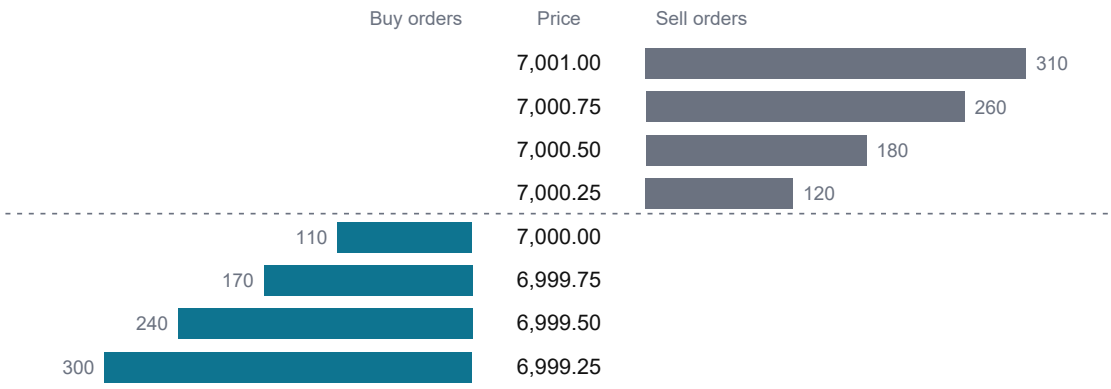


Fig. 3: The order book from the table. On the left the resting buy orders, on the right the resting sell orders. The gap in the middle is the spread.

3. Two kinds of orders

Limit order: "I buy 50 contracts, but at 6,999.75 at most." The order is placed in the order book and waits. The trader gets his price or nothing at all. He does not know whether and when he will be executed.

Market order: "I buy 200 contracts, now, at the best available price." The order is executed immediately against the resting sell orders, starting with the cheapest. The trader is certain to get his contracts. He does not know exactly at what price.

The market order for 200 contracts in the example:

- 120 contracts at 7,000.25. That empties this price level.
- 80 contracts at 7,000.50.

$$\text{Average price} = \frac{120 \times 7{,}000.25 + 80 \times 7{,}000.50}{200} = 7{,}000.35$$

After the trade the ask is at 7,000.50. The buyer has moved the price up by one level, simply by using up the cheapest offers. That is the basic mechanism of every price move: **the price moves when orders that want immediate execution empty a price level.**

4. Makers and takers

Two roles follow from the two kinds of orders:

- Whoever places a limit order in the book and waits **provides** liquidity. In this trade he is the **maker**.
- Whoever trades against a resting order with a market order **takes** liquidity. He is the **taker**.

Maker and taker are roles in a single trade, not fixed groups of people. A private investor who places a limit order is a maker at that moment. Nevertheless there are firms whose entire business is quoting prices. They are called **market makers**.

	Market maker	**Market taker**
Who that typically is	trading firms such as Citadel Securities, Virtu, Optiver, Jane Street, trading desks of banks	funds, asset managers, hedgers, private investors
What he does	continuously quotes buying and selling prices	buys or sells when he needs to
What he wants	to earn the spread, many times	to build or reduce a position
View on direction	none	usually yes, or a need to hedge
What he gives up for it	control over when and against whom he trades	half the spread on every trade

On some exchanges market makers have a contract: they commit to quoting prices almost continuously and get lower fees in return. On the large futures markets most quote voluntarily. This has a consequence that becomes important in chapter 3: they may withdraw their orders at any time.

5. What a market maker earns from

The market maker places a buy order at 7,000.00 and a sell order at 7,000.25 in the book at the same time. If first a seller comes and then a buyer, he has bought at 7,000.00 and sold at 7,000.25:

$$0.25 \text{ points} \times 50 \text{ USD} = 12.50 \text{ USD per contract}$$

Afterwards his inventory is back at zero. He has not earned from any price move, only from the spread. The amount per trade is tiny. The business pays through volume.

On stock exchanges there is a second source. Many exchanges pay the maker a small rebate per share and charge the taker a fee. This pricing model is called **maker-taker**. It was introduced in 1997 by an electronic trading platform that wanted to draw orders away from the established exchanges. In the US the taker fee is capped by regulation at a fraction of a cent per share, and the maker rebate lies slightly below it. At the CME, by contrast, both sides pay a fee per contract in the index futures.

6. What a market maker loses on

The market maker does not know who is trading against him. That is where his risk lies.

Inventory risk. He has bought at 7,000.00, but the buyer he wanted to sell on to does not come. Instead the price falls by one point. Loss: 50 USD per contract, as much as the income from four successful round trips.

Adverse selection. This is the bigger problem. Some takers know more than the market maker: they have read a piece of news faster, or they know that a much larger order follows behind theirs. Such traders only buy when the price is about to rise and only sell when it is about to fall. Against them the market maker loses every time. The term comes from insurance: whoever knows a risk that the insurer does not know is more likely to insure himself.

The market maker cannot tell the two groups apart. So he has to earn as much from the uninformed as he loses to the informed.

7. Derivation: how wide does the spread have to be?

The simplest model for this goes back to Lawrence Glosten and Paul Milgrom (1985). The assumptions:

- A share π of the takers is informed. The rest trade for other reasons, for example because a fund has to invest money.
- An informed taker knows that the price is about to jump by J points and trades in the right direction.
- The market maker sets his quotes half a spread $s/2$ away from the current value: buy at value minus $s/2$, sell at value plus $s/2$.

Against an uninformed taker the market maker earns half the spread:

$$\text{Gain} = \frac{s}{2}$$

Against an informed taker he sells, for example, at value plus $s/2$, and immediately afterwards the value has risen by J. He has sold too cheaply:

$$\text{Loss} = J - \frac{s}{2}$$

The condition. Under competition nothing is left for the market maker on average. Gain times frequency equals loss times frequency:

$$(1-\pi)\cdot\frac{s}{2} = \pi\cdot\left(J - \frac{s}{2}\right)$$

Multiplied out:

$$\frac{s}{2} - \pi\cdot\frac{s}{2} = \pi\cdot J - \pi\cdot\frac{s}{2}$$

The term $\pi\cdot s/2$ stands on both sides and drops out:

$$\frac{s}{2} = \pi\cdot J \quad\Longrightarrow\quad s = 2\,\pi\,J$$

So the spread is twice the share of the informed times the size of their advantage.

Worked example. Suppose every tenth taker is informed ($\pi = 0.1$) and his advantage is 1.25 points:

$$s = 2\times 0.1\times 1.25 = 0.25 \text{ points}$$

That is exactly one price level, the spread from the order book above. The check: out of ten trades, nine bring 0.125 points each, together 1.125. One costs 1.25 minus 0.125, that is also 1.125. The sum is zero.

Three things follow from the formula:

- **More informed traders, wider spread.** Before important economic data the danger of trading against someone who processes the number faster rises. Market makers set their quotes further apart or withdraw them.
- **Larger jumps, wider spread.** In turbulent markets J is large.
- **The spread is not a profit.** It is the price all takers pay so that the market maker can carry the losses against the few who are informed.

8. From the floor to the computer

Until the 1990s market makers stood on the exchange floor as people. On the New York Stock Exchange one "specialist" was responsible for each share, in the pits of the Chicago futures exchanges the "locals" traded for their own account. Spreads were wide. US shares were quoted in sixteenths of a dollar until 2001, so the smallest possible spread was 6.25 cents. Since the switch to decimal prices it has been 1 cent.

With electronic trading the profession became a computing problem: whoever adjusts his quotes to new information fastest is caught least often by informed traders and can quote the narrowest spread. This explains why a few highly technical firms dominate this market today and why speed matters to them. The formula from section 7 holds unchanged. Only J has become smaller, because quotes are updated faster.

Common misconception

Common misconception: "Market makers move the price to where they want it." A market maker earns the spread and wants as little inventory as possible at the end of the day. A price move is a risk to him, not a goal. The price moves when takers empty price levels. The market maker reacts to that and requotes. We have tested trading rules built on the idea of liquidity being deliberately collected: they hit their targets as often as chance does (to the study).

What you should understand now

- The order book is the list of all resting buy and sell orders. Bid and ask are its best prices, the spread is the distance between them.
- A limit order waits for its price, a market order takes the best available one. The price moves when market orders use up a price level.
- Maker and taker are roles in a trade. Market makers are firms that quote prices professionally and earn the spread.
- A market maker loses to traders who know more than he does (adverse selection) and to price moves while he holds inventory.
- In the simplest model the spread is $s = 2\,\pi\,J$: it grows with the share of informed traders and with the size of the price jumps.

Limits of this account

- **Made-up order book.** The quantities in the example show the mechanics. Real books change many thousands of times per second, and a large part of the orders is cancelled before being executed.
- **Simplest model.** The derivation knows only one price jump of fixed size and no inventory. Real market makers also manage their inventory, hedge it in related markets and adjust their quotes to the order flow.
- **Minimum step.** The spread cannot be smaller than one price level. In the E-mini it is almost always exactly one level; competition there takes place through size and queue position in the book.
- **Not covered:** hidden orders, off-exchange trading and the routing of customer orders to market makers for payment.

Continue with chapter 3: what liquidity is, how far an order moves the price and when the resting orders disappear.

Chapter 3

Liquidity: depth, price impact and when the resting orders disappear

In one sentence: Liquidity is not a fixed stock but an offer by market makers that they withdraw as soon as quoting prices becomes dangerous for them, and that is why it is smallest exactly when many need it.

What this chapter is: an explanation with worked examples. The order book is the made-up example from chapter 2. The square-root rule is an empirical value from the literature, not a law. The numbers on spreads, auctions and VWAP come from our own measurements and are marked as such. **Not a trading recommendation.**

Chapter 2 showed how an order is executed in the order book. This one is about what that costs and what those costs depend on.

1. What liquidity is: three measures

A market is called liquid when a lot can be traded without moving the price. This can be measured by three things:

- **Spread:** the distance between bid and ask. It is the price of a small trade.
- **Depth:** the quantity resting at the best price levels. It determines what a large trade costs.
- **Recovery:** how quickly the book is refilled after a large order.

Volume is not one of them. Volume counts how much was traded. Liquidity describes how much could be traded. On a day of panic volume is high and liquidity is poor.

2. A large order runs through the book

The order book from chapter 2, sell side:

Price	Resting contracts	Cumulative
7,000.25	120	120
7,000.50	180	300
7,000.75	260	560
7,001.00	310	870

The bid is at 7,000.00, so the midpoint between bid and ask is 7,000.125. A fund buys 500 contracts with a market order:

- 120 contracts at 7,000.25
- 180 contracts at 7,000.50
- 200 contracts at 7,000.75

$$\text{Average price} = \frac{120 \times 7{,}000.25 + 180 \times 7{,}000.50 + 200 \times 7{,}000.75}{500} = 7{,}000.54$$

Measured against the midpoint the fund has paid 0.415 points too much on average. One point is worth 50 USD:

$$0.415 \times 50 \text{ USD} \times 500 = 10{,}375 \text{ USD}$$

These costs are called **slippage**. They have two parts. Half the spread (0.125 points) would have been paid by a single contract as well. The rest (0.29 points) arises because the order was larger than the best price level.

An order that empties several price levels at once in this way is called a **sweep**. After the trade the ask is at 7,000.75, two levels higher than before. This shift is the **price impact** (market impact) of the order.

What happens next decides the true cost. If market makers refill the emptied levels immediately, the move was temporary. If they do not, because they suspect information behind the order, the price stays up. Then the market has learned from the order.

3. The square-root rule of thumb

Large investors therefore split their orders into many small pieces and spread them over hours. How strongly the price reacts in total then depends on the size of the whole order relative to daily turnover. There is a rule of thumb for this that has been measured similarly in many markets:

$$\text{Price impact} \approx Y \cdot \sigma_{\text{day}} \cdot \sqrt{\frac{Q}{V}}$$

Here σ_{day} is the usual daily fluctuation, Q the size of the order, V the daily turnover and Y a factor of the order of one.

This is an empirical value, not a derivation. The rule summarises measurements on orders of large investors. Depending on the study the exponent lies between about 0.4 and 0.7, and the factor Y differs from market to market.

Worked example. The index is at 7,000 and fluctuates by 1 % on a usual day, that is by 70 points. With $Y = 1$:

Order as a share of daily turnover	Price impact	In points
0.25 %	0.05 %	3.5
1 %	0.10 %	7.0
4 %	0.20 %	14.0
16 %	0.40 %	28.0

The message is in the root: an order four times as large moves the price only twice as far. The first contracts are the most expensive. Conversely this means: even a small order moves the price noticeably, and more than a glance at the visible depth would suggest.

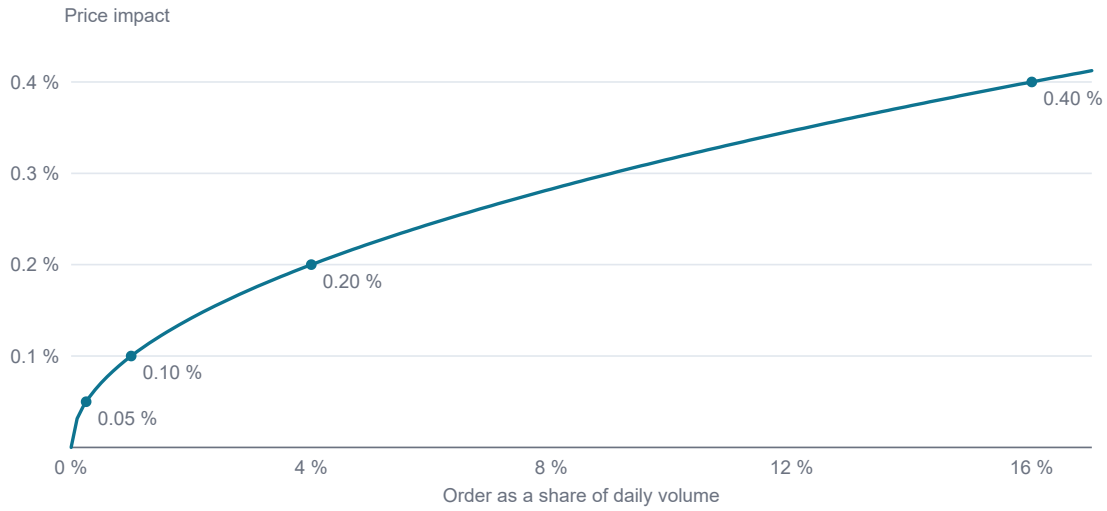


Fig. 4: Price impact by the square-root rule of thumb for a daily fluctuation of 1 % and Y = 1. The dots are the four rows of the table.

4. When liquidity disappears

The spread from chapter 2 was $s = 2\,\pi\,J$: share of the informed times the size of their advantage. From this it follows directly when a market maker widens his quotes or takes them out of the book altogether.

Before and after scheduled news. At 14:30 German time the US labour market data are released. In the second after that almost every order is "informed", because it already knows the number, and the jump J is large. Whoever quotes prices now loses. Market makers withdraw their orders beforehand.

This can be measured. On tick data for gold (April 2023 to December 2024, 452 trading days, morning and midday in New York) we looked for all moments in which the spread jumped to at least three times its usual value:

	Usual spread	Spread in the event	Back to normal after (median)	In 90 % of cases after
Outside scheduled releases	0.13 USD	0.45 USD	1 second	13 seconds
Around data releases	0.16 USD	0.65 USD	12 seconds	39 seconds

Outside scheduled releases such a jump occurred 2.5 times a day on average. A market order at that moment costs three to four times as much. No direction could be derived from it: whoever followed the move after such a jump or traded against it was at zero before costs (−0.07 basis points, t −0.26) and negative after costs (to the study).

At the edges of trading hours. At night and between sessions fewer market makers quote prices. The spread is wider, the depth smaller. The same order then moves the price further.

In fast moves. When a sweep empties several levels, the remaining market makers do not know whether there is information behind it. As a precaution they pull their orders. Depth falls exactly when further orders arrive. This is how medium-sized moves sometimes become large ones: the price falls through a thin book.

5. Auctions: one price for all

Twice a day trading on stock exchanges works differently: at the open and at the close. Instead of executing orders continuously, the exchange collects them for a few minutes and then computes a single price at which all matching orders are executed at once.

The rule: the price that applies is the one at which the largest quantity can be traded.

Worked example. For four possible prices it is known how much would be bought and sold in total:

Price	Willing to buy at this price or higher	Willing to sell at this price or lower	Tradable
6,999	900	200	200
7,000	700	500	500
7,001	400	800	400
7,002	150	1,000	150

Tradable is the smaller of the two quantities in each case. It is largest at 7,000, so 7,000 is the auction price. 500 units are traded, a surplus of 200 buy orders remains unexecuted. This surplus is called the **imbalance**. Some exchanges publish it shortly before the auction.

Why auctions matter:

- **The closing auction sets the closing price.** The value of every fund is computed from it. Index funds, which have to replicate an index exactly, therefore prefer to trade in the closing auction. On an ordinary day more liquidity is concentrated there than in any other minute.
- **The opening auction processes the night.** All orders that have accumulated since the previous evening meet in one price.
- **Settlement prices form in auctions.** The DAX future is settled on expiry day at the price of the midday auction at 13:00, the large US index futures at the opening prices of the shares (chapter 4).

We have measured that forced trading arrives there. On the days on which large indices change their composition, the price range in the auction window was 1.2 to 1.7 times as wide as on comparison days in all eight index markets examined (to the study). From the published imbalance before the closing auction, by contrast, nothing could be earned (to the study).

6. VWAP and TWAP: the yardstick of execution

A fund that buys 500 contracts spread over the day wants to know afterwards whether it was executed well. For that it needs a reference price. Two are common.

VWAP (volume weighted average price) is the average price of all trades of the day, weighted by their size:

$$\mathrm{VWAP} = \frac{\sum_i p_i \cdot v_i}{\sum_i v_i}$$

TWAP (time weighted average price) is the average of prices over time, without weighting by size.

Worked example. Three trades: 100 contracts at 7,000, 300 at 7,004 and 100 at 7,010.

$$\mathrm{VWAP} = \frac{100 \times 7{,}000 + 300 \times 7{,}004 + 100 \times 7{,}010}{500} = 7{,}004.4$$

The plain average of the three prices would be 7,004.67. The VWAP lies closer to 7,004 because the largest quantity was traded there.

The VWAP is a **yardstick**: whoever bought below the VWAP bought better than the average of the market. Many execution programmes are designed to come as close to it as possible. They therefore trade more when the market trades a lot, above all at the open and the close.

From this it is often concluded that the VWAP is a line from which the price bounces. We tested that, on the TWAP, because no true exchange volume is available for index CFDs. In the Nasdaq 100, S&P 500 and DAX (2015 to 2026, 17,858 touches of the TWAP of the running session) the line was broken in about half the cases on the first, second and third touch (52 %, 49 %, 50 %). A random market with the same fluctuation but no direction gave the same values (54 %, 49 %, 52 %, to the study). Rules with VWAP bands and VWAP bounces earned nothing after costs (to the study). The VWAP measures how well someone executed. As a price level it is of no use.

Common misconception

Common misconception: "High volume means high liquidity." Volume is what was traded. Liquidity is what would be tradable at the current price. In a sell-off volume rises because many have to sell at once, while the spread widens and depth falls. The most expensive minutes of a year are those with record turnover.

What you should understand now

- Liquidity has three sides: spread, depth and the speed at which the book is refilled. Volume is none of them.
- An order larger than the best price level pays more than half the spread. The difference to the midpoint is called slippage, the lasting shift price impact.
- By the square-root rule of thumb an order four times as large moves the price about twice as far. That is an empirical value.
- Market makers withdraw their orders when the share of informed orders or the size of the jumps rises: around news, at the edges of trading hours and in fast moves.
- An auction forms a single price for all, where the largest quantity is tradable. VWAP and TWAP are reference prices for the quality of an execution.

Limits of this account

- **Rule of thumb.** The square-root rule holds on average over many orders of large investors. For a single order the impact can be a multiple or zero. For very small and very large orders it fits less well.
- **Own measurements.** The numbers on the spread come from gold ticks of a single broker, not from the order book of an exchange, and cover 21 months. We have not measured this for index futures. The TWAP evaluation uses index CFDs on minute data.
- **Auction rule.** Real exchanges have additional rules for the case that several prices give the same quantity, and safeguards against extreme auction prices.
- **Not covered:** liquidity off the exchange, that is, the trading of large blocks between banks and funds that never appears in the order book.

Continue with chapter 4: what exactly a future is, who uses it for what and why its price is tied to the index.

Chapter 4

Futures: what they are, who uses them for what and why their price is tied to the index

In one sentence: A future is a contract to buy an asset at a later date, whose price is fixed today such that nobody can earn from it without risk by instead buying the asset itself and holding it until that date.

What this chapter is: an explanation with worked examples. Contract data come from CME and Eurex. Index levels, interest rate, dividend yield and margin in the example are assumptions of a realistic order of magnitude and named as such. **Not a trading recommendation.**

Chapter 1 told why the future was invented. This one is about how exactly it works.

1. What a future is

A future is a contract between two parties:

- The **buyer** (he is "long") commits to taking a certain asset on a fixed day at a price agreed today.
- The **seller** (he is "short") commits to delivering at that price.

With a future on a stock index nothing is delivered at the end. Only the difference is settled in money. If the index is higher than the agreed price on expiry day, the buyer receives the difference from the seller. If it is lower, the buyer pays.

So every future has two sides. What one gains, the other loses. The number of contracts currently open is called **open interest**. It plays a large role in chapter 9.

The exchange fixes everything except the price, so that everyone trades the same contract:

Fixed	For the E-mini on the S&P 500
What it is written on (underlying)	S&P 500
Value of one index point (multiplier)	50 USD
Smallest price step (tick)	0.25 points, that is 12.50 USD
Expiry	third Friday in March, June, September and December
Trading hours	Sunday evening to Friday evening, almost continuously

Only the price is negotiated. If the future is at 7,000 points, one contract moves a notional value of

$$7{,}000 \times 50 \text{ USD} = 350{,}000 \text{ USD}$$

The most important index futures at a glance:

Future	Exchange	Value per point	Notional at index level
E-mini S&P 500 (ES)	CME	50 USD	350,000 USD at 7,000
Micro E-mini S&P 500 (MES)	CME	5 USD	35,000 USD at 7,000
E-mini Nasdaq 100 (NQ)	CME	20 USD	580,000 USD at 29,000
Micro E-mini Nasdaq 100 (MNQ)	CME	2 USD	58,000 USD at 29,000
DAX future (FDAX)	Eurex	25 EUR	625,000 EUR at 25,000
Mini-DAX future (FDXM)	Eurex	5 EUR	125,000 EUR at 25,000
Micro-DAX future (FDXS)	Eurex	1 EUR	25,000 EUR at 25,000

The index levels in the last column are round example values.

2. Who uses futures and for what

Hedging. A pension fund holds shares worth 35 million USD that move roughly like the S&P 500. It fears a decline but does not want to sell the shares. It sells futures:

$$\frac{35{,}000{,}000\ \text{USD}}{350{,}000\ \text{USD}} = 100\ \text{contracts}$$

If the market falls by 5 %, that is by 350 points, the shares lose 1.75 million USD. The 100 futures sold gain:

$$350 \times 50\ \text{USD} \times 100 = 1{,}750{,}000\ \text{USD}$$

The loss is offset. If the market rises by 5 % instead, the shares gain 1.75 million USD and the futures lose the same amount. The pension fund has given up its price risk for the duration of the hedge, for better and for worse. That is the difference from the option in chapter 5.

Investing quickly. A fund receives 35 million USD of new money. Buying 500 shares takes time and costs money. Buying 100 futures is a single trade. The fund is in the market immediately and later swaps the futures for shares bit by bit.

Taking a view on direction. Whoever expects rising prices buys futures. Whoever expects falling prices sells them. Selling a future is just as easy as buying one. With shares one would first have to borrow them.

Aligning prices. Arbitrageurs trade the future against the shares as soon as the two no longer fit together (section 4).

Offsetting options. Option dealers hedge their books with futures. That is the subject of chapter 8 and the reason this chapter comes before the options.

Why the future and not the shares themselves? One trade instead of five hundred, low costs, trading almost around the clock, selling without borrowing, and very little capital tied up. The last point needs a section of its own.

3. Margin, leverage and daily settlement

Whoever buys a future does not pay 350,000 USD. He only posts collateral, the **margin**. The exchange sets it and adjusts it when the market becomes more turbulent. For index futures it is a few per cent of the notional value. For the example we assume 21,000 USD per contract, which is 6 %.

From this the **leverage** arises:

$$\text{Leverage} = \frac{350{,}000\ \text{USD}}{21{,}000\ \text{USD}} \approx 16.7$$

If the index moves by 1 %, that is by 70 points, the value of the contract changes by 3,500 USD. That is 16.7 % of the collateral posted. Leverage is not a feature you add on. It results from posting only a small collateral for a large notional value.

So that nobody loses more than he has posted in such moves, the clearing house settles **every day**. Gains are credited, losses deducted, based on the official daily settlement price of the future. An example for one contract bought at 7,000:

Day	Settlement price	Change in points	Credit or debit	Margin account balance
Purchase	7,000			21,000 USD
1	6,950	−50	−2,500 USD	18,500 USD
2	7,020	+70	+3,500 USD	22,000 USD
3	6,980	−40	−2,000 USD	20,000 USD

After three days the future is 20 points below the purchase price. Exactly 20 times 50 USD, that is 1,000 USD, are missing from the account. The loss did not fall due only on expiry day but day by day.

If the account falls below a minimum level, the broker demands additional funds (margin call). If the money does not arrive, he closes the position. That is the protection against the default risk from chapter 1. For the trader it means: a position can be closed by force before the market recovers.

4. The fair price: derivation from arbitrage

Why is the future at 7,047 when the index is at 7,000? It is tempting to see an expectation in that: the market reckons with rising prices. That is wrong. The price follows from a calculation in which the expectation does not appear.

There are two ways of owning the index in 90 days:

- **Way A:** buy the future today. No money flows, in 90 days settlement takes place at the price F.
- **Way B:** buy the shares of the index today, with borrowed money, and hold them for 90 days.

Both ways lead to the same result, so they have to cost the same. What does way B cost?

The assumptions: index $S = 7{,}000$, interest rate $r = 4\,\%$ per year, dividend yield $q = 1.3\,\%$ per year, term $T = 90/365$ years.

Step 1: take out a loan and buy the shares. For 7,000 (per index unit).

Step 2: pay interest. After 90 days:

$$7{,}000 \times 0.04 \times \frac{90}{365} = 69.04$$

Step 3: collect dividends. The shares belong to the holder, so he receives the dividends:

$$7{,}000 \times 0.013 \times \frac{90}{365} = 22.44$$

Step 4: add up. In the end way B costs the purchase price plus interest minus dividends:

$$7{,}000 + 69.04 - 22.44 = 7{,}046.60$$

That has to be the price of the future. In general:

$$F = S \cdot \big(1 + (r - q) \cdot T\big)$$

The expression $(r - q)$ is called the **cost of carry**: the cost of holding, interest minus dividend. With continuous compounding the formula reads $F = S \cdot e^{(r-q)T}$ and gives 7,046.76 here. The difference is negligible for short terms.

What happens if the price deviates? Suppose the future is at 7,060. An arbitrageur sells the future and at the same time goes way B: he buys the shares on credit. In 90 days he economically delivers the shares against the future. His profit is already certain today:

$$7{,}060 - 7{,}046.60 = 13.40 \text{ points} = 670 \text{ USD per contract}$$

He carries no price risk in doing so, because what the shares lose, the future sold gains. Many do this at the same time. Their selling pushes the future down, their buying lifts the shares, until the gap is closed. If the future is too low, it runs the other way: buy the future, sell the shares.

In practice a narrow band remains around the fair price in which arbitrage does not pay because of trading costs. For the large index futures it is a few tenths of a point wide.

5. Basis and expiry

The distance between future and index is called the **basis**:

$$\text{Basis} = F - S = S \cdot (r - q) \cdot T$$

In the example that is 46.60 points at 90 days to expiry. The basis depends on the time to expiry T. It gets smaller with every day:

Time to expiry	Fair price of the future	Basis
90 days	7,046.60	46.60
60 days	7,031.07	31.07
30 days	7,015.53	15.53
0 days	7,000.00	0.00

On expiry day the basis is zero. Future and index meet. This is called **convergence**, and it is enforced: at the end the future is settled at the index level.

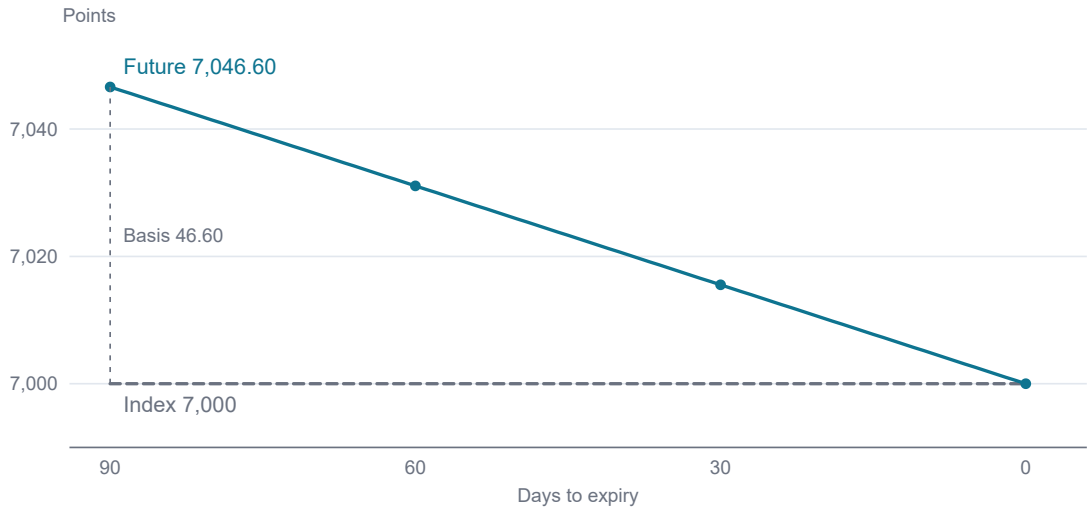


Fig. 5: Fair price of the future with the index unchanged at 7,000, an interest rate of 4 % and a dividend yield of 1.3 %. The basis melts away evenly and is zero on expiry day.

How settlement works. What matters is a single price on expiry day:

- **S&P 500 and Nasdaq 100:** settlement takes place on the third Friday of the expiry month at an index level computed from the opening prices of all index shares that morning.
- **DAX:** settlement takes place on the third Friday at the index level from the midday auction of Xetra trading at 13:00.

Whoever depends on that price trades in those auctions. That concentrates turnover there (chapter 3).

The DAX is a special case. The DAX counts dividends into the index level. So the holder of a DAX future does not miss a dividend, and in the formula $q = 0$. The basis is therefore larger than for a price index. An example from our own data: on 8 July 2026 the price with the highest volume in the DAX future was around 112 points above that of the CFD on the DAX (25,139 against 25,027), at 72 days to the September expiry. Annualised that is around 2.3 %, the order of magnitude of the euro money-market rate.

6. The roll

A future expires. Whoever wants to hold his position longer has to switch: close the expiring contract and open the next one. This is called the **roll**. For the large index futures turnover moves into the next contract about a week before expiry.

The next contract has three more months to expiry and is therefore higher, in the example by about 47 points. This difference is neither a gain nor a loss. It is the cost of carry for the additional three months, which would have been paid anyway.

Two practical consequences:

- **Continuous charts have jumps.** A chart that simply strings the contracts together shows a gap on roll day that never existed. For evaluations such series have to be adjusted.
- **Whoever is permanently long pays the carry continuously.** It sits in every new contract as a premium that melts away until expiry.

7. How index CFDs are tied to it

Many private investors trade indices not through futures but through CFDs (contracts for difference) with a broker. A CFD is not an exchange product. The broker is the counterparty and quotes his own prices. He derives them from the future, because that is where he hedges:

$$\text{CFD price} \approx \text{future} - \text{basis}$$

The CFD price therefore lies close to the index and keeps moving even when the stock exchange is closed and only the future trades.

The difference lies in the financing. In the future the cost of carry is in the price. With the CFD it is charged individually every night, as financing cost (swap), plus a mark-up by the broker. In our measurement, holding an index CFD overnight cost 1.5 to 2.1 basis points of swap, in addition to 2 basis points of spread. That was enough to consume the measurable overnight effect of the US indices completely (to the study).

For all further chapters this holds: when the future moves, the CFD moves. What option dealers do in the future arrives in the CFD price.

Common misconception

Common misconception: "The future is above the index, so the market expects rising prices." The basis is interest minus dividend times time to expiry. An expectation does not appear in the calculation. If the future were more expensive because of an expectation, anyone could earn from it without risk by selling it and buying the shares. The market's expectations are in the index level itself, not in the distance of the future from it.

What you should understand now

- A future is a standardised contract with two sides. With an index future only the difference is settled in money. One E-mini on the S&P 500 moves 50 USD per point.
- It is used for hedging, for investing quickly, for taking a view on direction, for arbitrage and for offsetting option books.
- Only a margin of a few per cent is posted. Leverage arises from that. Gains and losses are settled daily.
- The fair price is $F = S \cdot (1 + (r - q) \cdot T)$: index plus interest minus dividend. Arbitrage keeps the future at this price.
- The basis melts to zero by expiry. In the roll no additional costs are paid, only the carry for the next term. Index CFDs are derived from the future and charge the financing per night.

Limits of this account

- **Assumptions.** An interest rate of 4 %, a dividend yield of 1.3 %, a margin of 6 % and the index levels are example values. In reality dividends do not flow evenly but are bunched on certain dates. The basis then runs in steps.
- **Arbitrage is not free.** It needs capital, credit lines and the ability to sell shares one does not own. In crises the basis can therefore deviate from fair value for hours or days.

- **The DAX example** compares the highest-volume prices of two volume profiles, not two quotes made at the same moment, and is a single day.
- **Not covered:** futures on commodities, interest rates and currencies, in which storage costs and delivery play a role of their own, and tax treatment.

Continue with chapter 5: the option, the right without the obligation, and why call and put are linked by a simple equation.

Chapter 5

Options: call, put, payoff, time value and put-call parity

In one sentence: An option is insurance on a price, its buyer pays a premium up front and can only gain after that, and because a call can be rebuilt from a put and the index, the prices of call and put are firmly tied to each other.

What this chapter is: an explanation with worked examples. The option prices in the example are computed with the formula from chapter 6 (S&P 500 at 7,000, 30 days to expiry, 16 % volatility), not observed in the market. **Not a trading recommendation.**

Chapter 4 showed how a pension fund hedges its portfolio with futures and gives up gains in the process. The option solves exactly that problem.

1. What an option is

An option gives its buyer a **right**, but no obligation:

- A **call** is the right to **buy** an asset at a fixed price.
- A **put** is the right to **sell** an asset at a fixed price.

Five items describe an option completely:

Item	Meaning	In the example
Underlying	what the option is written on	S&P 500
Type	call or put	call
Strike	the fixed price	7,000 points
Expiry	until when the right holds	in 30 days
Multiplier	value of one point	100 USD

The price the buyer pays for this right is called the **premium**. He pays it immediately and never gets it back.

For every buyer there is a seller. He is called the **writer**. He receives the premium and carries the obligation in return: with a call he has to deliver, with a put he has to take delivery.

Options on the S&P 500 are settled in money at the end, like the future. They can only be exercised on expiry day. This design is called European. Options on individual shares can be exercised at any time, which is called American. For the mechanics of this series the difference does not matter.

2. The payoff at expiry

On expiry day the value of an option is simple. The index is then at S_T, the strike is K.

Call: The right to buy at 7,000 is only worth something if the index is above it. Then it is worth exactly the difference. Otherwise it expires worthless.

$$\text{Call payoff} = \max(S_T - K,\ 0)$$

Put: The right to sell at 7,000 is only worth something if the index is below it.

$$\text{Put payoff} = \max(K - S_T,\ 0)$$

The buyer's profit is the payoff minus the premium paid. In the example the call and the put with strike 7,000 each cost 128.09 points, that is 12,809 USD per contract.

Index at expiry	Call payoff	Call profit	Put payoff	Put profit
6,700	0	−128.09	300	+171.91
6,900	0	−128.09	100	−28.09
7,000	0	−128.09	0	−128.09
7,100	100	−28.09	0	−128.09
7,300	300	+171.91	0	−128.09

Three things stand out:

- **The buyer's loss is limited.** He cannot lose more than the premium, however far the market runs against him.
- **The gain is not.** With the call it grows with every point the index rises.
- **Being right is not enough.** At 7,100 the index has risen and the call buyer has still lost. He only earns above 7,128.09, the strike plus the premium. This level is called the **break-even**. For the put it is at 6,871.91.

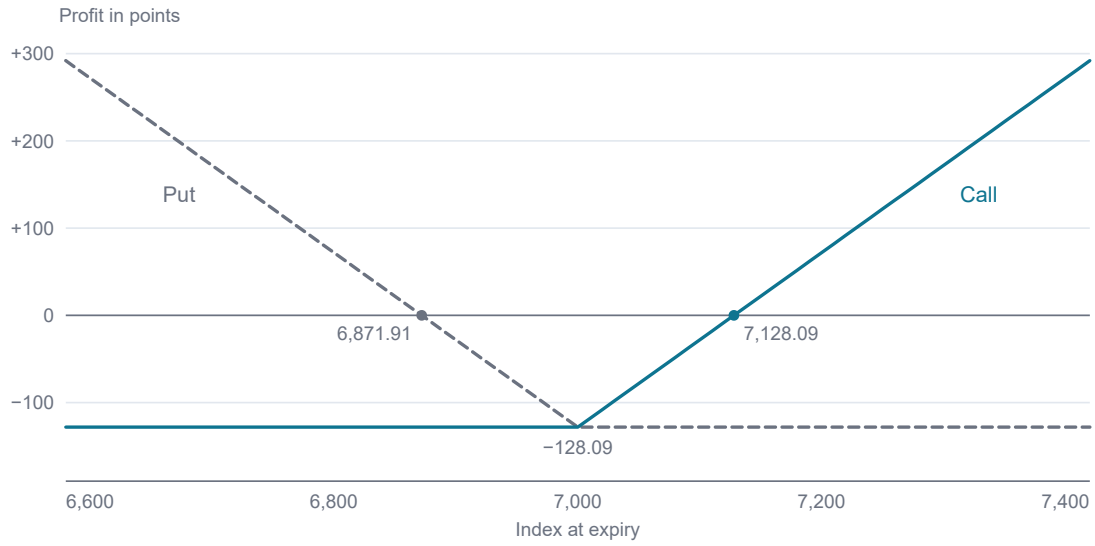


Fig. 6: Profit and loss on expiry day for a bought call and a bought put with strike 7,000 and a premium of 128.09 points. The kink is at the strike, the zero line is crossed at the break-even levels.

3. The four basic positions

Buy or sell, call or put: that makes four positions. The seller's view is the mirror image of the buyer's.

Position	Earns if the index	Largest gain	Largest loss
Call bought	rises sharply	unlimited	the premium
Call sold	does not rise sharply	the premium	unlimited
Put bought	falls sharply	strike minus premium	the premium
Put sold	does not fall sharply	the premium	strike minus premium

The buyer pays little and rarely gains a lot. The seller often takes in a little and rarely loses a lot. That is the relation between insured and insurer.

4. Hedging with puts instead of futures

The pension fund from chapter 4 holds shares worth 35 million USD. With 100 futures sold it was protected against losses, but also cut off from gains. Now it buys puts with strike 7,000 instead.

One contract covers 7,000 times 100 USD, that is 700,000 USD:

$$\frac{35{,}000{,}000}{700{,}000} = 50 \text{ contracts}, \qquad 50 \times 12{,}809 \text{ USD} = 640{,}450 \text{ USD premium}$$

That is 1.83 % of the portfolio for 30 days of protection. The comparison after one month, in million USD:

Index after 30 days	No hedge	With 100 futures sold	With 50 puts bought
−5 % (6,650)	−1.75	0.00	−0.64
unchanged (7,000)	0.00	0.00	−0.64
+5 % (7,350)	+1.75	0.00	+1.11

With puts the pension fund loses the premium in the worst case and keeps most of the gain in the good case. For that it pays 640,450 USD in every case. The future costs nothing and takes both sides away. The option costs something and takes only the bad side away. What this protection is worth, that is, why the premium is 128.09 points and not 50 or 300, is the question of chapter 6.

5. Intrinsic value and time value

Before expiry an option costs more than it would bring if exercised immediately. Its price has two parts.

- **Intrinsic value:** what the option would be worth if it expired now. For the call $\max(S - K, 0)$, for the put $\max(K - S, 0)$.
- **Time value:** the rest. It pays for the possibility that the price still moves favourably before expiry.

$$\text{Price} = \text{intrinsic value} + \text{time value}$$

Depending on where the strike lies relative to the current price, an option has one of three names:

- **In the money:** the option has intrinsic value. A call with a strike below the price, a put with a strike above it.
- **At the money:** strike and price are equal.
- **Out of the money:** the option has no intrinsic value and consists of time value only.

Three strikes at an index level of 7,000:

Strike	Call: price	of which intrinsic value	of which time value	Put: price	of which intrinsic value	of which time value
6,900	183.38	100	83.38	83.38	0	83.38
7,000	128.09	0	128.09	128.09	0	128.09
7,100	85.12	0	85.12	185.12	100	85.12

Two observations that become important later:

- **Time value is largest at the money.** That is where it is most uncertain whether the option will be worth something in the end.
- **Call and put with the same strike have the same time value.** That is no coincidence. It follows from put-call parity in the next section.

With every day that passes time value melts, until only the intrinsic value is left on expiry day. How fast is shown in chapter 7.

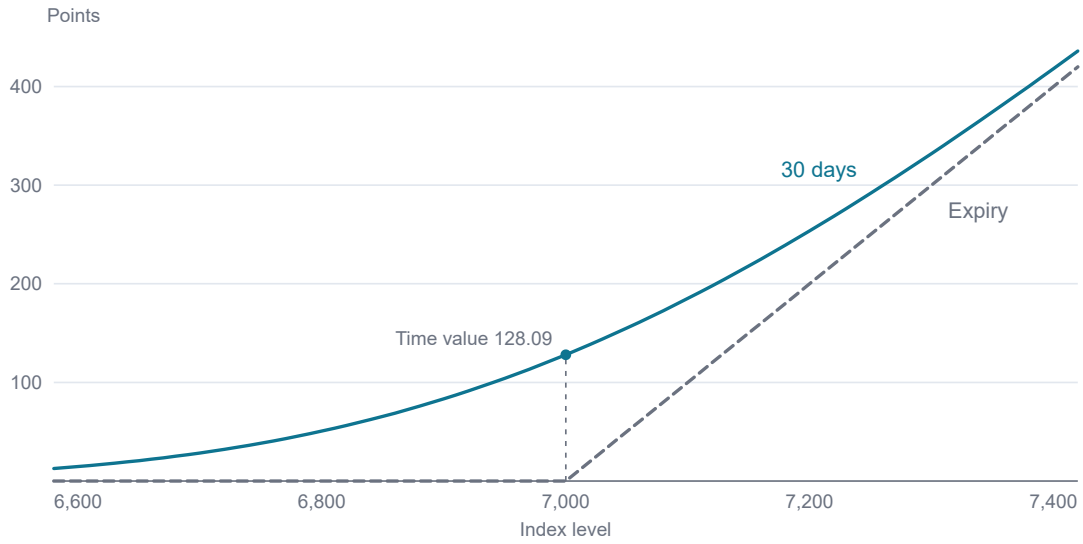


Fig. 7: Price of a call with strike 7,000 at 30 days to expiry (solid) and its payoff on expiry day (dashed), by index level. The distance between the two lines is the time value. It is largest at the strike.

6. Put-call parity: derivation

Call and put with the same strike and expiry are not independent products. The difference between their prices is fixed. The derivation needs only two portfolios and a comparison.

Portfolio A: a call with strike K, plus enough money in the account that with interest it amounts to exactly K on expiry day. Today that is the discounted amount $K \cdot e^{-rT}$.

Portfolio B: a put with the same strike, plus the index itself.

What are the two worth on expiry day?

On expiry day	Portfolio A: call and cash	Portfolio B: put and index
Index below K	call expires, cash is worth K: K	put pays $K - S_T$, index is worth S_T: K
Index above K	call pays $S_T - K$, cash is worth K: S_T	put expires, index is worth S_T: S_T

In both cases the portfolios are worth the same, namely the larger of the two values S_T and K. Two portfolios that are worth the same in every case at the end must cost the same today. Otherwise one could sell the dearer one, buy the cheaper one and keep the difference without risk.

$$C + K \cdot e^{-rT} = P + S$$

Rearranged:

$$C - P = S - K \cdot e^{-rT}$$

The difference between call and put price depends only on the index, the strike and the interest rate. Volatility does not appear, and neither does any view on direction.

Worked example without interest. In the examples of this series the interest rate is mostly zero. Then $C - P = S - K$. At the money ($S = K$) call and put cost the same: 128.09 and 128.09. At strike 6,900 the call is exactly 100 points dearer than the put: 183.38 against 83.38.

Worked example with interest. At 4 % interest and 30 days to expiry the discounted strike is

$$7{,}000 \times e^{-0.04 \times 30/365} = 6{,}977.02$$

So the call has to be $7{,}000 - 6{,}977.02 = 22.98$ points dearer than the put. Computed with the pricing formula the call costs 139.69 and the put 116.72. The difference is 22.98 (the prices are rounded to two decimals).

What happens if prices deviate? Suppose the call trades 5 points too high. An arbitrageur rebuilds the expensive side and sells it:

Trade today	Cash flow today
Sell call	+144.69
Buy put	−116.72
Buy index	−7,000.00
Take out loan	+6,977.02
Sum	+5.00 (rounded)

On expiry day everything cancels out, wherever the index is:

Index at expiry	Call sold	Put bought	Index	Loan repaid	Sum
6,800	0	+200	+6,800	−7,000	0
7,200	−200	0	+7,200	−7,000	0

The 5 points, that is 500 USD per contract, remain as a certain profit. Because anyone can do this, such deviations do not last.

7. Synthetic positions

Parity can be solved for each of its four quantities. So every position can be rebuilt from the other three. Such rebuilds are called synthetic. Without interest:

This position	can be rebuilt from
Index bought	call bought and put sold
Call bought	put bought and index bought
Put bought	call bought and index sold
Put sold	index bought and call sold

The first row says: whoever buys a call and sells a put with the same strike has the same position as with a future. If the index rises, the call gains. If it falls, the put sold loses. The right and the obligation together again make the contract that binds both sides.

The last row is the most important in practice. Holding shares and selling calls on them is called a **covered call** and is regarded as a cautious strategy for extra income. By parity it is the same position as a put sold: limited gain, and the full loss if the market falls.

For the further chapters one consequence matters above all: a call and a put with the same strike differ only by a position in the index. Everything that distinguishes an option from a future, the curvature of its payoff, is the same for both. That is why in chapters 8 and 9 it does not matter whether a dealer has calls or puts in his book. What counts is whether he has bought or sold them.

Common misconception

Common misconception: "A call is a bet on rising prices, a put a bet on falling ones." For direction no option is needed, a future is enough. By parity a call is a put plus a position in the index. So the difference between the two is only the direction. What an option is really about is the kink in its payoff: limited loss with open gain. That is what the premium is paid for, and that is why the buyer earns from large moves and not merely from having picked the direction.

What you should understand now

- A call is the right to buy, a put the right to sell, each at the strike. The buyer pays the premium, the seller carries the obligation.
- At expiry a call pays $\max(S_T - K, 0)$ and a put $\max(K - S_T, 0)$. The buyer can lose the premium at most and only earns beyond the break-even.
- In hedging, a future takes both sides away and costs nothing. A put takes only the bad side away and costs the premium.
- The price of an option consists of intrinsic value and time value. Time value is largest at the money and the same for call and put with the same strike.
- Put-call parity $C - P = S - K \cdot e^{-rT}$ ties the prices to each other. Every position can be rebuilt from the others.

Limits of this account

- **European exercise.** Parity holds in this form for options that can only be exercised on expiry day, and without dividends. With dividends the price reduced by the dividends takes the place of S. For American options it holds only approximately.
- **Trading costs.** Every option has a spread between bid and ask (chapter 2). The arbitrage from section 6 only pays once the deviation is larger than these costs.
- **Collateral.** Whoever sells options has to post margin, as with the future. The seller's unlimited risk of loss is visible in the tables only as a word.
- **Not covered:** combinations of several options such as spreads, straddles and strangles.

Continue with chapter 6: how the price of an option comes about, from a coin toss to the Black-Scholes formula.

Chapter 6

The option price: from a coin toss to the Black-Scholes formula

In one sentence: An option costs as much as the portfolio of index and loan that pays out exactly the same in the end, and because that portfolio is built without any assumption about the direction of the market, the option price contains no view on it, only an assumption about fluctuation.

What this chapter is: a derivation with worked examples. The transition from the tree to the formula is shown, not strictly proven. The numbers on the skew are example values. The numbers on the volatility premium come from our own measurement and are linked. **Not a trading recommendation.**

Chapter 5 left open why the call with strike 7,000 costs exactly 128.09 points. Here that price is derived.

1. The obvious calculation, and why it is wrong

The index is at 7,000. To simplify, there are only two possibilities in one month: it is at 7,350 (plus 5 %) or at 6,650 (minus 5 %). The interest rate is zero. A call with strike 7,000 then pays 350 or nothing.

What is it worth today? The obvious answer is the expected value: payoff times probability.

- Whoever regards both cases as equally likely computes $0.5 \times 350 = 175$.
- Whoever reckons with a rise at 70 % arrives at $0.7 \times 350 = 245$.

On that basis everyone would have his own price, depending on his view. In fact there is only one correct price, and to find it one does not need to know the probability at all.

2. The replication: a portfolio that pays the same

The idea: from two things whose price we know we build a portfolio that is worth exactly as much as the call in both cases. The two things are the index itself and a loan.

The portfolio consists of Δ units of the index, financed partly with a loan of size B. It should be worth as much as the call in both cases:

$$\begin{aligned} \text{Index rises:} \quad & \Delta \times 7{,}350 - B = 350 \\ \text{Index falls:} \quad & \Delta \times 6{,}650 - B = 0 \end{aligned}$$

Step 1: subtract the second equation from the first. The loan drops out:

$$\Delta \times (7{,}350 - 6{,}650) = 350 \quad \Longrightarrow \quad \Delta = \frac{350}{700} = 0.5$$

Step 2: determine the loan from the second equation.

$$B = 0.5 \times 6{,}650 = 3{,}325$$

Step 3: what does this portfolio cost today? Half an index unit costs 3,500, of which 3,325 is borrowed:

$$0.5 \times 7{,}000 - 3{,}325 = 175$$

The check. If the index rises, the half unit is worth 3,675, less 3,325 of loan leaves 350. If it falls, it is worth 3,325, less the loan leaves zero. The portfolio pays exactly like the call in both cases.

Two things that pay the same in every case must cost the same today. The call is worth 175.

3. Why it has to be this price

Suppose the call trades at 200 because many believe in rising prices. A dealer sells the call and rebuilds it for 175:

Trade today	Cash flow today
Sell call	+200
Buy 0.5 index units	−3,500
Take out loan	+3,325
Sum	+25

In the end his portfolio offsets the obligation from the call exactly in both cases. The 25 points remain with him, without risk. This works until the price has fallen to 175. If the call is below 175, it runs the other way.

The probability of a rise did not appear in this calculation. Whether it is 50 % or 70 % does not change the price. The reason: the market's view on direction is already in the index level of 7,000. Whoever is more optimistic than the market should buy the index. The option adds nothing to that.

What does determine the price is the **distance** between the two possibilities. If they were at plus and minus 10 %, the call would be twice as expensive. The option price depends on fluctuation, not on direction.

4. The general formula

The same calculation with letters. The index is at S and jumps to S_u or S_d. The call is then worth C_u or C_d.

$$\Delta \cdot S_u - B = C_u$$
$$\Delta \cdot S_d - B = C_d$$

Subtracting gives the number of units:

$$\Delta = \frac{C_u - C_d}{S_u - S_d}$$

That is the **delta**: the difference of the option values divided by the difference of the index levels. It says how much index is needed to rebuild the option. In chapter 8 it becomes the dealer's hedge.

The loan follows from the second equation: $B = \Delta \cdot S_d - C_d$. The price of the portfolio today is $\Delta \cdot S - B$. Substituted:

$$C = \Delta \cdot (S - S_d) + C_d = \frac{S - S_d}{S_u - S_d} \cdot (C_u - C_d) + C_d$$

Calling the fraction q gives a simple form:

$$C = q \cdot C_u + (1 - q) \cdot C_d, \qquad q = \frac{S - S_d}{S_u - S_d}$$

In the example $q = (7{,}000 - 6{,}650)/(7{,}350 - 6{,}650) = 0.5$ and $C = 0.5 \times 350 = 175$.

The formula looks like an expected value with the probability q. But q is not a forecast. It is a computed quantity that follows from the index levels alone: the probability at which the index itself earns nothing on average. It is called the **risk-neutral probability**. With an interest rate r little changes: the numerator of q then contains $S \cdot e^{rT}$ instead of S, and the result is discounted with e^{-rT}.

5. Two steps: the replication has to be adjusted

A single jump is too coarse. With two steps the picture already becomes richer. The index rises or falls per step by the factor 1.035 or 1/1.035:

	Start	after step 1	after step 2	Call payoff
risen twice			7,498.57	498.57
risen once		7,245.00		
middle	7,000		7,000.00	0
fallen once		6,763.29		
fallen twice			6,534.57	0

The calculation runs from back to front, each node with the formula from section 4. Here $q = 0.4914$.

- **After a rise (7,245):** the call is worth $0.4914 \times 498.57 = 245.00$.
- **After a fall (6,763.29):** both following states pay nothing. The call is worth 0.
- **At the start:** $0.4914 \times 245.00 = 120.39$.

More important than the price is what happens to the delta:

Node	Delta
Start (7,000)	0.51
after a rise (7,245)	1.00
after a fall (6,763.29)	0.00

Whoever rebuilds the call holds 0.51 index units at the start. If the index rises, he has to increase to a whole unit. If the index falls, he has to sell everything. **The replication is not a one-off trade. It has to be adjusted after every move**, and always in the direction of the move: buy after a rise, sell after a fall.

That is the core of everything that follows in part III. A dealer who has sold a call rebuilds it in exactly this way to offset his risk. His trading in the index is thereby dictated by the price path.

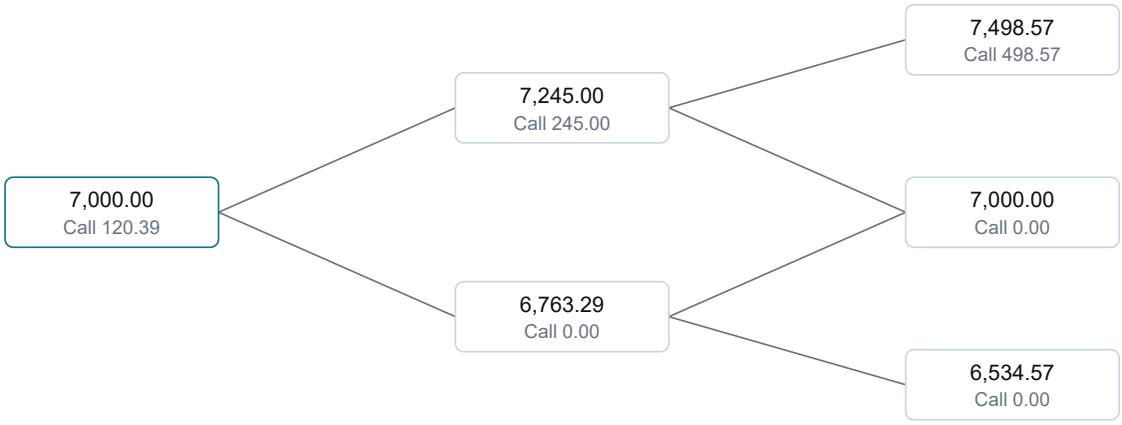


Fig. 8: The tree with two steps. At each node the index level is on top and below it the value of the call with strike 7,000.

6. Many steps

The more steps the tree has, the more finely it maps the possible paths of the index. The size of the steps is tied to the fluctuation one reckons with. For a volatility σ per year and steps of length Δt one chooses

$$u = e^{\sigma\sqrt{\Delta t}}, \qquad d = \frac{1}{u}$$

Smaller steps in time thus get smaller jumps in price, in the ratio of the square root. That is the same root as in chapter 3: fluctuation grows with the square root of time.

For the call from chapter 5 (index 7,000, strike 7,000, 30 days, 16 % volatility, zero interest) the tree gives:

Number of steps	Price of the call
1	160.52
2	113.51
5	134.61
10	124.93
50	127.45
100	127.77
1,000	128.05

The price settles down. If the number of steps goes to infinity, it lands at 128.09. That is the price from chapter 5.

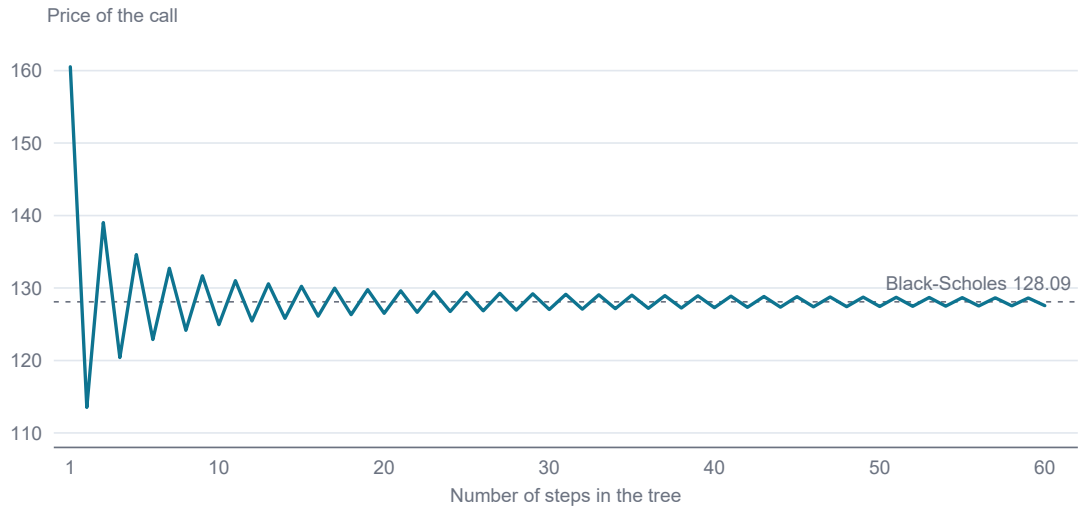


Fig. 9: Price of the call in the tree for 1 to 60 steps. The horizontal line is the limit 128.09 that the Black-Scholes formula gives.

7. The Black-Scholes formula

For the limiting case of infinitely many steps there is a closed formula. Fischer Black, Myron Scholes and Robert Merton published it in 1973. For a call:

$$C = S \cdot N(d_1) - K \cdot e^{-rT} \cdot N(d_2)$$

with

$$d_1 = \frac{\ln(S/K) + \left(r + \tfrac{1}{2}\sigma^2\right) T}{\sigma\sqrt{T}}, \qquad d_2 = d_1 - \sigma\sqrt{T}$$

Every quantity has a meaning:

- $N(\cdot)$ is the distribution function of the normal distribution. It gives a probability between 0 and 1 for every number.
- $\sigma\sqrt{T}$ is the fluctuation over the whole term. At 16 % per year and 30 days that is $0.16 \times \sqrt{30/365} = 4.59\ \%$.
- $N(d_2)$ is the risk-neutral probability that the option ends in the money. So $K \cdot e^{-rT} \cdot N(d_2)$ is today's value of the strike that is paid on exercise.
- $S \cdot N(d_1)$ is today's value of what is received on exercise.
- $N(d_1)$ is at the same time the **delta**: the number of index units in the replication, as in the tree.

So the formula is the same as in the tree: number of units times index minus loan. Only instead of two states there are infinitely many behind it.

Worked example. Index 7,000, strike 7,000, 30 days, 16 % volatility, zero interest:

$$\begin{aligned} \sigma\sqrt{T} &= 0.16 \times \sqrt{30/365} = 0.04587 \\ d_1 &= \frac{\ln(1) + \tfrac{1}{2} \times 0.16^2 \times 30/365}{0.04587} = 0.02294 \\ d_2 &= 0.02294 - 0.04587 = -0.02294 \\ N(d_1) &= 0.50915, \qquad N(d_2) = 0.49085 \\ C &= 7{,}000 \times 0.50915 - 7{,}000 \times 0.49085 = 128.09 \end{aligned}$$

The price of the put follows from put-call parity in chapter 5.

What the formula assumes: a volatility that does not change, a price without jumps, continuous adjustment of the replication at no cost. None of these assumptions holds exactly. Sections 9 and 10 show how the market deals with that.

8. Implied volatility

Of the five quantities in the formula four are known: index, strike, term and interest rate. The fifth, the volatility until expiry, nobody knows.

Dealers therefore turn the formula round. They do not ask: "What price follows from the volatility?" They ask: "What volatility do I have to put in so that the price comes out at which the option is trading right now?" That number is called **implied volatility**.

Worked example. The call does not trade at 128.09 but at 140. Substituting and trying until it fits gives an implied volatility of 17.49 %. As a rule of thumb: one volatility point more makes this option about 8 points dearer. So 12 points more in price are about 1.5 volatility points.

Implied volatility is thus another way of writing the price. It makes options comparable: 140 points for a call say little, 17.5 % volatility can be compared with yesterday, with other strikes and with other markets. The best-known value of this kind is the VIX, the implied volatility of the S&P 500 for the next 30 days.

9. The skew: insurance costs extra

If the formula were exactly right, all options on the same index with the same expiry would have to have the same implied volatility. That is not the case. For stock indices the following has held since the crash of 1987:

- Puts with a low strike, which protect against sharp losses, have a **higher** implied volatility.
- Calls with a high strike have a **lower** one.

This slope is called the **skew**. What it means in points is shown by an example with assumed volatility values:

Option	Price at 16 % volatility	Assumed market volatility	Market price
Put, strike 6,650 (5 % below the index)	20.71	20 %	39.67
Call, strike 7,350 (5 % above the index)	24.24	13 %	11.94

Protection against a fall of more than 5 % here costs almost twice what the formula gives with a uniform volatility. Two reasons work together:

- **Demand.** Funds and pension funds permanently buy puts as protection. Whoever sells them carries a risk that is hard to pass on, and demands more for it.
- **The assumptions fit worse to the downside.** Markets fall faster than they rise, sometimes in jumps, and fluctuation increases in sell-offs. The replication from section 5 then works worse than the formula assumes.

10. The volatility premium: implied lies above realised

Implied volatility is a price, and a price for insurance contains a mark-up. That is why implied volatility lies on average above the fluctuation that then actually occurs.

We have measured that. For 2015 to 2022 the fluctuation of the preceding 20 days was in the median only 0.76 (S&P 500), 0.84 (Nasdaq 100) and 0.74 (Dow) of the implied volatility. One consequence: the daily range priced in by the option market held the closing price on 80.1 % of days. With a normal distribution and no mark-up 68.3 % would be expected (to the study).

This difference is called the **volatility risk premium**. It is the pay of those who sell options, and it is not free money. The seller regularly takes in small amounts in quiet times and loses a great deal at once in a crash. How this shows up day by day in a dealer's book is worked through in chapter 8.

Common misconception

Common misconception: "Implied volatility is the market's forecast of the coming fluctuation." It is a price, converted into a volatility. That price contains the expected fluctuation and in addition a mark-up for the risk the seller carries. That is why it lies on average above the fluctuation that then occurs. Whoever concludes from a high VIX that the market "expects" exactly this move systematically overestimates it.

What you should understand now

- The price of an option is the price of the portfolio of index and loan that replicates its payoff. The probability of a rise does not appear in it.
- The number of index units in the replication is the delta: $\Delta = (C_u - C_d)/(S_u - S_d)$. It changes with the price, so the replication has to be adjusted continuously.
- With more and more steps the tree turns into the Black-Scholes formula. Its only unknown quantity is volatility.
- Implied volatility is the market price converted into a volatility. It differs by strike. Downside protection is dearer (skew).
- Implied volatility lies on average above realised volatility. Option sellers earn this premium for carrying crash risk.

Limits of this account

- **No proof.** That the tree converges to the formula is shown with numbers, not proven mathematically. The proof needs tools that are not assumed here.
- **Model assumptions.** Constant volatility, no jumps, free and continuous trading, no interest in the worked example. The skew is the market's visible correction for these assumptions.
- **Example values.** The volatility values in section 9 (20 % and 13 %) are chosen for illustration, not measured. How steep the skew is changes all the time.
- **Own measurement.** The ratio of realised to implied volatility comes from 2015 to 2022 and compares with the fluctuation of the 20 days before.
- **Not covered:** options with early exercise, dividends in the tree and models with changing volatility.

Continue with chapter 7: the Greeks, that is, how strongly the option price reacts to price, time and volatility.

Chapter 7

The Greeks: delta, gamma, theta, vega, vanna and charm

In one sentence: The Greeks say by how much the price of an option changes when the underlying price, time or volatility changes, and two of them, vanna and charm, say how the delta shifts in the process, that is, the dealer's hedge.

What this chapter is: a reference chapter with formulas and worked examples. All values are computed with the Black-Scholes formula (S&P 500 at 7,000, 30 days, 16 % volatility, zero interest), not observed in the market. Every formula is checked against a direct recomputation of the price. **Not a trading recommendation.**

Chapter 6 derived the price of an option. But a dealer wants to know more than what his options are worth today. He wants to know what happens to their value if something changes tomorrow.

1. What the sensitivities are for

The price of a call depends on three quantities that change all the time: the index level S, the time to expiry T and the implied volatility σ. For each there is a number that says how sensitively the price reacts to it. They carry Greek letters and are therefore called the Greeks.

Greek	Question it answers
Delta	By how much does the price change when the index rises by one point?
Gamma	By how much does the delta change in the process?
Theta	How much does the option lose in one day if nothing else happens?
Vega	By how much does the price change when volatility rises by one point?
Vanna	By how much does the delta change when volatility rises by one point?
Charm	By how much does the delta change when one day passes?

The first four describe the price. The last two describe the delta. They matter because the delta determines the dealer's hedge (chapter 8).

The example for all sections is three calls with 30 days to expiry at an index level of 7,000:

Strike	Position	Price in points
6,800	in the money	250.65
7,000	at the money	128.09
7,200	out of the money	53.67

In all formulas d_1 and d_2 stand for the two quantities from chapter 6, $N(\cdot)$ for the distribution function of the normal distribution and $\varphi(\cdot)$ for its density, the bell curve.

2. Delta: how closely the option follows the index

Meaning. Delta is the price change of the option per index point. A delta of 0.51 means: if the index rises by one point, the call rises by 0.51 points. At the same time delta is the number of index units with which the option can be replicated (chapter 6).

Formula.

$$\Delta_{\text{call}} = N(d_1), \qquad \Delta_{\text{put}} = N(d_1) - 1$$

The delta of a call lies between 0 and 1, that of a put between −1 and 0. The difference is always exactly 1. That is put-call parity from chapter 5 in another form.

Numbers.

Strike	Call delta	Put delta
6,800 (in the money)	0.744	−0.256
7,000 (at the money)	0.509	−0.491
7,200 (out of the money)	0.277	−0.723

A call deep in the money behaves almost like the index itself (delta close to 1). A call far out of the money hardly reacts (delta close to 0). At the money the delta is around one half.

3. Gamma: how fast the delta changes

Meaning. The delta is not fixed. If the index rises, a call becomes more likely to end with value, and its delta rises. Gamma measures by how much: the change in delta per index point. Gamma is thus the curvature of the price curve, that is, what distinguishes an option from a future. A future has delta 1 and gamma 0.

Formula.

$$\Gamma = \frac{\varphi(d_1)}{S \cdot \sigma \cdot \sqrt{T}}$$

Call and put with the same strike have the same gamma.

Numbers. At the money $\Gamma = 0.00124$ per point. Per 100 points that is 0.124. So if the index rises from 7,000 to 7,100, the delta of the call rises from 0.509 to around 0.63. The exact recomputation gives 0.630.

Two properties determine where gamma is large:

- **It is largest at the money.** There every point decides whether the option is worth something in the end. At 6,800 and 7,200 it is lower, at 0.00100 and 0.00104.
- **It grows towards expiry.** The denominator contains $\sqrt{T}$. The shorter the time to expiry, the larger the gamma at the money.

Time to expiry	Gamma at the money	Ratio to 30 days
30 days	0.00124	1.0
5 days	0.00304	2.4
1 day	0.00681	5.5

With one day to expiry the gamma at the money is 5.5 times as large as with 30 days. That is $\sqrt{30}$. An option that expires today has an even much larger gamma in the afternoon: its delta jumps between 0 and 1 as soon as the index crosses the strike. That is why options with very short terms play a role of their own in chapter 9.

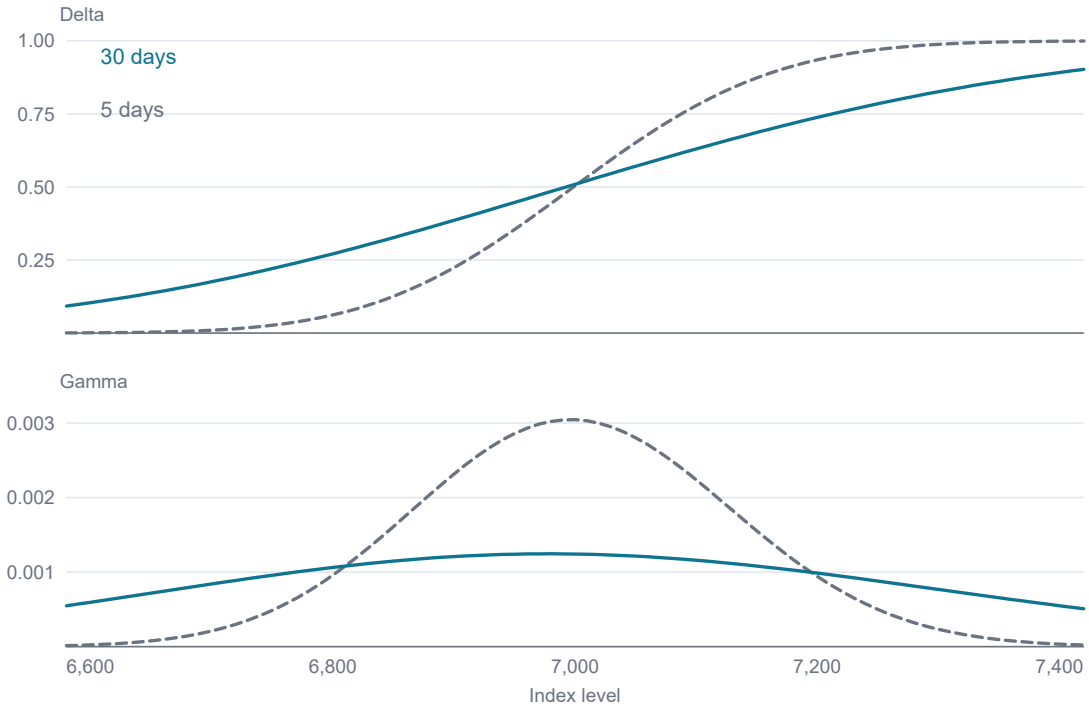


Fig. 10: Delta (top) and gamma (bottom) of a call with strike 7,000 by index level, at 30 days to expiry (solid) and at 5 days (dashed). Shortly before expiry the delta curve becomes steeper and gamma concentrates at the strike.

4. Theta: what a day costs

Meaning. Theta is the amount the option loses in one day if index and volatility do not change. It is the time value from chapter 5, melting away day by day.

Formula (without interest, per year; divide by 365 for one day).

$$\Theta = -\frac{S \cdot \varphi(d_1) \cdot \sigma}{2\sqrt{T}}$$

Numbers. At the money that is −779 points per year, that is −2.13 points per day. One contract loses 213 USD a day.

Time value does not melt evenly. It depends on $\sqrt{T}$, and a square root falls most steeply at the end:

Time to expiry	Price at the money	Theta per day
30 days	128.09	−2.13
5 days	52.30	−5.23
1 day	23.39	−11.69

In the first 25 days the option loses 76 points, in the last 5 days 52. The last day alone costs 23 points.

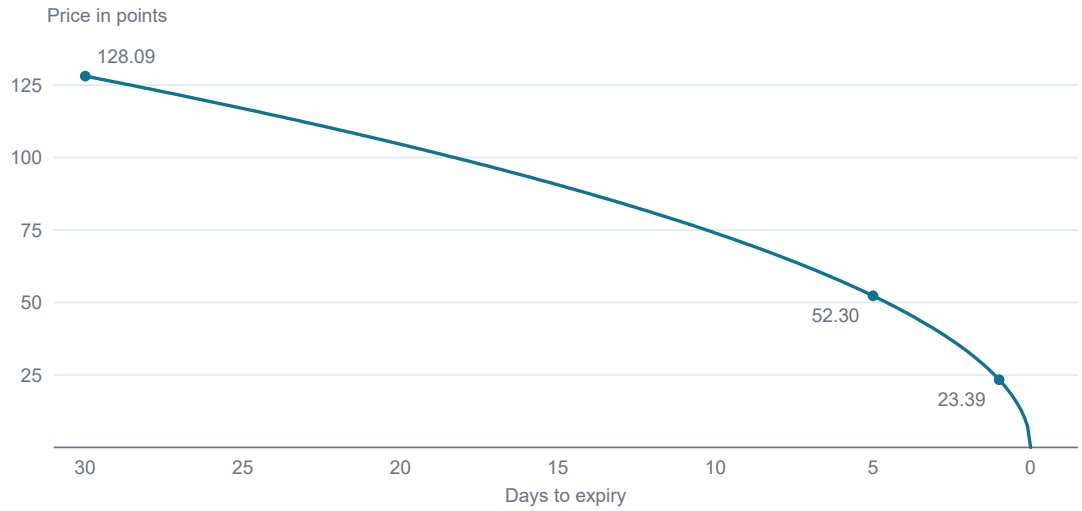


Fig. 11: Price of an at-the-money call with the index unchanged at 7,000 and 16 % volatility, from 30 days to expiry down to expiry. Time value melts faster the closer expiry comes.

5. Gamma and theta belong together

Gamma and theta are largest at the money and both grow towards expiry. That is no coincidence. If the formula for theta is divided by the one for gamma, almost everything drops out:

$$\frac{\Theta}{\Gamma} = -\frac{S \cdot \varphi(d_1) \cdot \sigma}{2\sqrt{T}} \cdot \frac{S \cdot \sigma \cdot \sqrt{T}}{\varphi(d_1)} = -\tfrac{1}{2}\,\sigma^2 S^2$$

So:

$$\Theta = -\tfrac{1}{2}\,\sigma^2 S^2\,\Gamma$$

The check with the numbers at the money:

$$\tfrac{1}{2} \times 0.16^2 \times 7{,}000^2 \times 0.00124 = 779 \text{ points per year}$$

Whoever holds gamma pays theta for it, in a fixed ratio. There is no such thing as a lot of gamma without a lot of theta. Here that is a rearrangement of two formulas. What lies behind it economically is shown in chapter 8: gamma brings money to whoever hedges his option on every price move, and theta is the price for that.

6. Vega: what a volatility point is worth

Meaning. Vega is the price change of the option when implied volatility rises by one point, for example from 16 % to 17 %. Whoever holds options earns from rising volatility. Whoever has sold them loses.

Formula (per 1.00 of volatility; divide by 100 for one volatility point).

$$\text{Vega} = S \cdot \varphi(d_1) \cdot \sqrt{T}$$

Numbers. At the money that is 8.00 points per volatility point, that is 800 USD per contract. This is the rule of thumb from chapter 6.

Vega behaves in the opposite way to gamma with respect to the term. In the formula $\sqrt{T}$ is in the numerator:

Time to expiry	Vega per volatility point	Gamma at the money
30 days	8.00	0.00124
5 days	3.27	0.00304
1 day	1.46	0.00681

Long-dated options react strongly to volatility and little to the individual price step. Short-dated ones react strongly to the price and hardly to volatility.

7. Vanna: how volatility shifts the delta

Meaning. Vanna is the change in delta when implied volatility rises by one point. The index does not move. The delta changes all the same, and with it the hedge a dealer needs.

Formula (per 1.00 of volatility; divide by 100 for one volatility point).

$$\text{Vanna} = -\frac{\varphi(d_1) \cdot d_2}{\sigma}$$

Numbers.

Strike	Delta at 16 % volatility	Vanna per volatility point	Delta at 17 %	Delta at 15 %
6,800 (in the money)	0.744	−0.0123	0.732	0.757
7,000 (at the money)	0.509	+0.0006	0.510	0.509
7,200 (out of the money)	0.277	+0.0133	0.290	0.263

The intuition behind it: more volatility means the index can travel further by expiry. For a call out of the money this makes it more likely to still get into the money. Its delta rises. For a call in the money it becomes less certain to stay there. Its delta falls. Rising volatility pulls all deltas towards the middle, falling volatility pushes them to the edges.

Why this matters: in quiet, rising markets implied volatility usually falls. The deltas of the options out of the money then shrink without the price having to move for it. Whoever has hedged such options has to adjust his hedge. Chapter 9 shows in which direction.

8. Charm: how time shifts the delta

Meaning. Charm is the change in delta when one day passes. Index and volatility stay the same.

Formula (without interest, per year; divide by 365 for one day).

$$\text{Charm} = \frac{\varphi(d_1) \cdot d_2}{2\,T}$$

Numbers.

Strike	Delta today	Charm per day	Delta tomorrow
6,800 (in the money)	0.744	+0.0033	0.747
7,000 (at the money)	0.509	−0.0002	0.509
7,200 (out of the money)	0.277	−0.0036	0.274

The intuition: on expiry day the delta is either 1 (in the money) or 0 (out of the money). It moves a little towards that with every day. A call out of the money loses delta, a call in the money gains it. Passing time acts on the delta like falling volatility: both push it to the edges.

Like gamma and theta, charm accelerates towards expiry, because T is in the denominator. In the last days before a large expiry the deltas therefore shift fastest, even if the price stands still.

9. All Greeks at a glance

The at-the-money call (strike 7,000, 30 days, 16 % volatility), per contract with 100 USD per point:

Greek	Value	Meaning in USD per contract
Price	128.09 points	12,809 USD
Delta	0.509	+50.90 USD per index point
Gamma	0.00124 per point	delta +0.124 per 100 points
Theta	−2.13 points per day	−213 USD per day
Vega	8.00 points per volatility point	+800 USD per volatility point
Vanna	+0.0006 per volatility point	almost zero at the money
Charm	−0.0002 per day	almost zero at the money

Vanna and charm are small at the money and large away from it. For gamma, theta and vega it is the other way round.

Common misconception

Common misconception: "Delta is the probability that the option ends in the money." As a rough estimate that will do, exactly right it is not. Delta is $N(d_1)$, the probability from the formula is $N(d_2)$. For the call with strike 7,200 these are 0.277 and 0.262. And $N(d_2)$ is not a prediction either: it is the risk-neutral computed quantity from chapter 6, which contains no view on direction. Delta is a hedging quantity, not a forecast.

What you should understand now

- Delta is the price change per index point and at the same time the number of index units in the replication. For a call it lies between 0 and 1, at the money around 0.5.
- Gamma is the change in delta per point. It is largest at the money and grows towards expiry with $1/\sqrt{T}$.
- Theta is the loss of time value per day. It accelerates towards expiry. Gamma and theta are firmly linked: $\Theta = -\frac{1}{2}\,\sigma^2 S^2\,\Gamma$.
- Vega is the price change per volatility point. It is large for long terms, gamma for short ones.
- Vanna and charm describe how volatility and time shift the delta. Falling volatility and passing time push the delta of out-of-the-money options towards zero.

Limits of this account

- **Small changes.** Each Greek holds for a small change in a single quantity. In reality price, time and volatility change at the same time, and all Greeks change with them.
- **Model.** The formulas assume Black-Scholes with a uniform volatility. With skew (chapter 6) volatility itself depends on the price, and the delta dealers actually use deviates from $N(d_1)$.
- **Without interest.** Theta and charm are given without interest and dividends. With interest small additional terms appear, the mechanics do not change.
- **Not covered:** rho (the effect of the interest rate) and further higher-order sensitivities.

Continue with chapter 8: what a dealer does with these numbers when he hedges his option book in the index.

Chapter 8

Hedging in the underlying: why option dealers with long gamma trade against the move and with short gamma trade with it

In one sentence: Whoever holds an option and offsets its price risk with futures has to trade again after every move, and whether he trades against the move or with it depends solely on the sign of his gamma.

What this chapter is: an explanation with a worked example, not a measurement. All numbers are computed (Black-Scholes, zero interest and dividends), not observed in the market. What of this mechanism shows up in real prices is in chapter 10. **Not a trading recommendation.**

This chapter needs three quantities from chapter 7. **Delta** says by how many points the option price changes when the index rises by one point. **Gamma** says by how much the delta changes in the process. **Theta** says how much the option loses in one day if nothing else happens.

1. What a dealer wants: the spread, not the direction

An option dealer (market maker) quotes buying and selling prices and earns the spread in between. A view on market direction is not part of his business. But every option a customer buys from him or sells to him puts exactly such a direction into his book. He gets rid of it by building the opposite position in the underlying. Underlying here means: what the option is written on, in practice the future on the index.

The example for this chapter:

Quantity	Value
S&P 500	7,000 points
Strike of the call	7,000 points
Time to expiry	30 days
Implied volatility	16 %
Price of the call	128.09 points
Price per contract (1 point = 100 USD)	12,809 USD
Delta	0.509
Gamma per point	0.00124
Theta per calendar day	−2.13 points

A customer sells the dealer 100 of these calls, for example a fund that writes calls on its share holdings. The dealer is now long 100 calls.

2. The delta hedge, step by step

Step 1: how strongly does the book depend on the index? One call gains 0.509 points when the index rises by one point. One point is 100 USD, and there are 100 contracts:

$$\begin{aligned}\text{Book delta} &= 0.509 \times 100 \times 100 \text{ USD} \\ &= 5{,}091 \text{ USD per index point}\end{aligned}$$

Step 2: what does he offset it with? With the E-mini future on the S&P 500. One future moves by 50 USD per index point.

$$\text{Futures} = \frac{5{,}091}{50} = 101.8 \approx 102 \text{ contracts, sold}$$

Step 3: the check. If the index rises by one point, the calls gain 5,091 USD and the 102 futures sold lose 5,100 USD. The sum is practically zero. The book is delta-neutral: small moves no longer change its value.

What is hedged is 5,091 USD per point times 7,000 points, that is around 35.6 million USD of index notional, for options that together cost 1.28 million USD.

3. The price moves: the hedge no longer fits

The hedge is only right for the price at which it was computed. If the index rises by 1 % to 7,070, the delta of the call rises. By how much is given by gamma:

$$\begin{aligned}\Delta_{\text{new}} &\approx \Delta + \Gamma \cdot dS \\ &= 0.509 + 0.00124 \times 70 = 0.596\end{aligned}$$

The exact recomputation gives 0.595. The book delta is now 5,948 USD per point, and 119 futures sold are needed. The dealer has 102. He has to **sell 17 futures after the price has risen.**

If the index then falls back to 7,000, the delta falls back to 0.509. The dealer only needs 102 futures and **buys the 17 back after the price has fallen.** Sold at 7,070, bought back at 7,000:

$$17 \times 50 \text{ USD} \times 70 \text{ points} = 59{,}500 \text{ USD}$$

That is not trading skill. The dealer has only adjusted his hedge. The profit drops out because his delta runs along with the price: after a rise he is too long and has to sell, after a fall he is too short and has to buy. Selling high and buying back low is built into long gamma.

4. What the move brings in: ½ · gamma · move²

How large is this profit in general? If the index moves by dS points, the delta runs evenly from Δ to $\Delta + \Gamma \cdot dS$. So on average over the stretch it was

$$\bar{\Delta} = \Delta + \tfrac{1}{2}\,\Gamma \cdot dS$$

The option changes its value by average delta times distance:

$$dV \approx \bar{\Delta} \cdot dS = \Delta \cdot dS + \tfrac{1}{2}\,\Gamma \cdot dS^2$$

The hedge of futures sold brings $-\Delta \cdot dS$. The first part cancels. What remains:

$$\text{Gamma gain} = \tfrac{1}{2}\,\Gamma \cdot dS^2$$

Two things are in this formula. The move enters squared: whether the index rises or falls by 70 points is the same, and a move twice as large brings four times as much. And the sign comes from gamma alone.

For the example:

$$\begin{aligned} \tfrac{1}{2} \times 0.00124 \times 70^2 &= 3.04 \text{ points} \\ &= 304 \text{ USD per contract} \\ &= 30{,}432 \text{ USD for 100 contracts} \end{aligned}$$

The exact recomputation gives 30,163 USD. The difference is the error of the approximation, because gamma also changes slightly along the way. There and back it is twice around 30,400 USD. That fits the 59,500 USD from section 3, which are computed with whole futures.

5. What waiting costs: theta

If that were all, a hedged option would be a machine: every move brings money, whichever way. There is no such machine for free. The price for it is time decay. If the index stays at 7,000 for a day, each call loses 2.13 points, the book around 21,300 USD. Recomputed exactly it is 21,525 USD, because theta increases slightly in the course of the day.

How large this loss has to be can be derived. The implied volatility of 16 % is the fluctuation contained in the option price. It says how large the squared move is on average in a short period dt (in years):

$$\mathbb{E}\big[dS^2\big] = \sigma^2 S^2\, dt$$

A fair price requires that a hedged book earns nothing on average if the market fluctuates exactly as much as priced in. So the average gamma gain and the time decay have to cancel:

$$\begin{aligned} \tfrac{1}{2}\,\Gamma\,\sigma^2 S^2\, dt + \Theta\, dt &= 0 \\ \Theta &= -\tfrac{1}{2}\,\sigma^2 S^2\,\Gamma \end{aligned}$$

That is the relation between gamma and theta. Whoever holds gamma pays theta for it, and the more so the higher implied volatility is. The check with the example numbers:

$$\begin{aligned} \tfrac{1}{2} \times 0.16^2 \times 7{,}000^2 \times 0.00124 &= 779 \text{ points per year} \\ &= 2.13 \text{ points per day} \end{aligned}$$

The two together give the daily result of a hedged book:

$$\text{Result} \approx \tfrac{1}{2}\,\Gamma\,\big(dS^2 - \sigma^2 S^2\, dt\big)$$

The bracket contains a comparison: the move that actually came against the move that was in the price. A hedged option book is therefore not a bet on direction. It is a bet on whether the market fluctuates more or less than priced in, that is, realised against implied volatility.

The break-even is where the bracket is zero:

$$\begin{aligned} |dS| = \sigma\, S\, \sqrt{dt} &= 0.16 \times 7{,}000 \times \sqrt{1/365} \\ &= 58.6 \text{ points} = 0.84\ \% \end{aligned}$$

If the index moves by more than 0.84 % in a day, the long-gamma book earns, below that it loses. The calculation here is in calendar days, because the time to expiry is given in calendar days. Counting in trading days (252 per year), the threshold is at 70.6 points or 1.01 %. That is the same calculation with a different calendar and exactly the quantity known as the expected move.

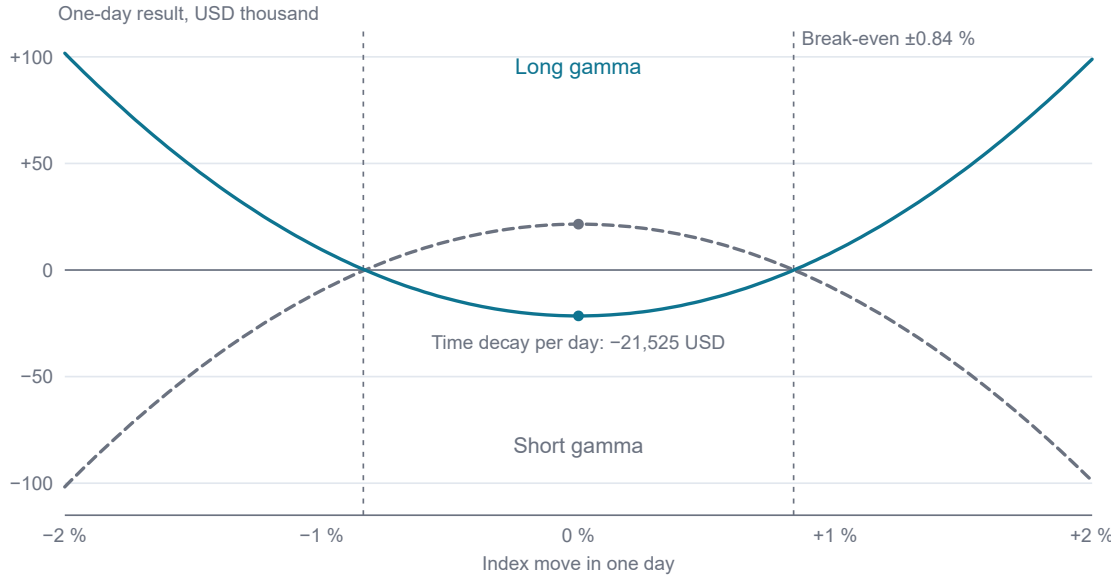


Fig. 12: Daily result of the hedged book of 100 calls by index move, recomputed exactly. Solid: the dealer holds the options (long gamma). Dashed: he has sold them (short gamma). Without a move the day costs or brings around 21,500 USD, the curves cross the zero line at ±0.84 %.

6. Five days worked through

The dealer hedges once a day at the closing price, with whole futures. The option is recomputed exactly every day.

Day	Index	Move	Delta	Futures sold	Trade	Options, USD	Futures, USD	Sum, USD
0	7,000		0.509	102	sells 102	0	0	0
1	7,070	+1.00 %	0.596	119	sells 17	+365,446	−357,000	+8,446
2	7,010	−0.85 %	0.522	104	buys 15	+8,102	0	+8,102
3	6,930	−1.14 %	0.417	83	buys 21	−389,697	+416,000	+26,303
4	6,990	+0.87 %	0.495	99	sells 16	−138,615	+167,000	+28,385
5	7,000	+0.14 %	0.508	102	sells 3	−111,584	+117,500	+5,916

The columns "Options" and "Futures" show the cumulative result since day 0.

Three things can be read off.

Every trade runs against the last move. After every rise he sells, after every fall he buys. The dealer does not decide that, his delta forces him to.

Large days earn, small days cost. Day 3 with −1.14 % brings around 18,200 USD, day 2 with −0.85 % lies almost exactly on the threshold and brings nothing, day 5 with +0.14 % costs around 22,500 USD.

In the end it is almost zero. After five days the index is back at 7,000, the options have lost 111,584 USD of time value, the futures have earned 117,500 USD. What remains is 5,916 USD. That is no coincidence: annualised, the five days fluctuated at 16.6 %, and 16 % was priced in. Realised and implied volatility were almost the same, so the book earns almost nothing.

The cross-check is a quiet week with the closing prices 7,015, 7,005, 6,990, 7,000 and 7,000. Annualised that is 3.1 % of fluctuation. The same book then loses 107,334 USD, almost exactly five times the daily time decay.

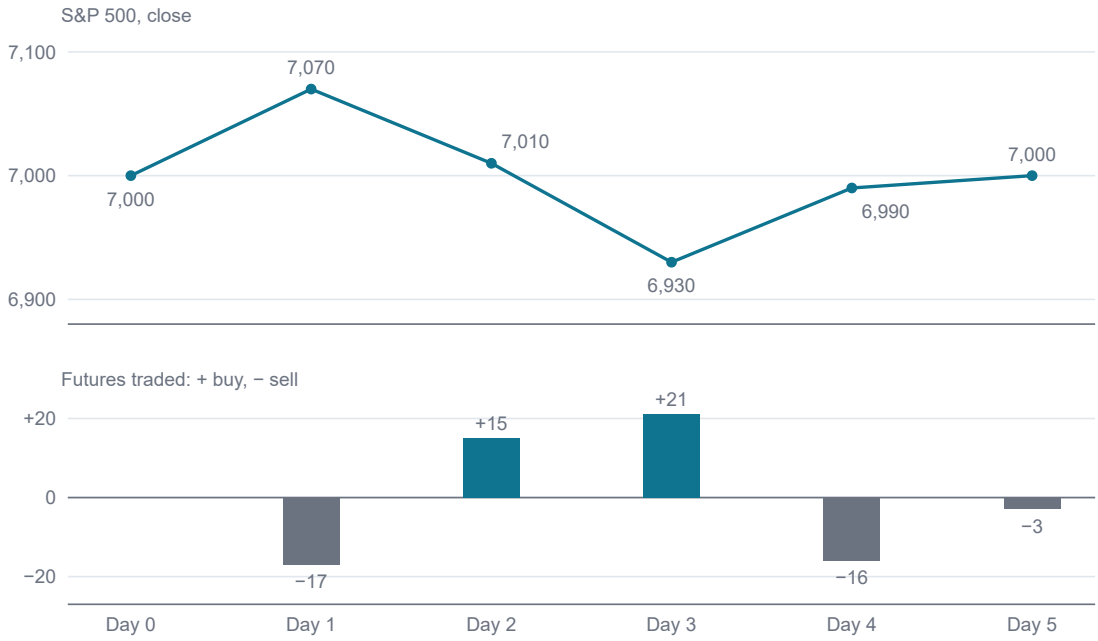


Fig. 13: On top the index path from the table, below it the futures the dealer trades at the close. Every bar points in the opposite direction to the day's move.

7. The other side: short gamma

Now a customer buys 100 puts with the same strike and the same expiry from the dealer, for example to protect his portfolio. The dealer has sold the puts and is short. Gamma and theta of these puts are as large as for the call, only the dealer carries them with the opposite sign.

A put here has a delta of −0.491. Puts sold have the opposite sign: the book gains 4,909 USD per index point to the upside. Against that the dealer sells 98 futures.

If the index falls by 1 % to 6,930, the put delta is −0.578. The book now depends on the index with 5,778 USD per point, and 116 futures sold are needed. The dealer has to **sell 18 futures after the price has fallen.** If the index rises again, he has to buy them back at a higher price.

Everything from sections 3 to 6 holds in mirror image:

	Long gamma	Short gamma
The dealer has	bought the options	sold the options
After a rise	he sells futures	he buys futures
After a fall	he buys futures	he sells futures
Effect on the price	dampening	amplifying
Move of 1 %	+30,400 USD	−30,400 USD
Day without a move	−21,300 USD	+21,300 USD
Earns if realised volatility	is above implied	is below implied

The amounts hold for the example book of 100 contracts and are rounded. Whether it is about calls or puts does not matter for this table. All that matters is whether the dealer holds the options or has sold them.

8. From the single book to the market

How much does a dealer have to trade after a move of 1 %? The move in points is $0.01 \cdot S$. The delta changes by $\Gamma \cdot 0.01 \cdot S$ per option. In USD of index notional, with N contracts and the multiplier m:

$$\text{Hedge flow per } 1\,\% = \Gamma \cdot S^2 \cdot 0.01 \cdot N \cdot m$$

For the example:

$$0.00124 \times 7{,}000^2 \times 0.01 \times 100 \times 100 = 6.09 \text{ million USD}$$
$$\approx 17 \text{ futures}$$

Those are the 17 futures from section 3: 6.09 million USD divided by the notional of one future of 350,000 USD (50 USD times 7,000 points). Summing this quantity over all strikes and expiries in which dealers hold positions gives the gamma exposure, GEX for short. Chapter 9 derives it and shows which assumptions it depends on. The most important one up front: the formula needs the sign of the dealer position, and exactly that is not publicly known.

The mechanics predict that a market in which dealers are short gamma overall has wider days. That can be measured. On a freely published gamma series for the S&P 500 the range of the next US session in the lower gamma tercile was 0.25 typical daily ranges larger than in the upper one (t 6.5), in the one-off holdout 0.20 (t 3.2). No direction could be derived from it (to the study).

Common misconception

> **Common misconception:** "Short gamma means the market falls." Gamma has no direction. A dealer with short gamma sells into falling prices and buys into rising ones. He amplifies the move that is running anyway, down as well as up. Short gamma is a statement about the width of a day, not about its sign.

A second misconception follows from the same confusion: that dealers "defend" a price level. The hedge is a reaction to a move that has already happened. Nobody decides in the process where the price should be.

What you should understand now

- A delta hedge makes an option book neutral only for small moves. After every larger move it has to be adjusted, and gamma determines the quantity needed.
- Whoever holds options (long gamma) sells after rises and buys after falls. Whoever has sold them (short gamma) does the opposite.
- A move of dS brings the long-gamma book $\frac{1}{2}\,\Gamma\,dS^2$, regardless of direction. Against that stands the daily time decay theta.
- Gamma and theta belong together: $\Theta = -\frac{1}{2}\,\sigma^2 S^2\,\Gamma$. The break-even of a day is the priced-in daily move $\sigma S\sqrt{dt}$.
- A hedged option book earns or loses on the difference between realised and implied volatility, not on market direction.

Limits of this account

- **Model.** The calculation uses Black-Scholes with constant volatility, without interest and dividends. In real markets implied volatility changes with the price, and with it the delta. This effect (vanna) and that of passing time on the delta (charm) are covered in chapter 7.
- **Hedging.** The dealer in the example hedges once a day and pays no costs. Real dealers hedge their whole book net, not every option individually, often only beyond a deviation and frequently with other options instead of futures. The flow in the underlying is therefore smaller and more irregular than the formula.
- **Approximation.** $\frac{1}{2}\,\Gamma\,dS^2$ holds for small moves. At 1 % it deviates from the exact value by just under 1 % in the example, for large jumps or shortly before expiry by considerably more.
- **Not shown:** who in the market actually holds which side. The example takes the dealer's position as known. In reality it is an estimate (chapter 9).

Continue with chapter 9: how this calculation is done for the whole market, and what open interest, GEX, the gamma flip and walls rest on.

Chapter 9

Dealer positioning: open interest, GEX, gamma flip and walls

In one sentence: GEX is the calculation from chapter 8 for the whole market, that is, how many USD option dealers have to trade per 1 % price move, and it stands and falls with an assumption that cannot be checked from public data: who bought the options and who sold them.

What this chapter is: a derivation with a worked example. The option chain in the example is made up. Where our own measurements appear, they are marked and linked. What of this shows up in real prices is collected in chapter 10. **Not a trading recommendation.**

Chapter 8 showed what a single dealer with 100 options has to trade in the index. This one is about all dealers and all options together.

1. The question

From chapter 8 we know: a dealer with long gamma trades against the move and dampens it. A dealer with short gamma trades with it and amplifies it. If one knew for the whole market whether dealers together are long or short gamma, and by how much, one would know something about the character of the coming trading day.

Three things are needed for that: how many options are open, who holds which side of them, and how much gamma each option has. The first is public. The third can be computed (chapter 7). The second is the weak point.

2. Open interest and volume

Volume counts how many contracts were traded today. **Open interest** (OI) counts how many contracts are open at the end of the day. The two are easily confused. Four trades in the same option show the difference:

Trade	Volume	Open interest afterwards
Anna buys 10 contracts new, Ben sells them new	10	10
Clara buys 5 new, David sells 5 new	5	15
Anna sells her 10 to Emil, who enters new	10	15
Clara and David close their 5 against each other	5	10

The day's volume is 30. Open interest is 10. Open interest only rises when both sides enter new, and only falls when both exit. If a position merely changes owner, it stays the same.

Two properties are decisive for everything that follows:

- **Open interest comes once a day.** The clearing houses compute it after the close and publish it the next morning. During the trading day the number does not change, whatever happens.
- **Open interest knows no sides.** Every open contract has a buyer and a seller. The number does not say which of the two is a dealer.

3. Who has which side?

GEX needs an answer to the question that open interest does not answer. The usual answer is a rule of thumb about the behaviour of customers:

- **Customers buy puts** to protect portfolios (the pension fund from chapter 5). So dealers have **sold** these puts.
- **Customers sell calls** on their holdings to take in premium (the covered call from chapter 5). So dealers have **bought** these calls.

From this comes the standard assumption: **dealers are long calls and short puts.** By chapter 8 this means: from calls they have long gamma, from puts short gamma.

That is an assumption, not a finding. Customers also buy calls and also sell puts. On some days and at some strikes the rule holds, on others it does not. Data on who actually opened which side exists at the exchanges, but not for free and not in the usual GEX tools.

4. GEX derived

Chapter 8 derived how much a dealer has to trade after a move of 1 %, in USD of index notional, for N contracts with multiplier m:

$$\text{Hedge flow per } 1\,\% = \Gamma \cdot S^2 \cdot 0.01 \cdot N \cdot m$$

Step 1: per strike. For every strike K the open interest of the calls and of the puts is known. Call and put with the same strike and expiry have the same gamma Γ_K (chapter 7). With the standard assumption calls count positive and puts negative:

$$\text{GEX}_K = \Gamma_K \cdot (\text{OI}_{\text{call}} - \text{OI}_{\text{put}}) \cdot m \cdot S^2 \cdot 0.01$$

Step 2: the sum. Over all strikes and expiries:

$$\text{GEX} = \sum_K \text{GEX}_K$$

The unit is USD per 1 % move. A GEX of +600 million USD means: if the index rises by 1 %, dealers by this calculation sell futures with a notional of 600 million USD. If it falls by 1 %, they buy that much. With negative GEX it is the other way round.

5. Worked example: a small option chain

The index is at 7,000. There is one expiry in 30 days, implied volatility is 16 % everywhere, the multiplier is 100 USD. Five strikes with made-up open interest:

Strike	OI calls	OI puts	Calls minus puts	Gamma	GEX in million USD per 1 %
6,800	5,000	30,000	−25,000	0.00100	−1,228
6,900	8,000	22,000	−14,000	0.00117	−805
7,000	20,000	20,000	0	0.00124	0
7,100	26,000	6,000	+20,000	0.00119	+1,169
7,200	32,000	3,000	+29,000	0.00104	+1,482
Sum					**+618**

The calculation for the 7,100 row:

$$0.001193 \times 20{,}000 \times 100 \times 7{,}000^2 \times 0.01 = 1.169 \text{ billion USD}$$

The picture is typical: below the price there are many puts, above it many calls. In total the dealers in this example are long gamma, with 618 million USD per 1 %. That corresponds to around 1,760 E-mini futures that would be sold after a rise of 1 % and bought after a fall of 1 %.

Two well-known levels can be read from the table:

- **Call wall:** the strike with the largest positive GEX. Here 7,200.
- **Put wall:** the strike with the largest negative GEX. Here 6,800.

6. The profile and the gamma flip

The sum of +618 million USD holds only at an index level of 7,000. Open interest stays the same, but the gamma of every option depends on the distance of the price from its strike (chapter 7). If the index falls, the puts at 6,800 and 6,900 move closer to the money and their gamma grows, while the calls at 7,100 and 7,200 move away. Recomputing the sum for every conceivable index level gives a profile:

Index level	GEX in million USD per 1 %
6,700	−1,016
6,800	−683
6,900	−97
7,000	+618
7,100	+1,284
7,200	+1,745

Between 6,900 and 7,000 the sign changes. The exact zero crossing is at 6,914. This point is called the **gamma flip**. The usual reading: above the flip dealers are long gamma and dampen moves, below it they are short gamma and amplify them.

In the example the flip is 86 points, that is 1.2 %, below the price. If the index falls to it, nothing has changed in the positions. All that has changed is which options are close to the money.

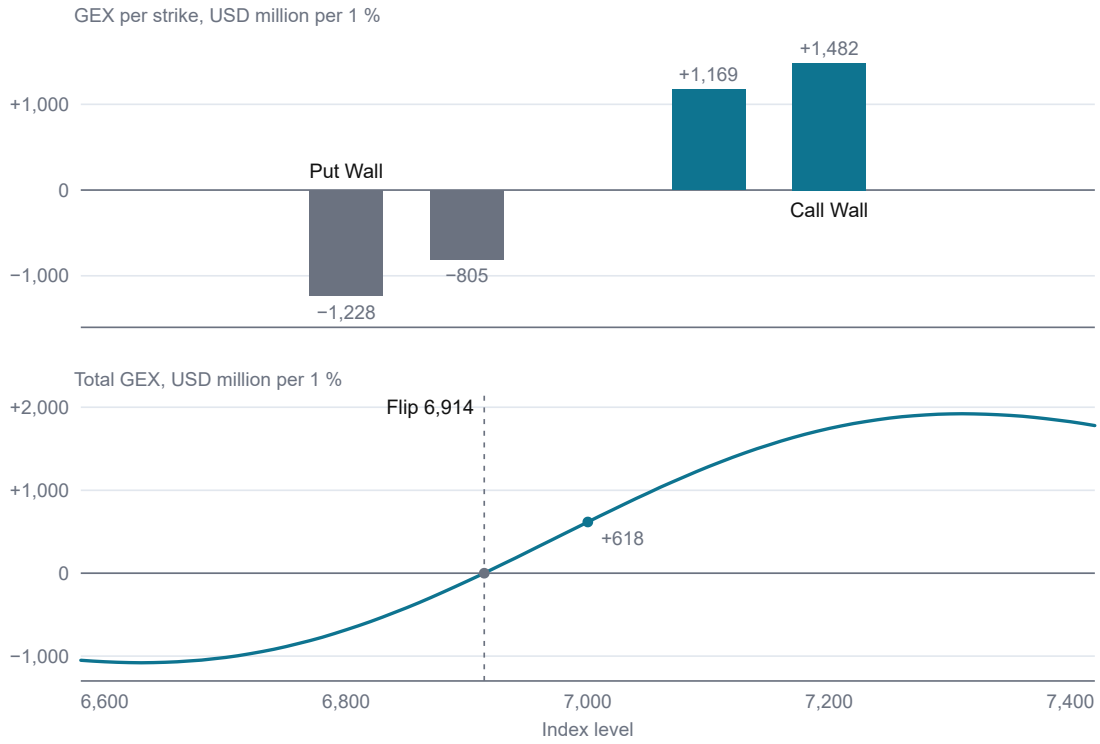


Fig. 14: On top the GEX per strike at an index level of 7,000 with put wall (6,800) and call wall (7,200). Below it the sum for every index level between 6,600 and 7,400. It crosses the zero line at the gamma flip at 6,914.

7. What walls are supposed to mean

Call wall and put wall are often read as resistance and support. The reasoning from the mechanics:

- **At the call wall** dealers have, by the standard assumption, bought many calls. As the price approaches, their long gamma grows. They sell into the rise. That slows it.
- **At the put wall** they have sold many puts. As the price approaches, their short gamma grows. They sell into the fall. That accelerates it, until the price has left the level behind.

By this logic the call wall is a brake and the put wall rather a place where things speed up, not a support. Both are conclusions from a model. They are not orders in the order book. Whether prices actually behave differently at these levels we have not been able to measure, because a reliable history of open interest per strike is missing (chapter 10).

8. 0DTE: the gamma that is missing from open interest

Chapter 7 showed: gamma at the money grows towards expiry, with one day to expiry to 5.5 times the value at 30 days. Options that expire today (0DTE) have the largest gamma of all in the afternoon.

Two problems follow for every GEX calculation:

- **Today's positions are in no open interest.** An option bought in the morning and expired in the evening was never part of the number published the next morning. Precisely the part with the largest gamma is invisible in the data.
- **The nearest expiry drowns out all the others.** If all terms are thrown into one sum, the shortest ones determine the result because their gamma is largest. These options no longer exist in the evening.

We observed the second problem on our own tool. On 10 June 2026 the call wall in the Nasdaq 100 jumped in the course of the day, with unchanged method and unchanged open interest, from 29,091 to 30,390, that is by around 1,300 points. Nobody had repositioned. The market had moved away from the strikes of the day's expiry, their gamma collapsed, and the weight tipped to more distant expiries (to the study). A map that mixes all terms measures different things in the morning and in the afternoon.

9. Expiry and pinning

Pinning is the idea that on expiry day the price "sticks" to a strike with a lot of open interest. The mechanics behind it: if dealers are long gamma at that strike, they sell above it and buy below it. Shortly before expiry gamma there is very large, and the trading correspondingly strong. It pushes the price towards the strike from both sides.

This again presupposes that dealers are long gamma there. If they are short gamma, the same mechanics work in the opposite direction and drive the price away from the strike.

We measured it like this: is the closing price on expiry days closer to round strikes than on other Fridays? In the S&P 500, Nasdaq 100, Dow and DAX the distance was the same (t 0.57, in the DAX −1.23). In index prices pinning was not visible to us (to the study).

10. Vanna and charm flows

Gamma describes what dealers have to trade after a **price move**. Chapter 7 showed two further reasons why their delta shifts without the price moving: a change in implied volatility (vanna) and the passing of time (charm).

Both act the same way on options out of the money: **falling volatility and passing time make their delta shrink.** What this means for the dealer, with the standard assumption:

Dealer's position	His hedge	The option's delta shrinks	He has to
Out-of-the-money puts sold	futures sold	he needs fewer futures sold	buy futures back
Out-of-the-money calls bought	futures sold	he needs fewer futures sold	buy futures back

In both cases dealers buy. Two widespread narratives arise from this:

- **Vanna rally:** in a quiet, rising market implied volatility falls. Dealers buy futures back, that supports the market, volatility falls further.
- **Charm into expiry:** in the days before a large expiry the delta of out-of-the-money options melts fastest. Dealers buy futures back, even without a price move.

The flip side: if volatility suddenly rises, the same thing runs backwards. The deltas of the out-of-the-money puts grow, dealers have to sell futures, and that in a falling market.

Both flows depend on the same assumption as GEX, and in addition on the skew from chapter 6. Our own measurement series on this is too short and was distorted by data errors. A substitute test, the open above the previous day's high in expiry week, showed no selling pressure: the first hour there ran better on average, not worse (+6.3 against +1.5 basis points, to the study).

11. The assumptions everything depends on

Every GEX number, every flip and every wall is the result of a calculation with assumptions. The most important ones:

- **The sign.** Dealers are long calls and short puts. If that is not true for a strike, its GEX has the wrong sign. For the flip it is enough if the assumption is wrong for a few large positions.
- **The data vintage.** Open interest is from the previous evening. Everything opened today is missing, including almost all 0DTE trading. Whoever recomputes the map during the day computes with old positions and new prices. The levels then trail behind the price.
- **Which expiries count.** The result depends heavily on which terms are included. We measured this on our own tool: a calculation from the option book of the S&P 500 ETF with 7 to 45 days to expiry showed short gamma on 60 % of days over 42 trading days. A freely published reference that uses all expiries of the S&P 500 showed short gamma on not a single day in the same period.
- **The model.** Gamma comes from the Black-Scholes formula with the implied volatility per strike. With skew, the delta dealers really use deviates from it (chapter 7).
- **Data quality.** Freely available option chains contain gaps and placeholders. An earlier reconstruction of our own over 1,302 trading days was not related to the reference at all (rank correlation −0.01). The formula was right, the input data were not.
- **Behaviour.** Dealers hedge net, not every option individually, often only beyond a deviation and frequently with other options instead of futures (chapter 8). The real flow in the index is smaller than the formula.

A simple rule for dealing with GEX values follows: two providers that compute with different expiries, different chains or a different sign rule can show an opposite regime for the same day, and neither of them has miscalculated.

Common misconception

> **Common misconception:** "GEX measures how dealers are positioned." GEX measures nothing. From yesterday's open interest, a formula and an assumption about buyers and sellers it computes an estimate. Only the clearing house knows the dealers' positions. A call wall is therefore not a wall of orders but a strike at which, by a model calculation, a particularly large amount would have to be hedged, if the assumption holds.

What you should understand now

- Volume counts trades, open interest counts open contracts. Open interest appears once a day and does not say who is buyer and who is seller.
- GEX per strike is $\Gamma_K \cdot (\mathrm{OI}_{\mathrm{call}} - \mathrm{OI}_{\mathrm{put}}) \cdot m \cdot S^2 \cdot 0.01$, in USD per 1 % move. The sign rests on the assumption that dealers have bought calls and sold puts.
- The gamma flip is the index level at which the sum changes sign. Call wall and put wall are the strikes with the largest positive and negative contribution.
- Options with very short terms have the largest gamma and are missing from open interest. Which expiries are included can reverse the sign of the result.

- Vanna and charm shift the hedge even without a price move. By the standard assumption falling volatility and passing time lead to buying by dealers, rising volatility to selling.

Limits of this account

- **Made-up chain.** Five strikes, one expiry, a uniform volatility. Real chains have hundreds of strikes, dozens of expiries and a skew.
- **One sign rule.** There are others, such as estimating the side from the price of each individual trade. That needs data not available here.
- **Own findings.** The comparison with the reference covers 42 trading days in a single market environment. The example of 10 June 2026 is a single day.
- **Not covered:** options on futures, options on individual shares and their effect on the index, and structured products whose hedging also arrives in the index.

Continue with chapter 10: what of all this survives a measurement.

Chapter 10

What can be measured: the mechanics against our own tests

In one sentence: Of the mechanics in this series, our measurements show what concerns the size of a move, that is, range, costs and fluctuation, and next to nothing that concerns its direction.

What this chapter is: a summary of our own measurements, not a new study. Every number comes from a linked study on this website. Data basis, sample and limits are in the studies themselves. **Not a trading recommendation.**

Chapters 1 to 9 explained what market participants have to do. This chapter asks what of that can be seen in prices.

1. Three stages that have to be kept apart

Between "dealers have to sell here" and "money can be made here" lie two steps that are often skipped.

- **Stage 1, the mechanics:** someone has to trade, and it can be derived in which direction. That was the content of chapters 2 to 9.
- **Stage 2, the effect:** this trading is measurable in the price, for example as a wider range or as a move in one direction.
- **Stage 3, the rule:** a trading rule fixed in advance earns from it, after costs, in data that was not used in the search.

A mechanism can be correct and still fail at stage 2, because the flow is too small or because others offset it. And a measurable effect can fail at stage 3, because it is smaller than trading costs or because it concerns only the size of a move and not its sign.

2. How we measure

Every rule is fixed before the data is evaluated. The search runs in one period, the check runs once in a second, untouched period (holdout). Results are after costs. The comparison is a random market: the same prices with random direction, that is, the same fluctuation and no information at all. And because among many tests some always look good by chance, a correction is applied across all tests together. The workflow is described here.

As of November 2026: 108 rules have been through an untouched holdout. Two passed.

3. The overview

Chapter	Statement from the mechanics	Effect visible in the price?	Does a rule earn from it?
2, 3	Market makers withdraw liquidity when it becomes dangerous for them	yes: spread at three to four times	no: direction afterwards like a coin toss
3	Index funds have to trade in the closing auction on rebalancing days	yes: auction window 1.2 to 1.7 times as wide	yes, in the FTSE 100: +11.5 basis points in the holdout
3	The published auction imbalance shows the direction	not usable	no
3	VWAP and TWAP are price levels	no: as in the random market	no
4	Stock indices rise mainly overnight	yes: +3.22 basis points per night	not with CFDs: financing eats it up
6	Implied volatility lies above realised	yes: 0.74 to 0.84	the band is not a price limit
8, 9	Little dealer gamma means wider days	yes, also in the holdout	no: no direction
9	Gamma flip and walls are price levels	not testable for us	not testable for us
9	On expiry day the price sticks to round strikes	no	not tested
9	Charm flows push or support before expiry	not in the substitute test	no

4. What is confirmed

Three findings survived the measurement. All three concern the size of a move or a flow with a known direction and a known time.

Dealer gamma and the width of the day. Chapter 8 predicts: if dealers are short gamma overall, their trading amplifies moves. On a freely published gamma series for the S&P 500 the range of the next US session in the lower gamma tercile was 0.25 typical daily ranges larger than in the upper one (t 6.5). In the one-off holdout from 2023 to June 2026 it was 0.20 (t 3.2). Moves there also reverted less. The finding holds at the same VIX level. There is no direction in it: two rules that followed the day's direction in the lower gamma tercile came in at −0.50 and −3.05 basis points per trade (to the study).

The volatility premium. Chapter 6 says: the option price contains a mark-up. For 2015 to 2022 the fluctuation of the preceding 20 days was in the median 0.76 (S&P 500), 0.84 (Nasdaq 100) and 0.74 (Dow) of the implied volatility. The priced-in daily range held the closing price on 80.1 % of days; with a normal distribution and no mark-up it would be 68.3 % (to the study).

Forced trading in the closing auction. Chapter 3 says: index funds trade changes to the index in the closing auction before the effective date. The range in the auction window was 1.2 to 1.7 times as wide on rebalancing days as on comparison days in all eight markets examined. In the FTSE 100 a rule held as well: sell in the last hour of trading, buy back after the auction price. In the holdout from 2023 to 2026 that brought +11.5 basis points per rebalancing day (t 3.9, 22 of 27 days positive). It is one of the two patterns that passed our joint correction, and it is small: 27 days, around eight opportunities a year (to the study).

5. What remains mechanics

For the core of chapter 9 we have no measurement that would carry the verdict.

Gamma flip, call wall, put wall. To check whether prices behave differently at these levels, one needs the open interest per strike for many years, day by day. We do not have that data. Only the current state is freely available.

Our own attempt to reconstruct the series failed, and that belongs here. Over 1,302 trading days our reconstructed gamma series was not related to the published reference (rank correlation −0.01). It showed short gamma on 99 % of days, the reference on 12 %. The formula was the same as in chapter 9. What was wrong were the input data and the selection of expiries. We discarded all evaluations we had previously made on that series (to the study).

The lesson from it is more general than our mistake: a GEX number is as good as its option chain and its assumptions. Before attributing an effect to a level, one has to know whether the calculation behind it is right at all.

0DTE, vanna and charm. The mechanics are coherent (chapters 7 and 9). They can only be measured with data on who opened which option when. Our own charm series is too short and was distorted by the same data errors. The only substitute test went against the widespread narrative: in expiry week the first hour after an open above the previous day's high ran better on average, not worse (+6.3 against +1.5 basis points, to the study).

Pinning. On expiry day the closing price in the S&P 500, Nasdaq 100, Dow and DAX was no closer to round strikes than on other Fridays (t 0.57, in the DAX −1.23). That does not refute pinning in individual shares, where it is described in the literature. In the index it was not visible to us (to the study).

6. What failed as a signal

- **The spread jump as a directional signal.** On gold ticks (April 2023 to December 2024) the spread jumped to more than three times its usual value 2.5 times a day on average outside scheduled releases. Whoever then followed the move or traded against it was at zero before costs (−0.07 basis points, t −0.26). The signal indicates costs, not direction (to the study). The order book of the ES future shows the same: after losses, market makers pull depth (t −4.4 over 40 days), and after 817 one-sided flow bursts the direction is a coin toss (to the study).
- **The auction imbalance.** Whoever followed the published imbalance of the closing auction in the S&P 500 ETF lost 3.86 basis points per trade after costs (t −3.6, to the study). The sum over 20 heavyweights of the Nasdaq 100 passed in the search period with +3.15 basis points (t 2.6) and came in at −0.15 in the holdout (to the study).
- **VWAP and TWAP as a level.** 17,858 touches of the TWAP in the Nasdaq 100, S&P 500 and DAX: on the first, second and third touch the line was broken in 52 %, 49 % and 50 % of cases, in the random market in 54 %, 49 % and 52 % (to the study). Rules with VWAP bands earned nothing after costs (to the study).
- **The expected move as a limit.** After the first touch of the band the price first ran back in 38.5 % of cases, in the random market in 42.9 %. All 15 rules against the band lost after costs (to the study).
- **The overnight return with CFDs.** From 2000 to 2018 the return of the US indices arose almost entirely overnight (+3.22 basis points per night, t 4.13). With a CFD nothing of it is left after spread and financing (to the study).
- **Further forced flows.** Pension funds at month-end, trend followers and leveraged ETFs: several visible, none tradable (to the study). Volatility-targeting funds: no effect (to the study).

7. The pattern behind it

Putting the confirmed and the failed findings side by side shows a regularity.

Forced trading changes the size of moves, not their sign. A dealer with short gamma sells into falling and buys into rising prices. He makes the day wider. In which direction is decided by the first impulse, and the mechanics do not know that. A market maker who withdraws his orders makes the next order more expensive, in both directions.

There is a direction only if the flow itself has one and comes at a known time. That is the case on index rebalancing days: shares leaving the index are sold, in the closing auction, on a date fixed weeks in advance. That is exactly where we have our only finding from this subject area.

Small effects lose against costs. Most visible effects are a few basis points. Spread and financing cost just as much. What would be a rule before costs is none after costs.

8. What can be done with it

No signal follows from what is confirmed. Three sober uses follow:

- **The expected width of the day.** Implied volatility and dealer gamma say something about how large the coming move is likely to be. Stop distance and position size can be aligned with that. Whether this earns more after costs we have not tested.
- **Avoiding costs.** Around data releases a market order costs a multiple. In our measurement the spread normalised within seconds to less than a minute. Waiting is cheaper there than trading.
- **Putting levels in their place.** A gamma flip, a call wall or a VWAP are results of calculations. They describe where something could happen if the assumptions hold. As the sole reason for a trade they did not carry in any of our tests.

Common misconception

Common misconception: "Whoever understands the mechanics has an edge in trading." Understanding the mechanics protects against wrong conclusions: against the call wall as resistance, against the future as a forecast, against short gamma as a signal to the downside. An edge only arises from it when a rule earns after costs in unseen data. Of 108 rules we have tested in this way, two managed that, and only one of them comes from the subject area of this series.

What you should understand now

- Between a correct mechanism and a rule that earns lie two stages: the measurable effect in the price and the test after costs in unseen data.
- Three things are confirmed for us: little dealer gamma goes with wider days, implied volatility lies above realised, and on index rebalancing days forced trading arrives in the closing auction.
- We could not test gamma flip and walls as price levels because the history of open interest is missing. Our own reconstructed series was unusable.
- What failed as a signal: the spread jump, the auction imbalance, VWAP and TWAP as a level, the expected move as a limit and the overnight return with CFDs.

- Forced trading changes the size of moves. It delivers a direction only if the flow itself has one and comes at a known time.

Limits of this account

- **Data basis.** Almost all tests use index CFDs on minute data, not futures. We have evaluated an order book once so far: the ES future on 40 days, order by order (to the study). The gamma series comes from a single provider whose calculation is not fully published.
- **What is missing.** No test with open interest per strike, none with data on who opened options, none with positions from the current day's trading. The central claims of chapter 9 thus remain untested, not refuted.
- **Periods.** The search periods mostly end in 2022, the holdouts run to 2026. A finding over four years is not a law of nature.
- **No disproof for other markets.** That we did not find something in the index says nothing about individual shares or other asset classes.

Back to the overview of the series.

An explanation of market mechanics with worked examples. Not investment advice and not a trading recommendation. Trading futures and options involves the risk of loss up to and beyond the capital employed.