


Gold Order Flow

From Zero to Volume Trading

The complete beginner-to-advanced guide to reading gold with volume, VWAP, volume profile, footprint and the order book.



Contents

PART I • FOUNDATIONS

Chapter 0 · Trading from Zero: Markets, Gold, CFDs and Futures

Chapter 1 · The Gold Futures Contract and the Trading Day

Chapter 2 · Why Volume? Price Is the Result

PART II • THE TOOLS OF VOLUME

Chapter 3 · Volume and Delta

Chapter 4 · VWAP

Chapter 5 · Volume Profile I: Building and Reading a Profile

Chapter 6 · Volume Profile II: Shapes, Yesterday's Map and Value Migration

PART III • THE AUCTION

Chapter 7 · Auction Market Theory

PART IV • INSIDE THE BAR

Chapter 8 · Footprint Charts: Inside the Bar

Chapter 9 · Footprint Patterns, Honestly

PART V • THE ORDER BOOK

Chapter 10 · Liquidity, the Order Book and the Tape

PART VI • PRACTICE

Chapter 11 · Putting It Together on Gold

Chapter 12 · Honest Research: Testing Ideas Without Fooling Yourself

REFERENCE

Appendix · Glossary, Pre-Session Checklist, GC→XAUUSD Conversion and Contract Cheat Sheet

Before you start

This book is education. It does not tell you what to buy or sell, it gives no signals, entries, stops or targets, and nothing in it is financial advice. Trading futures, CFDs and other leveraged products involves substantial risk of loss, and you can lose more than you deposit. Examples use real COMEX gold futures data recorded by AK Order Flow; XAUUSD equivalents are approximate. Contract specifications change: always check the exchange and your broker.

© AK Order Flow. Free to read and share in its complete, unmodified form.

PART I • FOUNDATIONS

0

Trading from Zero: Markets, Gold, CFDs and Futures

In this chapter: We start from nothing. You will learn what a market and a price really are, what it means to buy and sell, how traders can profit (or lose) from a falling price as easily as from a rising one, and what a broker does for you. We will separate the three ways most people trade gold on a screen (spot, CFDs and futures) and explain why they look identical on a chart but are completely different deals underneath. We will explain leverage and margin in plain words, look honestly at why most beginners lose money, and define risk management in a way you can actually use. Finally, we will explain why this book focuses on gold, and what "volume trading" and "order flow" will add to your understanding in later chapters. No prior knowledge is assumed. No promises are made.

Before we begin, one honest sentence that applies to this whole book: trading carries real risk, most beginners lose money, and nothing in these pages tells you what to buy or sell. This book teaches how the gold market works and how to read the information it produces. What you do with that understanding is your own decision.

What Is a Market, and What Is a Price?

A market is a meeting place for disagreement

A market is any place, physical or electronic, where people who want to buy something meet people who want to sell it. A farmers' market, a car auction, an online marketplace for second-hand phones and the global gold market all work on the same basic idea: buyers and sellers come together, and when they agree, a trade happens.

Here is the part that surprises many beginners. Every single trade needs two people who see the same thing differently. When someone buys an ounce of gold, someone else sells that same ounce at that same moment, at that same price. The buyer thinks gold is worth holding at that price; the seller thinks it is time to let it go. Both might have perfectly good reasons. One might be a jewellery company stocking up, the other a fund taking profits. One might be a short-term trader expecting a bounce, the other a short-term trader expecting a drop. A market is a machine that turns all of this disagreement into a single, visible number.

A price is the last agreement

That number is the **price**. In the simplest sense, the price of gold is **the price at which the most recent trade happened**. It is not a fact about what gold is "really worth". It is a record of the last point where a buyer and a seller agreed.

This is worth repeating, because it will matter enormously later in the book: price is the result of trades. It does not move by itself. It moves because buyers and sellers keep making new agreements at new levels. If buyers become more eager than sellers, they have to offer more to get their trades done, and the price rises. If sellers become more eager than buyers, they have to accept less, and the price falls.

Key idea: A price is not a judgment of value. It is the last place a buyer and a seller agreed. Every move on a chart is made of thousands of these small agreements.

Waiting prices: the bid and the ask

At any moment, there are also people who are not trading yet but are waiting to trade at a particular price. Some buyers are saying, "I will buy, but only at \$3,999.90." Some sellers are saying, "I will sell, but only at \$4,000.10." These waiting offers give us two more prices you will see on every trading screen:

- The **bid** is the best (highest) price that a waiting buyer is currently offering.
- The **ask** (also called the **offer**) is the best (lowest) price that a waiting seller is currently willing to accept.

The ask is always a little above the bid. The gap between them is called the **spread**. If the best bid is \$3,999.90 and the best ask is \$4,000.00, the spread is \$0.10.

If you want to buy *right now*, you cannot wait for a seller to come down to you; you must accept the best seller's price, the ask. If you want to sell *right now*, you must accept the best buyer's price, the bid. So a trader who buys and immediately sells, with no change in the market at all, loses the spread. That is why the spread is a **cost**, the price of being impatient. We will come back to this many times.

Gold note: Gold is quoted in US dollars per troy ounce. A **troy ounce** is the traditional unit for precious metals and weighs about 31.1 grams, a little more than the everyday (avoirdupois) ounce of about 28.3 grams. When this book says "ounce", it always means troy ounce.

Volume: how much actually traded

The last of our basic words is **volume**: how much actually changed hands in a given period. If, in one minute, buyers and sellers agreed on trades totalling 500 contracts, the volume for that minute is 500. Volume tells you how much activity sat behind a price move. A price that rises on very little trading and a price that rises on enormous trading look the same on a simple line chart, but they are different events. That difference is the subject of this entire book.

What Does Trading Gold Actually Mean?

You can profit, or lose, without owning any gold

When most people hear "trading gold", they picture someone buying a gold bar and selling it later. That is one way to do it, but it is not what most traders on a screen are doing. Most of them never touch, store or even legally own any gold at all.

Trading means taking a **position** on where a price will go. If you are right, you keep the difference between where you got in and where you got out. If you are wrong, you pay that difference. The

gold itself is just the thing whose price everyone is watching.

Two directions: long and short

There are two ways to take a position.

Long means you buy first. You make money if the price rises, and lose money if it falls. This is the intuitive direction: buy low, sell higher.

Short means you sell first. You make money if the price falls, and lose money if it rises. Later, you buy back to finish the trade. This feels strange at first: how can you sell something you do not have? In the markets this book discusses (futures and CFDs), a short position is simply a contract that gains when the price falls. You never have to own the gold first. The exchange or broker records that you have "sold" a certain amount, and you finish the trade by buying the same amount back.

Two more words complete the picture:

- A **position** is the trade you have open right now, for example "long 10 ounces of gold from \$4,000.00".
- To **close** a position means to end the trade by doing the opposite: a long is closed by selling, a short is closed by buying back. Only when you close is your profit or loss final. Before that, it is "open" or "unrealised" and keeps changing with every tick of the price.

Worked example: a long and a short on 10 ounces

Imagine a position that represents 10 troy ounces of gold. (Later you will meet a real contract of exactly this size, the Micro Gold future.)

Long example. You buy 10 ounces at \$4,000.00. Later you sell at \$4,012.00. Profit per ounce = $\$4,012.00 - \$4,000.00 = \$12.00$. Total = $\$12.00 \times 10 \text{ ounces} = \textbf{\$120 profit}$ (before costs).

If instead gold had fallen and you sold at \$3,991.00, the result would be $(\$3,991.00 - \$4,000.00) \times 10 = -\$9.00 \times 10 = \textbf{\$90 loss}$.

Short example. You sell 10 ounces at \$4,000.00, expecting a fall. Gold drops and you buy back at \$3,990.00. Profit per ounce = $\$4,000.00 - \$3,990.00 = \$10.00$. Total = $\$10.00 \times 10 = \textbf{\$100 profit}$ (before costs).

If gold rose instead and you bought back at \$4,008.00, the result would be $(\$4,000.00 - \$4,008.00) \times 10 = \textbf{\$80 loss}$.

Notice the symmetry. Long and short are mirror images. Neither is "safer" by nature; both lose money when the market goes the other way.

Common mistake: Thinking of short selling as "betting against the market" in some moral or unusual sense. In futures, going short is exactly as routine as going long. Every futures trade has a buyer and a seller, so for every long position that exists, there is a short position somewhere else.

Four ways to get exposure to gold's price

"Exposure" means that your money goes up or down when the price of gold goes up or down. There are four common ways to get it.

1. **Physical gold (bars and coins).** You own real metal. This is the oldest and most direct form of ownership. The downsides are practical: you must store it safely, possibly insure it, and dealers buy back at a lower price than they sell (a wide spread). Physical gold suits long-term holders far better than short-term traders.
1. **Gold ETFs.** An **exchange-traded fund** is a fund listed on a stock exchange. A gold ETF holds gold (usually in vaults) and issues shares that track its value. You own shares in the fund, not the metal directly. ETFs are popular for investors who want gold in a normal brokerage account.
1. **CFDs (for example, the symbol XAUUSD).** A **contract for difference** is a private agreement with your broker to settle the *difference* in price between when you open and when you close. You own nothing except that agreement. XAUUSD is the standard symbol for "gold (XAU) priced in US dollars (USD)", and it is the symbol most retail traders on platforms such as MetaTrader or TradingView know.
1. **Futures (GC and MGC).** A **futures contract** is a standardised contract traded on an exchange. For gold, the main exchange is **COMEX**, part of **CME Group** in the United States. **GC** is the standard gold futures contract (100 troy ounces) and **MGC** is the Micro Gold contract (10 troy ounces). CME also lists a smaller one-ounce gold contract; check CME's website for its current details.

Short-term traders mostly use the last two: CFDs and futures. On a chart they look almost identical, because they track the same underlying gold price. Behind the screen, as we will see shortly, they are completely different deals.

The first six words, collected

Here are the six words every gold trader uses daily, gathered in one place:

Word	Meaning	Gold example
Price	The last traded price	Last trade at \$4,000.00
Bid	The best price waiting buyers offer	\$3,999.90
Ask	The best price waiting sellers accept	\$4,000.00
Spread	Ask – bid; a cost you pay to trade immediately	\$0.10
Volume	How much actually traded	500 contracts in a minute
Leverage	Controlling a large position with a small deposit	\$2,000 controlling \$40,000 of gold

One rule to burn into memory: **you buy at the ask and sell at the bid**. Whenever you trade immediately, you cross the spread.

Brokers: Your Door Into the Market

What a broker does

Ordinary people cannot walk into an exchange and trade directly. They need a **broker**: a regulated company that opens an account for you, holds your money, gives you a trading platform and carries your orders to the market (or, in some cases, takes the other side itself).

What exactly a broker does depends heavily on which product you trade:

- A **futures broker** (in the United States, the firm that holds your money is called a **futures commission merchant**, or FCM) passes your order to the exchange. Your trade is matched with another trader's order in the exchange's central order book. The broker earns a commission; it does not care whether you win or lose.
- A **CFD broker** creates the contract with you directly. The broker sets the prices you see (based on prices from its own sources), and it is the other side of your trade, either keeping that risk itself or passing it on to a bank or another partner.

Questions worth asking about any broker

You do not need to become an expert in financial regulation, but you should be able to answer these questions about any firm before sending it money:

1. **Who regulates it?** Look the firm up on the regulator's own public register, not just the logo on the broker's website. Examples include the CFTC and NFA in the United States, the FCA in the UK, regulators applying EU (ESMA) rules in Europe, and ASIC in Australia.
2. **What am I actually trading?** A futures contract on an exchange, or a CFD with the broker?
3. **What does it cost?** Commissions, spreads, overnight financing (swap), data and platform fees.

4. **What happens if I lose more than my deposit?** Some products and some regions offer negative-balance protection; others do not.
5. **How is my money held?** Regulated firms usually must keep client money separate from their own.

Common mistake: Choosing a broker because of a large deposit bonus, a promise of very high leverage or a friendly "account manager" who offers to trade for you. Unregulated firms, guaranteed returns and pressure to deposit more are warning signs, not features.

Futures vs CFD: What Are You Really Buying?

The same gold chart can sit on top of two completely different deals. This section explains the difference, because it affects your costs, your risks and, most importantly for this book, the data you can see.

First, a word on "spot"

You will often hear the phrase **spot gold**. "Spot" means the price for (almost) immediate delivery, as opposed to delivery at some later date. The global spot gold market is centred in London and is traded **over the counter** (OTC): banks and dealers trade directly with one another rather than through one central exchange. Because there is no single central book, there is no single, complete, official record of every spot trade as it happens.

Retail traders almost never trade the true interbank spot market. When a retail platform shows "XAUUSD", it is usually offering a **CFD** whose price is based on spot gold, as quoted by that broker.

CFD: a deal with your broker

With a **CFD (contract for difference)**, you and your broker agree to settle the difference between the price when you open and the price when you close. If you buy a CFD for 1 ounce at \$3,971.00 and close it at \$3,981.00, the broker pays you \$10. If you close at \$3,961.00, you pay the broker \$10. No gold moves anywhere; only money does.

The broker is the counterparty. The prices you see come from the broker's own price feed. Two different brokers can show slightly different prices, different spreads and different candle highs and lows at the same moment.

Futures: a contract on an exchange

With a **futures contract**, your broker only routes your order to the exchange (COMEX, for gold). There, your order meets other traders' orders in **one central order book** that everyone in the world shares. When your buy order matches someone's sell order, a trade happens and is published to every market participant at once. A **clearing house** then steps between the two sides and guarantees the trade, so you do not have to worry about whether the unknown trader on the other side can pay. (Chapter 1 explains all of this in detail.)




Figure 1. Futures trade in one central order book, so the exchange publishes every real trade. CFD/spot prices come from many separate broker books.

How they are built

	Futures (GC, MGC)	CFD (XAUUSD)
The other side of your trade	Another trader, via the clearing house	Your broker (or its partner)
Where the price comes from	One central order book on the exchange	The broker's own price feed
Volume you can see	Real volume of the whole market	Tick count at that broker only
Same data for everyone?	Yes: one market	No: differs from broker to broker

That third row, "volume you can see", is the reason this book exists in its current form. We will return to it at the end of this section.

What it costs

	Futures	CFD
Main cost	Commission + exchange and regulatory fees + spread	Spread (sometimes plus commission)
Holding overnight	No daily swap charge, but contracts expire and must be "rolled"	A daily swap or financing charge, usually
Position size	Fixed sizes: GC = 100 oz, MGC = 10 oz (plus CME's one-ounce contract)	Flexible lots, e.g. 0.01 lot $\approx$ 1 oz (varies by broker)
Expiry	Yes, in set contract months	Usually none

A **swap** (or **overnight financing**) is a small charge or credit a CFD broker applies for each night you hold a position, reflecting the cost of financing it. Futures do not charge a swap, but that cost does not disappear: it is built into the futures price itself, which is one reason the futures price usually sits above the spot price. That leads us to something you will see on your own screens.

Why GC and XAUUSD show different numbers

If you put a GC futures chart next to an XAUUSD CFD chart, you will notice that the futures price is usually higher, by something like \$27–30 at the time of writing. On a given afternoon, GC might be trading at \$4,000.0 while XAUUSD shows roughly \$3,971.



Figure 2. The same gold, two prices: COMEX futures trade above spot/CFD by a gap (the basis) that changes slowly. Convert every futures level before using it on an XAUUSD chart.

This gap is not an error and not a trick. A futures contract is a promise to deliver gold at a future date. Someone who holds physical gold until then gives up interest they could have earned on their

money and pays for storage and insurance. The futures price includes that **cost of carry**, so it usually sits above spot. The gap depends on interest rates and on how far away the delivery month is. It changes slowly over time and jumps when traders move from one contract month to the next. It also varies slightly between CFD brokers, because each one quotes its own price.

For a CFD trader reading this book, the practical lesson is simple: **a level on the GC chart corresponds to a level roughly \$27–30 lower on an XAUUSD chart**, give or take, and you should measure the current gap yourself rather than trust any fixed number. Throughout this book, when we mention a GC price, we will try to give an approximate XAUUSD equivalent beside it.

Try it: If you have access to both charts, note the GC price and your broker's XAUUSD price at the same minute on three different days. Subtract. You will see that the gap is fairly stable from day to day but not constant, and that it is your own broker's gap that matters for your own charts.

If you lose, who wins?

This is one of the most important and least discussed differences.

With **futures**, your loss is another trader's gain. The exchange and your broker earn their fees whether you win or lose. They have no reason to care which side is right.

With **CFDs**, it depends on the broker's business model. Some brokers pass client orders on to banks or liquidity providers (often called an "A-book" model). Others keep the opposite side of their clients' trades themselves (a "B-book" model), in which case your loss can be the broker's gain. This is legal, common and regulated in many countries, but you should know which model your broker uses, because it shapes the relationship between you.

Regulation differs by country

The rules depend on where you live. A few examples, as general education and not legal advice:

- **United States:** retail traders can trade futures, but CFDs are not offered to US retail clients.
- **European Union and United Kingdom:** retail CFD trading is allowed but restricted. Leverage on gold CFDs for retail clients is capped at 20:1, retail accounts have negative-balance protection (you cannot lose more than your deposit), and brokers must publish the percentage of their retail clients who lose money. Those published figures commonly show that a majority of retail CFD accounts lose.
- **Futures, everywhere:** losses can exceed the money in your account. If the market moves far enough against you, you can owe your broker money.

Rules change, and they differ from country to country. Check your own regulator's current rules before you trade anything.

Why this book reads futures data

Here is the heart of the matter. The style of analysis this book teaches, which we call **volume trading** or **order flow analysis**, depends on seeing:

- **every trade** in the market,
- its **real size**,
- **which side was aggressive** (the buyer who paid the ask, or the seller who hit the bid), and
- **the same data that everyone else sees**.

A single CFD broker can only show what happened among its own clients and in its own price feed. What a CFD platform labels "volume" is usually **tick volume**: a count of how many times the broker's price changed, not how many ounces or contracts traded. Tick volume can loosely resemble real activity, but it is not the same thing, and it differs from broker to broker.

Only an exchange with one central order book can show the whole market. For gold, that exchange is COMEX, and the data is the data of GC and MGC futures. So, in this book:

- We **read futures data**, because only futures show real, whole-market volume.
- Footprint charts, delta, volume profile and other order-flow tools (later chapters) work properly on that data, and mislead on tick volume.
- CFDs are not "bad". Many traders use them for perfectly sensible reasons, such as flexible position sizes. But if you want to *read* volume, you need futures data, even if you end up trading something else.

Key idea: CFD = a deal with your broker. Futures = a contract on an exchange. Only futures show the real volume of the whole gold market, which is why every volume chart in this book is built from GC futures data.

Leverage and Margin in Plain Words

Controlling a big position with a small deposit

Suppose gold is at \$4,000 an ounce. A position of 10 ounces is worth \$40,000. A position of 100 ounces is worth \$400,000. Very few people trade with that much cash, yet traders routinely hold positions of that size. How?

The answer is **leverage**: the ability to control a large position with a much smaller deposit. The deposit is called **margin**. Margin is not a price you pay for the gold, and it is not a fee. It is closer to a **security deposit**, like the deposit you leave when renting a flat: money set aside to show you can cover losses. If the trade goes well, the margin is released back to you when you close. If it goes badly, losses are taken from your account.

The full value of the position is called its **notional value**:

Notional value = price × quantity

So 10 ounces at \$4,000 has a notional value of \$40,000, whatever deposit you put down.

Worked example: what leverage does to a 1% move

Assume, for illustration only, that you need a deposit of \$2,000 to hold a position of 10 ounces of gold worth \$40,000. Your leverage is $\$40,000 \div \$2,000 = \mathbf{20 \text{ to } 1}$ (written 20:1).

Now gold moves 1%, from \$4,000 to \$4,040, or from \$4,000 to \$3,960.

- The position gains or loses 1% of \$40,000 = **\$400**.
- Relative to your \$2,000 deposit, that is **20%**.

A 1% move in gold is not rare; it can happen within a single day. With 20:1 leverage, that ordinary move changes your deposit by a fifth. A 5% move against you would wipe out the entire \$2,000. That is what leverage really is: it does not change how much gold moves; it changes **how much each move matters to you**.

Gold move	P&L on \$40,000 position	As % of \$2,000 deposit
0.25% (\$10)	\$100	5%
0.5% (\$20)	\$200	10%
1% (\$40)	\$400	20%
2.5% (\$100)	\$1,000	50%
5% (\$200)	\$2,000	100%

(If you are a CFD trader: these dollar moves are the same on XAUUSD, which tracks the same gold price a little lower, around \$3,971 in this example.)

Margin calls and forced closing

Because your losses come out of your account, brokers keep watching. If your account falls below a certain level, called the **maintenance margin** in futures (CFD brokers use similar terms such as "margin level"), one of two things happens:

- You get a **margin call**: a demand to deposit more money, or
- The broker **closes your position for you**, often automatically and at the worst possible moment, locking in the loss.

In futures, if the market moves very fast, your losses can be larger than everything in your account, and you will owe the difference. Chapter 1 shows exactly how futures margin works, including the daily settlement of gains and losses.

Common mistake: "Low margin means low risk." It is the opposite. A lower margin requirement means higher leverage, which means each dollar of gold movement has a bigger effect on your account. The deposit is small; the position is not.

Key idea: Leverage is a magnifier. It magnifies gains and losses equally, and it does not care whether you are right.

Why Most Beginners Lose Money

It would be easy to skip this section. Please do not. Understanding why most people fail is the most useful protection you can have.

Reason 1: costs are certain, profits are not

Every trade has costs: the spread, commissions and fees, slippage (getting filled at a worse price than expected) and, for CFDs, overnight financing. You pay these on every single trade, winning or losing. Profits, on the other hand, are uncertain.

Here is a simple thought experiment. Imagine a game where you flip a fair coin. Heads, you win \$100; tails, you lose \$100. Over many flips, you would expect to break even. Now add a \$10 cost to every flip. Over 100 flips, you would expect to lose about **\$1,000**, even though the coin is perfectly fair. The cost turned a neutral game into a losing one.

Short-term trading without a genuine, tested advantage behaves a lot like that coin. The more often you trade, the more the costs accumulate. Chapter 1 shows how to calculate the full cost of a gold trade in ticks and dollars.

Reason 2: leverage turns normal noise into disaster

Prices wobble constantly. On gold, moves of several dollars in a few minutes are completely ordinary and mean nothing. With modest position sizes, these wobbles are an irritation. With large leverage, the same ordinary wobble can hit a margin limit, trigger a forced close or wipe out a week's work. Many beginners do not lose because their ideas are terrible, but because their positions are too big to survive normal movement while the idea plays out.

Reason 3: no plan

A trade without a plan is a decision made in the moment, under pressure, with money at stake. A plan answers in advance: Why am I entering? Where would I accept that I am wrong? How much am I risking? What would make me exit? Without those answers written down beforehand, every price tick becomes a new emotional decision.

Reason 4: human nature works against us

Research in behavioural finance has documented a few patterns that show up again and again:

- **Cutting winners early, letting losers run.** People feel the pain of a loss more strongly than the pleasure of an equal gain, so they grab small profits quickly and hold losing positions hoping they will "come back".
- **Revenge trading.** After a loss, the urge to "win it back" leads to bigger, faster, worse trades.

- **Overconfidence after a few wins.** A short lucky streak feels like skill and leads to larger positions just before the luck runs out.

Reason 5: overtrading and tiny samples

Beginners often trade too often and judge a method on far too few trades. Ten trades tell you almost nothing about whether a method works; random luck easily dominates such a small sample. Yet many people abandon one method after three losses and fall in love with another after three wins.

Reason 6: not understanding the instrument

Confusing tick size and tick value, misunderstanding how much a one-dollar move is worth, not knowing when the contract expires or when the market is thin: these technical misunderstandings cause losses that have nothing to do with market direction. That is why Chapter 1 spends so much time on the gold futures contract itself.

Reason 7: the search for a holy grail

Courses, channels and indicators that promise a simple, sure way to win are everywhere. There is no such thing. Markets are competitive: if a simple pattern reliably produced profits after costs, other traders would trade it until it stopped working. That applies to the methods in this book as well; we will be honest about it throughout.

Key idea: Most beginners do not lose because markets are rigged. They lose because of costs, oversized positions, no plan and ordinary human psychology. All four are things you can understand and control.

What Risk Management Really Means

Risk management sounds like a corporate phrase, but the idea is simple: **decide in advance how much you are willing to lose, and arrange things so that no single mistake, or even a long run of mistakes, can end your trading.**

It has several parts.

Part 1: know where you are wrong before you enter

Every trading idea has a point where the idea is simply invalid. If you buy because you think a level will hold and the price falls well through it, your reason for being in the trade no longer exists. The price at which your reason disappears is your **invalidation point**. Many traders place a **stop-loss order** there: an instruction to the broker to close the position automatically if the price reaches that level. (Chapter 1 explains why a stop-loss order limits your loss but does not guarantee the exact exit price.)

This book does not tell you where to place any such level. What matters here is the principle: **the decision about where you are wrong comes before the trade, not during it.**

Part 2: size the position from the risk, not from hope

Once you know how far away your invalidation point is, you can calculate how big your position should be. The question is not "How many contracts can I afford?" but "How many contracts can I hold so that, if I am wrong, I lose only the amount I decided in advance?"

The formula is:

Position size = money you are willing to risk ÷ (distance to invalidation × dollar value per unit of movement)

Worked example. You have a \$10,000 account and decide in advance to risk no more than 1% of it on any single idea: \$100. (Keeping the risk per idea to a small fraction of the account, often quoted in trading education as 1% or less, is a widely used guideline; it is a choice, not a law.) Suppose your invalidation point is \$5 away from where you would enter.

- A position of 10 ounces (one Micro Gold contract) gains or loses \$10 for every \$1 gold moves. A \$5 move costs **\$50**. Two such contracts would risk \$100, which fits the plan.
- A position of 100 ounces (one standard GC contract) gains or loses \$100 for every \$1 gold moves. A \$5 move costs **\$500**, five times the planned risk. That position is simply too large for this account and this idea.

The same idea, the same chart, the same invalidation point, yet one size is reasonable and the other is reckless. Position size is where risk management actually happens.

Part 3: limits beyond the single trade

Good risk management also sets limits above the level of individual trades:

- A **daily loss limit**: an amount after which you stop trading for the day, no matter what.
- A **maximum drawdown** for the account: a point at which you stop, step back and review everything before trading again. (A **drawdown** is the fall from your account's highest value to a later low.)
- **Never trading with money you need** for rent, bills or anything else in your life.

Why small losses matter so much: the recovery table

Losses and gains are not symmetrical. If your account falls by a certain percentage, you need a *larger* percentage gain to get back to where you started.

Loss	Gain needed to recover
10%	11.1%
20%	25%
30%	42.9%
50%	100%
75%	300%
90%	900%

A 50% loss needs a 100% gain just to break even. This is why protecting capital comes first. Small losses are recoverable; large ones often are not.

A run of losses happens to everyone

Even a sound approach will produce runs of losing trades. Look at what ten losses in a row do to an account at different levels of risk per trade:

- Risking 1% of the current balance each time, ten straight losses leave you about **9.6%** down. Uncomfortable, but survivable.
- Risking 10% each time, the same ten losses leave you about **65%** down, needing almost a 190% gain to recover.

Same trades, same market, same bad luck. The difference is entirely position size.

Practise without real money first

Most platforms offer a **simulated** (demo or paper) account that uses real market prices but fake money. Use it. Learn the platform, the order types, the contract sizes and your own reactions before risking a cent. A simulator does not perfectly reproduce real fills or real emotions, but mistakes there are free.

Common mistake: Treating risk management as something you add after you find a good strategy. It is the other way round. Without risk management, even a good approach can destroy an account during an ordinary bad run.

Why This Book Focuses on Gold

There are thousands of markets. This book studies one of them in depth, for several reasons.

Gold is deep and global. Gold is traded around the world, nearly around the clock: by central banks, jewellers, miners, investment funds, banks and individual traders. The COMEX gold futures market is one of the most actively traded commodity futures markets in the world. Deep markets have many participants and plenty of activity to study.

Gold trades almost 23 hours a day. Gold futures trade from Sunday evening to Friday afternoon, New York time, with a short pause each day. Activity flows from Asia to London to New York. That gives us a rich, varied trading day to study, with quiet periods and busy periods that behave differently (Chapter 1 maps the whole day).

Gold reacts to big, understandable forces. Interest rates, the US dollar, inflation data, central bank buying, geopolitical tension and risk appetite all move gold. Major US economic releases, such as the monthly jobs report, often produce sharp, high-volume moves that make the mechanics of the market unusually visible.

Gold has a clean, public, centralised futures market. Because GC and MGC trade on one exchange with one order book, we can see real volume, real trade sizes and which side was aggressive. That is exactly the raw material volume trading needs.

Many retail traders already watch gold. XAUUSD is one of the most popular CFD symbols in the world. If you trade it, everything in this book applies to you: you can read the futures data and translate the important levels to your own chart using the GC-XAUUSD gap.

Honesty also requires the downsides. Gold can be **very volatile**. After the large rallies of recent years, daily moves have often been wide, and in late 2025 CME raised margin requirements on precious metals more than once as volatility rose. A market that moves a lot offers more to study, and also more ways to get hurt. Everything in the previous section about position size applies with full force to gold.

Gold note: Focusing on one market lets you learn its personality: when it is busy, when it is quiet, how it behaves around news, how its contract months roll. That depth is worth more than shallow familiarity with twenty markets.

What Volume Trading and Order Flow Will Add Later

You now know what a price is: the last agreement between a buyer and a seller. Most charts show you only the result of those agreements, compressed into lines or candles. The rest of this book is about the information that ordinary price charts leave out.

Candles show what happened; volume shows who, where and how hard

A standard candle on a chart shows four numbers for a period of time: where the price opened, its highest and lowest points, and where it closed. That tells you **what** happened to the price. It does not tell you:

- **who** drove the move: buyers aggressively paying the ask, or sellers stepping away;
- **where** the trading actually happened: at the top of the candle, the bottom or the middle;
- **how hard**: whether thousands of contracts traded or only a handful.

Volume-based tools answer those questions. Over the following chapters you will meet them one by one:

- **Volume and delta** (Chapter 3): how much traded, and how much of it was aggressive buying versus aggressive selling.
- **VWAP** (Chapter 4): the average price weighted by volume, a reference many participants watch.
- **Volume profile** (Chapters 5 and 6): how much traded at each price, revealing where the market spent its effort.
- **Auction market theory** (Chapter 7): a framework for thinking about the market as a continuous two-sided auction searching for fair value.
- **Footprint charts** (Chapters 8 and 9): opening up each candle to see buying and selling at every price inside it.
- **The order book and the tape** (Chapter 10): the orders waiting to trade and the trades that actually happened, in real time.
- **Putting it together on gold** (Chapter 11) and **honest research** (Chapter 12): how to combine these views and how to test ideas without fooling yourself.

More information, not a magic edge

It is important to set expectations now, before the excitement of new tools takes over. Order flow is **more information**, and better information about what has already happened. It is **not** a crystal ball. In our own testing, order-flow signals on their own did not predict the direction of gold after costs. What volume information does offer is **context**: a clearer picture of where the market has done business, where it has shown effort without result, and where activity has been thin. Many experienced practitioners find it most useful for understanding the situation, timing decisions and managing risk, rather than as a stand-alone signal.

If anyone tells you that a volume tool shows "where the banks are buying" or gives you a sure edge, be sceptical. We cannot see who is behind a trade; we can only see its size, its price, its time and which side was aggressive. Everything else is inference, and inference can be wrong.

Key idea: Price charts show you the result. Volume and order flow show more of the process behind the result. Understanding the process does not guarantee profits, but it does stop you from mistaking a guess for a fact.

Chapter summary

- A **market** is where buyers and sellers meet; every trade needs both, at the same price, at the same moment.
- A **price** is the last agreement between a buyer and a seller, not a statement of "true value".
- The **bid** is the best waiting buyer's price, the **ask** the best waiting seller's price, and the **spread** (ask – bid) is a cost: you buy at the ask and sell at the bid.
- **Volume** measures how much actually traded.

- **Long** = buy first, profit if the price rises. **Short** = sell first, profit if it falls. Profit or loss is final only when you **close**.
- You can gain exposure to gold through physical metal, ETFs, **CFDs** (XAUUSD) or **futures** (GC, MGC).
- A **CFD** is a private deal with your broker on the price difference; a **futures** contract trades on an exchange through one central order book, with a clearing house guaranteeing both sides.
- GC futures usually trade around **\$27–30 above** XAUUSD at the time of writing because of the cost of carry; the gap drifts and differs slightly by broker, so measure it yourself.
- Only futures show **real, whole-market volume**; CFD platforms usually show **tick volume**. That is why this book reads GC futures data.
- **Leverage** lets a small deposit (**margin**) control a large **notional value**. It magnifies gains and losses equally.
- Most beginners lose because of **costs, oversized positions, no plan, psychology, overtrading** and **misunderstanding the instrument**.
- **Risk management** means deciding in advance where you are wrong and how much you can lose, then sizing the position to fit, and setting daily and overall limits.
- Gold is deep, global, active nearly 23 hours a day, driven by understandable forces and traded on a centralised futures market, but it can be very volatile.
- Volume trading adds **who, where and how hard** to the **what** shown by candles. It is more information, not a guaranteed edge.

Checklist

- I can explain in one sentence what a price is.
- I can define bid, ask, spread and volume, and I know that I buy at the ask and sell at the bid.
- I can calculate the profit or loss of a long or short position on a given number of ounces.
- I can explain the difference between a CFD and a futures contract, including who is on the other side of my trade.
- I know why GC and XAUUSD show different prices and that I must measure the gap myself.
- I know why tick volume on a CFD platform is not the same as real futures volume.
- I can calculate notional value and explain what a 1% gold move does to a leveraged deposit.
- I know whether my own country allows CFDs, what leverage limits apply and whether losses can exceed my deposit.
- I can name at least four reasons beginners lose money.
- I can size a position from a fixed amount of risk and a distance to my invalidation point.
- I understand that order flow is more information, not a crystal ball.

Quiz

1. You buy 10 ounces of gold at \$4,000.00 and sell at \$3,994.50. What is your profit or loss before costs?
2. The best bid is \$3,999.90 and the best ask is \$4,000.10. You buy at market and immediately sell at market, with no price change. How much do you lose per ounce, and why?
3. Which of the following shows the real traded volume of the whole gold market: (a) tick volume on a CFD platform, (b) COMEX gold futures volume, (c) the number of price changes in your broker's XAUUSD feed?
4. With 20:1 leverage, roughly what percentage of your deposit is gained or lost when gold moves 2%?
5. An account has fallen 50% from its starting value. What percentage gain is needed to get back to the start?

Quiz answers

1. A **\$55 loss**. $(\$3,994.50 - \$4,000.00) \times 10 = -\$5.50 \times 10 = -\$55$.
2. You lose **\$0.20 per ounce**, the full spread. You bought at the ask (\$4,000.10) and sold at the bid (\$3,999.90). Trading immediately means crossing the spread.
3. **(b)** COMEX gold futures volume. Both (a) and (c) describe tick volume: a count of price changes at one broker, not the volume of the whole market.
4. About **40%**. With 20:1 leverage, every 1% move in gold changes the deposit by about 20%, so $2\% \times 20 = 40\%$.
5. **100%**. Half of the original account must double to return to the starting value. That is why limiting losses comes before everything else.

1

The Gold Futures Contract and the Trading Day

In this chapter: You will learn exactly what a futures contract is and who stands on the other side of your trade; how margin, leverage and daily mark-to-market move money in and out of a futures account; the two gold contracts this book uses, GC and MGC, and how to turn any price move into dollars; why there are several gold contracts at once and what "rolling" does to volume and to your charts; when gold trades across its nearly 23-hour day and when it officially settles; what a candle shows and, more importantly, what it hides; what liquidity really means; the full cost of a trade in ticks and dollars; and, finally, what "order flow" and "volume trading" actually mean. By the end, you will have every piece of groundwork needed to start reading volume.

Everything in this book, from footprint charts to volume profiles, is built on the data that gold futures produce. You cannot read that data well without knowing how the contract works, how it is priced in ticks, when its day begins and ends, and what each trade costs. All contract details below were checked in October 2026; exchanges change margins, fees and sometimes rules, so check CME Group's current contract specifications before relying on any number.

What Is a Futures Contract?

A standard deal: price today, delivery later

A **futures contract** is a **standardised** agreement to buy or sell a set amount of something (for example, 100 troy ounces of gold) on a **set date in the future**, at a price agreed **today**.

Think of it as a pre-order with a fixed price. A jeweller who will need 100 ounces of gold in December and fears a rise can buy a December contract today; if gold rises, the gain on the futures offsets the dearer metal. A miner fearing a fall can do the opposite and sell. This is called **hedging**, the original reason futures markets exist. Traders with no interest in owning gold use the same contracts to take positions on the price, closing them before delivery; they provide much of the day-to-day activity.

Everything is fixed, except the price

"Standardised" is the key word. The exchange decides every detail of the contract in advance:

- the **contract size** (100 troy ounces for GC),
- the **quality** of the gold that can be delivered,
- the **delivery months** available,
- the **minimum price change**, called the **tick size**,
- the trading hours and the rules for delivery and settlement.

The only thing left for buyers and sellers to decide is **the price**. That is a powerful idea. Because every GC contract is identical, every trader in the world is trading exactly the same thing, and they

compete only on price. A December GC contract bought in Singapore is perfectly interchangeable with one sold in Chicago.

Key idea: In a futures contract, everything is standardised except the price. That is why one order book can serve the whole world: everyone is trading the identical product.

One exchange, one book, one clearing house

Gold futures trade on **COMEX**, which is part of **CME Group**, through CME's electronic trading platform, **Globex**. All buy and sell orders for a given contract go into **one central order book**. When a buy order and a sell order meet at the same price, a trade happens.

After the trade, a **clearing house** steps in between the two sides. The clearing house becomes the buyer to every seller and the seller to every buyer. You never have to trust or even know the person on the other side of your trade. If they fail to pay, the clearing house (backed by the margin system you will learn about in the next section) makes sure you are paid. Failing to meet one's obligations in a contract is called **default**, and the clearing house exists to remove that worry.

Why futures data is special

For this book, the most important consequence of the single central book is this: the exchange **records and publishes every real trade**, including its price, its size in contracts, its exact time and which side initiated it. That full record is the raw material of volume trading.

Spot gold (traded privately between banks and dealers) and CFDs (each broker with its own book) offer no such record of the whole market. Chapter 2 returns to this when it compares real volume with tick volume.

Most traders never take delivery

Many newcomers worry about this, so let us settle it. Gold futures are **physically delivered**: if a contract is held into its delivery period, real gold changes hands through exchange-approved vaults, following a precise process. A truck does not arrive at your house.

In practice, the vast majority of futures traders never take or make delivery. They close before the delivery period or **roll** into a later month (explained later in this chapter), and most retail brokers require this.

Long, short and an honest limit

In futures, going short is as simple as going long: you just enter a contract to sell. Because every contract has a buyer and a seller, total long contracts always equal total short contracts; the number still open is called **open interest**. Futures are leveraged, and losses can exceed your deposit; this chapter explains how the contract works, not whether to trade it.

Common mistake: "Futures are only for big institutions." Micro contracts such as MGC were designed specifically for smaller accounts. The price is the same; only the size is smaller.

Margin, Leverage and Daily Mark-to-Market

Margin is collateral, not the price

When you buy a futures contract, you do not pay its full value. The exchange sets an **initial margin**: an amount of money that must be in your account, set aside as collateral, to show you can cover likely losses. CME's own educational material describes futures margins as typically a small percentage of the contract's value, often in the range of roughly 3% to 12%.

The full value of the contract is its **notional value**:

Notional value = price × contract size

With gold at \$4,000.0 per ounce (roughly \$3,971 on XAUUSD, given the gap explained in Chapter 0):

- One **GC** contract (100 oz) has a notional value of $\$4,000 \times 100 = \text{\$400,000}$.
- One **MGC** contract (10 oz) has a notional value of $\$4,000 \times 10 = \text{\$40,000}$.

Controlling \$400,000 or \$40,000 of gold with a deposit that is a small fraction of that amount is **leverage**. And leverage, as Chapter 0 explained, magnifies both gains and losses.

Worked example: a 1% move

Gold moves 1%, which at \$4,000 is \$40 per ounce.

- On one GC: $\$40 \times 100 \text{ oz} = \text{\$4,000}$ gained or lost.
- On one MGC: $\$40 \times 10 \text{ oz} = \text{\$400}$ gained or lost.

Relative to a margin deposit that is only a few percent of notional value, a 1% move represents a large share of the deposit. If the margin on MGC were, hypothetically, \$2,000, a \$400 change would be 20% of it.

Mark-to-market: settling up every day

Futures have a feature that surprises many CFD traders: profits and losses are **not** left to pile up until you close. Every trading day, the exchange sets an official **settlement price** for each contract (for gold, based on trading around 13:30 New York time, explained later in this chapter). The clearing house then compares each open position with that settlement price and moves money accordingly. Accounts that lost value that day pay; accounts that gained value receive. This daily settling is called **mark-to-market**. In plain words: **losers pay winners every single day**.

Worked example: four days of mark-to-market

All margin figures below are **hypothetical**, chosen only to make the arithmetic clear. Real margins change; check CME and your broker.

You deposit \$3,000 and buy one MGC at \$4,000.0 (XAUUSD $\approx$ \$3,971). Suppose the initial margin is \$2,000 and the **maintenance margin** (the minimum your account may fall to before action is required) is \$1,800.

Day	Settlement price	Change vs previous	Daily P&L (× \$10 per \$1)	Account balance
Start	4,000.0 (entry)	—	—	\$3,000
Day 1	3,970.0	−\$30.0	−\$300	\$2,700
Day 2	4,020.0	+\$50.0	+\$500	\$3,200
Day 3	3,890.0	−\$130.0	−\$1,300	\$1,900
Day 4	3,870.0	−\$20.0	−\$200	\$1,700

After Day 4, the account is \$1,700, below the \$1,800 maintenance margin. You receive a **margin call**: you must deposit enough to bring the account back up to the initial margin (\$2,000 in this example, so \$300), or reduce or close the position. If you do nothing, the broker will close the position for you.

Notice: on Day 2 you were \$200 ahead, yet two days later you faced a margin call, after one large but far from unheard-of \$130 day. Leverage turns ordinary volatility into urgent account problems. In practice, many retail brokers do not wait for settlement; their systems watch your account in real time and may close positions automatically during the session.

Margins change, and brokers add their own rules

Margin is not a fixed number. Exchanges raise margins when markets become more volatile; in late 2025, CME raised precious-metals margins more than once. Brokers also set their own **day-trade margins** (lower requirements for positions closed within the session), which are internal policy, not exchange rules, and can change at any moment. Whenever you see a margin figure, ask "as of when, and from whom?"

Common mistake: "My maximum loss is the margin I deposited." False. Margin is a minimum deposit, not a cap on losses. A fast move can produce a loss larger than your whole account, and you will owe the difference.

Gold note: Because a standard GC contract moves \$100 for every \$1 change in gold, a typical day's range can be worth thousands of dollars per contract. That is why MGC, one-tenth the size, is much more suitable for learning and for managing position size in smaller accounts.

Gold Futures: GC vs MGC, Tick Size and Tick Value

The two contracts

This book works with two COMEX gold futures contracts:

- **GC**, the standard gold futures contract: **100 troy ounces**.

- **MGC**, Micro Gold futures: **10 troy ounces**, exactly one-tenth of GC.

Both are quoted in **US dollars per troy ounce**, and both have the same **tick size**: the smallest price change allowed is **\$0.10 per ounce**. Gold can trade at \$4,000.0 or \$4,000.1, but never at \$4,000.05.

Tick value: what one tick is worth

The **tick value** is how many dollars one tick of movement is worth for one contract:

Tick value = tick size × contract size

- GC: $\$0.10 \times 100 \text{ oz} = \textbf{\$10 per tick}$
- MGC: $\$0.10 \times 10 \text{ oz} = \textbf{\$1 per tick}$

	GC	MGC
Contract size	100 troy oz	10 troy oz
Tick size	\$0.10 per oz	\$0.10 per oz
Tick value	\$10	\$1
Value of a \$1 move	10 ticks = \$100	10 ticks = \$10
Notional value at \$4,000	\$400,000	\$40,000

A \$1 move in gold, say from \$4,000.0 to \$4,001.0 (XAUUSD roughly \$3,971 to \$3,972), is ten ticks. That is \$100 on one GC and \$10 on one MGC.

The P&L formula

For any futures position:

P&L = (price change ÷ tick size) × tick value × number of contracts

The sign depends on direction: a long gains when the price rises; a short gains when it falls.

Worked example 1. Long 1 MGC from 4,000.0 to 4,002.5. Price change = +2.5. Ticks = $2.5 \div 0.10 = 25$. P&L = $25 \times \$1 \times 1 = \textbf{+\$25}$.

Worked example 2. Short 2 GC from 4,015.3 to 4,011.8. Price change = -3.5, which is a gain for a short. Ticks = 35. P&L = $35 \times \$10 \times 2 = \textbf{+\$700}$.

Worked example 3. Long 3 MGC from 4,008.0 to 4,001.6. Price change = -6.4. Ticks = 64. P&L = $64 \times \$1 \times 3 = \textbf{-\$192}$.

Try it: Gold falls from 4,020.0 to 4,012.5. What is the P&L on one short GC, and on one short MGC? (Answer: 75 ticks, so +\$750 on GC and +\$75 on MGC.)

Separate books, separate volume

GC and MGC are **two separate markets**. Each has its own order book, its own traders and its own volume. Their prices move almost exactly together, because if they drifted apart, traders would immediately buy the cheaper one and sell the dearer one, pulling them back in line. This kind of trading is called **arbitrage**.

But their volumes are very different. GC is usually deeper (more orders waiting at each price) and is where larger participants tend to trade. For that reason, many order-flow analysts **read GC data even if they trade MGC**. Throughout this book, our charts are built from GC data by default.

Tick size also matters because footprint charts and volume profiles, which you will meet later, display price in steps of one tick or a small group of ticks; you need to know what each step is worth.

Common mistake: "A tick in gold is \$1." No. The tick **size** is \$0.10 (ten cents per ounce). The tick **value** is \$10 on GC or \$1 on MGC. Keep the two ideas separate.

Common mistake: "MGC volume is the same as GC volume." They are two different books with different participants. Always know which contract your volume data comes from.

Contract Months, Rollover and Continuous Charts

More than one gold contract at once

Because every futures contract has an expiry date, several gold contracts trade at the same time, one for each listed delivery month. Each has a short code made of the product symbol, a month letter and a year. For example, **GCZ26** is the GC contract for December 2026.

The standard futures month codes are:

Month	Code	Month	Code
January	F	July	N
February	G	August	Q
March	H	September	U
April	J	October	V
May	K	November	X
June	M	December	Z

For gold, the main contract months are **February (G), April (J), June (M), August (Q) and December (Z)**. October (V) also trades, and CME's settlement rules have at times treated it as a

lead month too; check CME's current notices for details.

The active month holds the volume

Although several months trade at once, almost all of the activity concentrates in one of them: the **active month** (also called the **front month** or **lead month**). When this book was written in October 2026, the active gold contract was December 2026, GCZ26. The other months trade, but much more thinly.

Rolling: moving to the next month

Some weeks before a contract's delivery period begins, traders who want to keep their positions close them in the expiring month and open the same positions in the next active month. This is called **rolling** (or **rollover**). For gold, traders typically roll in the weeks before the delivery month starts; the December contract, for example, gives way to February late in November.

During the roll, a very visible thing happens to volume. Activity in the old contract shrinks quickly, and activity in the new contract grows quickly. If you plotted daily volume for both contracts, you would see two lines crossing like an X: the moment when the new month's volume overtakes the old one's.

If you keep watching the old contract during roll week, you will see volume collapse. It is tempting to read this as "interest in gold has dried up", but the traders simply moved next door. This **volume migration** is a mechanical event, not a market signal.

Common mistake: "Volume suddenly halved, so traders have lost interest in gold." Check the calendar first. During roll periods, volume moves to the next month; it does not vanish.

Continuous charts and their artificial gaps

Most charting platforms offer **continuous charts** (for example, symbols such as GC1! or @GC, depending on the platform). A continuous chart joins consecutive contracts end to end so that you can look at years of history on one chart.

The problem is that different months trade at different prices. Because of the cost of carry (the same effect behind the GC-XAUUSD gap in Chapter 0), a later month usually trades above an earlier one. When a continuous chart switches from December to February, the price can jump by several dollars even though nothing happened in the market. Unless the data is adjusted, that jump appears as an **artificial gap**.

Worked example. Suppose on roll day, December gold trades at 4,000.0 and February gold trades at 4,028.0 at the same moment. An unadjusted continuous chart that switches from December to February that day will show a sudden \$28 jump, which is 280 ticks, or \$2,800 per GC contract, that no trader actually experienced.

Some platforms fix this with **back-adjustment**: shifting all older prices by the gap so the joins are smooth. That fixes the price line, but old prices are then no longer the prices that actually traded,

and back-adjustment does **not** fix volume or delta: each day's volume belonged to whichever contract was active that day.

The practical rule: for footprint charts, volume profiles and cumulative delta, read **the specific, active contract** (for example, GCZ26), not a continuous chart; be suspicious of volume jumps on roll days; and learn your platform's roll rules, since platforms switch on slightly different dates.

Common mistake: "There is a gap on the continuous chart, so there must have been news." It may simply be where two contracts were glued together. Likewise, a multi-month volume profile on a continuous chart has shifted or joined prices, so its levels are not exact.

When Gold Trades: Globex Hours, Sessions and Settlement

The nearly 23-hour clock

COMEX gold trades electronically on CME Globex almost around the clock:

- From **Sunday 18:00 to Friday 17:00, New York time (ET)**,
- with a **daily 60-minute halt from 17:00 to 18:00 ET**.

(Chicago is one hour behind New York, so in Chicago time the halt is 16:00 to 17:00. This book uses New York time throughout.)

One detail confuses almost everyone at first: the CME **trading day starts at 18:00 ET on the previous evening**. So the trading that happens on Monday evening in New York, which is Tuesday morning in Asia, belongs to **Tuesday's** trading day. Sunday evening's open begins Monday's trading day. When a chart shows "Tuesday's" volume profile or delta, it usually starts at 18:00 ET on Monday.




Figure 3. One trading day starts at 18:00 New York time and runs through Asia, London and New York; the busiest hours are London and the New York morning.

Three sessions

Traders divide the gold day into three broad **sessions**. The boundaries are conventions and vary slightly between sources:

- **Asia** runs from the 18:00 ET reopen until around London's opening. Activity comes from Tokyo, Singapore, Hong Kong and Shanghai (home of the Shanghai Gold Exchange). Volume is usually lower.
- **London** begins around 03:00 ET. London is the centre of the physical gold market, and volume typically rises when it opens.
- **New York** is often counted from 08:20 ET, the old opening time of the COMEX trading floor (when gold was traded face to face in a "pit"), through to the settlement at 13:30 ET. Major US economic data is often released at 08:30 ET. The **London–New York overlap**, the morning in New York, is usually the busiest and most volatile part of the day.

After the settlement and into the late afternoon, activity usually thins out until the 17:00 halt.

Settlement: gold's official daily price

Each day the exchange sets an official **settlement price** for every gold contract. For GC, the daily settlement is based on trading in the active month on Globex during the window from **13:29 to 13:30 ET**: in essence, the volume-weighted average price of trades in that minute. (CME publishes the exact procedure; check the current version.)

Notice the time: settlement happens around 13:30, **hours before** the market's 17:00 daily close. Trading continues after settlement, but the official price used for mark-to-market (see the earlier section) has already been set.

Common mistake: "The daily close is the price at 17:00." For the exchange's official profit and loss, the settlement at about 13:30 ET is what counts. Your chart's "daily close" may be something else depending on how your platform defines the day.

Why the clock matters for volume

Because activity varies so much through the day, a volume bar that looks "huge" at 02:00 ET might be ordinary at 09:30 ET. Compare volume **against the same time of day or the same session** on other days. Volume tools also need to know when "the day" starts: a profile or cumulative delta that resets at midnight gives different results from one that resets at 18:00 ET.

Daylight saving time

The United States and Europe switch their clocks on different dates. In 2026, for example, Europe moves its clocks back on 25 October, while the United States does so on 1 November. For that week, the time difference between London and New York is four hours instead of the usual five, so London's open shifts by an hour relative to New York time. A similar mismatch happens each spring. If session-based patterns look slightly "off" on those days, the calendar may be the reason.

Common mistake: "The London open always makes the high or low of the day." This is a popular claim, not proven. Treat it as folklore unless you test it yourself.

Common mistake: "Low volume in Asia means Asian moves do not matter." Low volume means the market is thinner and more fragile, not that it is irrelevant. Thin markets can move further on less trading.

Try it: On a 15-minute GC chart of one full trading day (18:00 ET to 17:00 ET), mark 03:00, 08:20, 08:30 and 13:30 ET, and watch the volume panel rise into London, peak in the New York morning and fade after settlement.

What a Candle Shows — and What OHLC Hides

A candle is four numbers

A **candlestick** (or simply **candle**) summarises all the trading in a period of time, for example five minutes, using just **four numbers**:

- **Open:** the first traded price in the period.
- **High:** the highest traded price.
- **Low:** the lowest traded price.
- **Close:** the last traded price.

Together these are called **OHLC**. The thick part of the candle, the **body**, spans from the open to the close. The thin lines above and below, the **wicks** (or **shadows**), reach up to the high and down to the low. If the close is above the open, the candle is usually drawn in green or white (a rising candle); if below, in red or black (a falling candle).




Figure 4. A candle records only open, high, low and close for its period. The body spans open to close; the wicks reach the high and low.

Worked example. A five-minute GC candle opens at 4,003.2, trades up to 4,006.0, down to 4,001.4 and closes at 4,005.1 (XAUUSD roughly \$29 lower on each of those numbers). The body runs from 4,003.2 to 4,005.1, a rising candle of 19 ticks. The upper wick runs from 4,005.1 to 4,006.0 (9 ticks); the lower wick from 4,001.4 to 4,003.2 (18 ticks). The full range is 46 ticks, or \$460 on one GC.

This compression lets you see days of trading at a glance, but it throws a lot away.

Hidden #1: the path

A candle with a given OHLC could have been formed in many different ways. In our example, did the price go down to 4,001.4 first and then rise to 4,006.0? Or did it rise to 4,006.0 first, collapse to 4,001.4 and then recover to 4,005.1? The candle looks identical either way. It cannot tell you whether the high came before the low.




Figure 5. Two very different sessions can print the exact same candle. Volume tools exist to recover what OHLC throws away.

Hidden #2: how much traded at each price

Did most of the trading happen near the high, near the low or in the middle? The candle does not say. That long lower wick might have been built by three contracts or by three thousand. Those are completely different events: in one case, a handful of trades briefly touched a low price; in the other, a large amount of business was done there.

Hidden #3: who was aggressive, and how hard

Did buyers push the price up by aggressively paying the ask with market orders? Or did sellers simply step back, so the price drifted up through an empty space with little real buying? A candle draws both situations the same way. (A standard candle does not even include volume; platforms show it as a separate panel, and only as one total per candle.)

Same candle, different story

Picture two five-minute GC candles from two New York mornings with almost identical OHLC. In the first, most volume traded near the low, with aggressive buyers dominating. In the second, the price rose on thin volume and most trading happened near the high, with aggressive sellers active there. Twins on a candle chart; opposite stories underneath.

Key idea: Candles show **what** happened. Volume shows **who**, **where** and **how hard**. Tools such as footprint charts (later chapters) reopen the candle and restore the information it threw away.

An honest limit

More information is not automatically better prediction. Information from inside the candle has to be read correctly, and even then it describes the past more precisely; it does not reveal the future. In our own testing, volume-based signals on their own did not predict gold's direction after costs. Think of this information as context, not as a fortune-teller.

Common mistake: "A big green candle means buyers were strong." It might mean sellers were simply absent. The candle cannot tell the difference.

Common mistake: "A long wick always means strong rejection." This is a popular claim, not proven. You need to see the volume behind the wick before you can say anything about it.

What Is Liquidity?

A practical definition

Liquidity is one of the most used and least understood words in trading. A widely used definition comes from Larry Harris's book *Trading and Exchanges*. In essence, liquidity is the ability to trade a large size, quickly, at low cost, whenever you want to.

That definition has four dimensions, and it helps to picture a swimming pool:

- **Width** is the cost of trading, measured mainly by the **bid-ask spread**. In the pool picture, it is how far you have to reach to get to the water. A narrow spread means cheap trading.
- **Depth** is how much is waiting to be traded at each price: the **limit orders** sitting in the order book. It is how deep the water is.
- **Immediacy** is how quickly you can execute a given size. It is how fast the pump can draw water.
- **Resiliency** is how quickly prices and depth return to normal after a large trade. It is how fast the pool refills after someone removes a bucket of water.

Who provides liquidity and who consumes it

This is one of the most important ideas in the whole book, so it is worth stating plainly:

- **Limit orders** (orders to buy or sell at a specific price or better, which wait in the book until someone trades with them) **provide** liquidity. They are **passive**: they sit and wait.
- **Market orders** (orders to buy or sell immediately at the best available price) **consume** liquidity. They are **aggressive**: they take what is available.

A market order is the bucket that removes water from the pool. Price moves when aggressive orders eat through all the liquidity waiting at a level, or when the passive orders at that level are pulled or moved away. Chapter 2 builds directly on this mechanism.

Worked example. The best ask in GC is 4,000.2 with 8 contracts waiting. An aggressive buyer sends a market order for 8 contracts. All 8 are filled at 4,000.2, and that level is now empty. The new best ask becomes the next level up, say 4,000.3. The price has moved up one tick, not because anyone "decided" it should, but because the liquidity at 4,000.2 was used up.

Liquidity changes through the day

In gold, the London–New York overlap usually brings the deepest book and tightest spread; after settlement the book is usually thinner. Around major releases, such as the US jobs report or inflation data, many liquidity providers pull their limit orders seconds before the number. The book goes thin, and the price can jump several dollars on comparatively little volume.

Visible depth is a promise, not a guarantee

The limit orders you see in the book are intentions and can be cancelled in a fraction of a second. Placing orders you intend to cancel in order to mislead others (**spoofing**) is illegal, but ordinary cancellation is legal and extremely common. Visible depth is real at that instant, not a commitment.

Common mistake: "A liquid market is one where the price does not move." No. Liquidity means trading is cheap and easy. Liquid markets can still move a lot.

Common mistake: "High volume means high liquidity." Volume is what **has traded**. Depth is what is **waiting** to trade. They are related but different.

Common mistake: "A large order shown in the order book will definitely trade." It may be cancelled before the price ever reaches it.

The Real Cost of Trading: Spread, Slippage, Commission

Every trade comes with a bill. Many beginners never add it up. This section shows you how.

Cost 1: the spread

The **spread** is the distance between the best **bid** and the best **ask**. If you buy with a market order, you pay the ask. If you then sell immediately with a market order, you receive the bid. In GC during busy hours, the spread is typically **one tick**. So simply entering and exiting immediately costs you one tick: **\$10 on GC, \$1 on MGC**.

The spread is the price of immediacy: the fee you pay for not waiting.

Cost 2: slippage

Slippage is the difference between the price you expected and the price you actually got. It happens when:

- your order is larger than the size available at the best price,
- the market is moving fast (for example, at a news release), or
- a **stop order** is triggered. A stop order waits until a set price is reached and then becomes a market order, filling at whatever price is available.

Worked example. You want to buy 5 GC at market. The best ask is 4,000.1. But the book above looks like this:

Ask price	Contracts waiting
4,000.3	3
4,000.2	2
4,000.1	2

Your order takes everything at 4,000.1 (2 contracts), then everything at 4,000.2 (2 contracts), then 1 contract at 4,000.3.

- 2 contracts at 4,000.1: 0 ticks of slippage
- 2 contracts at 4,000.2: 1 tick each = 2 ticks
- 1 contract at 4,000.3: 2 ticks
- Total slippage: 4 ticks × \$10 = **\$40**, and your average fill is 4,000.18 rather than the 4,000.1 you expected (XAUUSD equivalent roughly \$3,971).

On a news release, with a much thinner book, the same order could slip much further.

Cost 3: commissions and fees

You also pay:

- your **broker's commission**,
- the **exchange fee** charged by CME,
- **regulatory fees** (for example, a small NFA fee in the United States),
- and sometimes **platform and market data fees**.

Commissions and exchange fees are charged **per contract, per side** (so once to enter and once to exit). They differ between brokers and change over time; get exact numbers from your broker's fee schedule and from CME.

The round-trip cost

A **round trip** is a complete trade, in and out:

Round-trip cost = spread (if you enter and exit with market orders) + average slippage + fees on both sides

Worked example. The fee figures below are **hypothetical**, for illustration only.

	GC	MGC
Spread (1 tick)	\$10	\$1
Average slippage (assume 0.5 tick)	\$5	\$0.50
Fees, both sides (hypothetical)	\$5	\$1.50
Total in dollars	\$20	\$3
Total in ticks	2 ticks	3 ticks

Notice that the micro contract can cost **more in ticks**, because fixed fees are a larger share of a smaller contract. Always convert your costs into ticks; that is the number you compare with how far you expect the price to move.

Small targets, big cost share

Suppose a trader aims to capture 4 ticks and the round trip costs 2 ticks: half of the hoped-for gain is gone before the market has done anything. A pattern whose average advantage on paper is a fraction of a tick is wiped out by a single tick of spread.

This is exactly where our own research reached its honest conclusion: **order-flow signals on their own did not predict gold's direction after costs**. Many "patterns" show a tiny paper edge that disappears once spread and fees are counted. That does not make order flow useless; it means costs belong in every evaluation from the start.

Limit orders: no spread, but other risks

A **limit order** does not pay the spread, because it waits at its price rather than crossing to the other side. But it has two costs of its own:

- **Fill risk:** the market may never come to your price, and you miss the trade entirely.
- **Adverse selection:** limit orders tend to get filled precisely when the price is moving against you. If you place a limit buy at 3,995.0 and it fills, it is because sellers pushed the price down

to you, and they may keep pushing. Trades you most wanted often run away without filling you; trades that fill you easily are often the ones that keep going the wrong way.

Common mistake: "A backtest without costs is still valid." No. This is where most strategies fail. Always test after costs.

Common mistake: "A stop order guarantees my exit price." It does not. Once triggered, it becomes a market order and can slip, sometimes badly in fast markets.

Try it: Look up your own broker's commission and CME's current exchange fee for GC and MGC. Add one tick of spread and half a tick of slippage. What is your round-trip cost in ticks for each contract?

What "Order Flow" and "Volume Trading" Actually Mean

Two definitions

Order flow is the study of the stream of orders and trades in a market: who traded (in the sense of size and side), at what price, how large, and with what type of order.

Volume trading (or **volume analysis**) is a broader umbrella. It includes every method based on volume: volume per candle, **volume profile** (volume at each price level), **delta** (aggressive buying minus aggressive selling), **footprint charts** and more. Order flow is a part of volume analysis focused on the finest detail: individual orders and trades.

Intent versus actual

A very useful way to organise order-flow data, popularised by order-flow educators such as Jigsaw Trading, is to split it into two kinds:

Intent: what people say they will do. These are the **limit orders waiting in the order book**. They show where participants say they are willing to trade. They can be cancelled at any time. The tools that display them are the **DOM** (depth of market, a ladder showing waiting orders at each price) and the **heatmap** (a picture of the order book over time, as shown by platforms such as Bookmap).

Actual: what people did. These are the **trades that really happened**. They cannot be cancelled. The tools that display them are **Time & Sales** (also called **the tape**, a running list of every trade), **footprint charts**, **delta** and **cumulative delta (CVD)**, and **volume profile**.

	Intent	Actual
What it is	Orders waiting to trade	Trades that happened
Can it be cancelled?	Yes, at any moment	No, it is final
Typical tools	DOM, heatmap	Time & Sales, footprint, delta, CVD, volume profile
Question it answers	Where do people say they will trade?	Where did they trade, and who was aggressive?

Part of the skill is noticing when the two disagree: a large order that vanishes as the price approaches is intent that never became actual.

Volume trading means looking past the shape of the price and asking: who built this move, and how hard did they push? Did aggressive orders drive it, or did the other side step back? Where did heavy trading happen that might matter again later, and where did price pass through with little trading at all?

What "who" really means

When order-flow traders say "who", they do not mean names. Exchange data does not reveal the identity of traders. What we can see is the **size** of each trade, its **price**, its **time** and which side was **aggressive**. From those facts we make **inferences**, for example, "a lot of aggressive selling met resting buy orders here and the price did not fall". Inferences can be wrong, and an honest analyst treats them as hypotheses, not facts.

Context, not a crystal ball

The honest frame of this book bears repeating: **candles show what happened; volume shows who, where and how hard**. That is more information, not a crystal ball. Our own testing, and the experience of serious educators in this field, suggest that order flow is most useful for **context, timing and managing trades**, rather than for predicting direction on its own.

Common mistake: "Order flow means seeing the banks' orders." Identities are invisible. You see the footprints of trades, not the names behind them.

Common mistake: "Order flow shows the future." It shows the very recent past in great detail, and it is a whole family of data and methods, not a single indicator. The future remains a matter of probabilities.

Key idea: Intent can be cancelled; actual cannot. Learning to read both, and the gap between them, is the core skill that the rest of this book develops.

Chapter summary

- A **futures contract** is a standardised agreement to trade a set amount at a future date at a price agreed today; only the **price** is negotiated.
- Gold futures trade on **COMEX** (CME Group) via **Globex**, in **one central order book**, with a **clearing house** guaranteeing both sides. Every trade is published, which makes real volume analysis possible.
- **Margin** is collateral, not the price. **Leverage** magnifies gains and losses. **Mark-to-market** settles them daily against the **settlement price**; below **maintenance margin** comes a **margin call**. Losses can exceed your deposit, and margins change.
- **GC** = 100 oz, **MGC** = 10 oz, **tick size** \$0.10, **tick value** \$10 and \$1. **P&L = (price change ÷ tick size) × tick value × contracts**. The two have separate books and volume.
- Volume concentrates in the **active month** and migrates at the **roll**. **Continuous charts** can show artificial gaps; read volume on the specific contract.
- Gold trades **Sunday 18:00 to Friday 17:00 ET**, halting **17:00–18:00 ET** daily; the CME day starts at **18:00 ET the previous evening**. Sessions: **Asia, London, New York**. GC settles **13:29–13:30 ET**. Compare volume with the same time of day.
- A **candle** shows OHLC but hides the **path**, the **volume at each price** and **who was aggressive**.
- **Liquidity** = **width, depth, immediacy, resiliency**. **Limit orders provide it**; **market orders consume it**.
- Every trade pays **spread, slippage and fees**; measure the **round-trip cost in ticks**.
- **Intent** (the book) can be cancelled; **actual** (trades) cannot. Order flow is context, not a crystal ball.

Checklist

- I can explain what a futures contract is and what the clearing house does.
- I can walk through a mark-to-market example and say when a margin call happens.
- I know the tick size and tick values for GC and MGC and can calculate P&L for any position.
- I can name the active gold contract right now, and I use the specific contract, not a continuous chart, for volume work.
- I know gold's trading hours, the daily halt, when the CME day starts and when settlement happens, in New York time.
- I can list three things a candle hides and the four dimensions of liquidity.
- I have calculated my own round-trip cost in ticks for GC and MGC.
- I can explain intent versus actual order-flow data, with tools for each.

Quiz

1. In a futures contract, which of these is **not** standardised by the exchange: contract size, delivery month, price or tick size?
2. Gold rises from 4,000.0 to 4,002.5. What is the P&L on one long MGC? And on one long GC?
3. During roll week, volume in the old gold contract collapses. What is the most likely reason?
4. When does the CME trading day for gold begin, and when is GC's daily settlement price set (New York time)?
5. You buy 1 GC at the ask and immediately sell at the bid. The spread is 1 tick. Before fees, how much have you lost?

Quiz answers

1. **The price.** Everything else is fixed by the exchange; traders compete only on price.
2. The move is $2.5 \div 0.10 = 25$ ticks. On one MGC: $25 \times \$1 = \mathbf{+\$25}$. On one GC: $25 \times \$10 = \mathbf{+\$250}$.
3. **Traders have moved their positions to the next contract month.** Volume migrated; it did not disappear, and it says nothing about interest in gold.
4. The trading day begins at **18:00 ET on the previous evening**. GC's daily settlement is based on trading during **13:29–13:30 ET**, hours before the 17:00 ET daily close.
5. **\$10**: one tick on one GC contract. On MGC, the same round trip would cost \$1 before fees.

2

Why Volume? Price Is the Result

In this chapter: You will learn what volume-based trading adds to ordinary price-chart trading, and what it does *not* add. We follow the chain from decisions to orders to trades to price; see why moving averages, RSI and MACD always lag; treat the market as a two-sided auction; read the order book; and learn why every trade has one buyer and one seller, of whom only one was in a hurry. Then we compare a pattern trader and a volume trader on the same gold chart, clear away popular myths, separate real exchange volume from "tick volume", and map the platforms and data feeds. The theme in one sentence: *candles show WHAT happened; volume shows WHO, WHERE and HOW HARD.*

Price Is the Result, Volume Is the Cause

Most people meet markets through a price chart. A line or a row of candles wiggles up and down, and it is natural to think of price as something that moves by itself, like the weather. That picture is wrong, and fixing it is the first step towards understanding volume.

Every price on your chart is a trade

On a futures exchange, the number your chart plots is the **last traded price**, often just called the **last**. It is the price at which the most recent trade took place. A **trade** happens only when two orders **match**: someone willing to buy at a price meets someone willing to sell at that same price. If nobody trades, the last price does not change, no matter how many people are watching, hoping or posting opinions on social media.

So when you see gold futures (GC) print 4000.1 after 4000.0, what you are actually seeing is this: at least one contract changed hands at 4000.1. The chart did not "decide" to go up. Two parties agreed on a price, and the exchange recorded it.

Gold note: Throughout this chapter we use round GC prices near 4,000 for the examples. Spot gold CFDs (XAUUSD) usually trade a little below the futures price, because a futures contract includes a "carry" cost until delivery. In recent months the gap has been roughly \$27–30, so GC 4000.0 corresponds to approximately 3,970–3,973 on XAUUSD. The gap changes over time and from broker to broker, so always treat the equivalent as approximate.

The chain: decisions → orders → trades → price

Think of what happens behind every printed price as a chain of four links:

1. **Decisions.** People and algorithms decide to buy or sell. These decisions come from many places: a central-bank statement, a fund rebalancing, a jewellery manufacturer hedging, a trader's stop being hit, an algorithm following a rule.

2. **Orders.** Each decision becomes an **order**: an instruction sent to the exchange to buy or sell a number of contracts, either immediately or at a specific price.
3. **Trades.** When a buy order and a sell order meet at the same price, a trade happens. The sum of contracts traded is **volume**.
4. **Price.** The price of the latest trade becomes the new "last", and that is what your chart draws.

Candles and indicators live at the very end of this chain. They are summaries of the price link. Volume and order flow live one step earlier, at the trades link, and the order book (which you will meet later in this chapter) shows part of the orders link. Each step back along the chain brings you closer to the reasons price moved, although you never reach the decisions themselves.

Key idea: Price is the *result* of orders meeting. Volume, the record of those meetings, sits one step closer to the cause. That is the entire reason volume analysis exists.

Two ways to move up one tick

Here is a concrete example that shows why the trades link matters.

Imagine GC is quoted with buyers waiting at 4000.0 and sellers waiting at 4000.1. (In the next sections you will learn to call these the **bid** and the **ask**.) Suppose 60 contracts are offered for sale at 4000.1. Now the price "ticks up": the best price at which you can buy moves from 4000.1 to 4000.2. There are two very different ways this can happen.

Story A — the push. Impatient buyers arrive and buy everything on offer at 4000.1. Over a few seconds, aggressive buy orders totalling 200 contracts hit the market. The first 60 consume all the sellers at 4000.1, and the rest start filling at 4000.2. The offer at 4000.1 has been "eaten", so the best ask moves up. Price rose because buyers *pushed* it.

Story B — the pull. Nobody buys much at all. Instead, the sellers resting at 4000.1 change their minds. Perhaps a news headline makes them nervous, or an algorithm decides to step back. They cancel their sell orders. Now the cheapest seller is at 4000.2. A trader buys 5 contracts there, and the last price prints 4000.2. Price rose because sellers *pulled* away.

On a candle chart, both stories look identical: one tick up. But they mean different things. In Story A there were 200 contracts of aggressive buying; in Story B there were 5. Volume, and the split between aggressive buying and aggressive selling (called **delta**, which Chapter 3 covers in detail), tells these stories apart. The candle cannot.

	Story A: the push	Story B: the pull
What moved price	Buyers consumed all offers at 4000.1	Sellers cancelled their offers at 4000.1
Contracts traded	~200	~5
Aggressive side	Strongly buyers	Barely anyone
What the candle shows	+1 tick	+1 tick

Try it: Next time you watch a gold chart with a Time & Sales window open, find two moments where price moved by one tick. Note how many contracts traded in each move. You will often see that some ticks are "bought" with hundreds of contracts and others happen almost silently.

What "cause" does and does not mean

It is tempting to jump from "volume is closer to the cause" to "volume predicts the future." Be careful. There are three honest limits.

First, *explaining is not predicting*. Volume tells you what forced the latest move; it does not tell you what the next people will decide.

Second, *part of the cause is invisible in volume*. Cancelled orders, like the sellers in Story B, never become trades and leave no mark on a volume bar.

Third, *the biggest decisions are made away from the chart*. Rates, the dollar, geopolitics and fund flows drive gold over days and weeks; order flow shows how those decisions hit the market, not the decisions themselves.

In our own testing, order-flow signals on their own did not predict gold's direction once realistic costs were included. So this book treats volume as **a window into the mechanism**, not a crystal ball: more information and better questions.

Common mistake: "Volume always moves before price." This is a popular claim, not proven. Sometimes heavy volume comes first; often it arrives together with the move or after it.

Common mistake: "Price can only go up if people buy." As Story B shows, price can rise when sellers simply withdraw.

Why MA, RSI and MACD Lag (and What They Throw Away)

Most traders start with indicators: a moving average on the chart, RSI or MACD in a panel underneath. These tools are popular for good reasons, they are simple and they smooth out noise, but they have a built-in property that every volume trader should understand. They are mathematical transformations of past closing prices. That makes them late by design, and it means they ignore most of what happened in the market.

The moving average

A **simple moving average (SMA)** of length N is the average of the last N closing prices. A 20-period SMA on a 5-minute chart averages the closes of the last 20 five-minute candles, which is

100 minutes of history.

Here is a small worked example on gold. Suppose GC has closed at exactly 4000.0 for five candles in a row. Then it jumps to 4010.0 and stays there. Watch a 5-period SMA:

Candle	Close	5-period SMA
1-5	4000.0 each	4000.0
6	4010.0	$(4000 \times 4 + 4010) \div 5 = 4002.0$
7	4010.0	$(4000 \times 3 + 4010 \times 2) \div 5 = 4004.0$
8	4010.0	4006.0
9	4010.0	4008.0
10	4010.0	4010.0

The price moved \$10 in one candle (that is 100 ticks, worth \$1,000 on one GC contract or \$100 on one MGC). The average needed five candles to catch up fully. With a 20-period average it would need 20 candles. This delay is called **lag**, and it is not a flaw in a particular indicator: it is what averaging *is*. You cannot remove noise by averaging without also delaying the signal.

An **exponential moving average (EMA)** gives more weight to recent closes, so it reacts faster. But faster is not the same as instant. An EMA still blends the present with the past, so it still lags.

RSI in one formula

The **Relative Strength Index (RSI)** compares the average size of up-closes with the average size of down-closes over N candles (14 by default):

- $RS = \text{average gain} \div \text{average loss}$
- $RSI = 100 - 100 \div (1 + RS)$

Suppose over the last 14 five-minute candles, the average up-move in gold's close was \$2.00 and the average down-move was \$1.00. Then $RS = 2.0$ and $RSI = 100 - 100 \div 3 \approx 66.7$. If the average gain grows to \$3.00 with the same average loss, $RS = 3$ and $RSI = 75$.

Notice what the formula uses: closing prices only, and only how much each close changed from the previous one.

MACD

MACD (Moving Average Convergence Divergence) is the difference between a 12-period EMA and a 26-period EMA of the close. A 9-period EMA of the MACD line itself is drawn as the "signal line". MACD is sometimes described as a "leading" indicator. It is not: it is built entirely from moving averages, so it inherits their lag, twice.

What these indicators throw away

All three reshape one series, the closes; they create no new information. Picture a funnel: at the top are **all the trades** in a candle (every price, size, timestamp and aggressive side); one step down, the candle compresses them into **open, high, low, close (OHLC)**; one more step, most indicators keep only the close. Each step throws information away. In particular, classic price indicators ignore:

1. **Volume.** A close reached with 10 contracts and a close reached with 10,000 contracts have exactly the same weight.
2. **The path inside the candle.** The open, high, low and the order in which they happened are dropped.
3. **The aggressive side.** Whether buyers were lifting offers or sellers were hitting bids is invisible.
4. **Location.** At which prices most of the trading took place inside the candle is invisible.

Volume tools put those four things back. That is the honest case for them: not that they are magic, but that they use data the classic indicators never look at.

Key idea: Price indicators are the smoothed past. Order flow is the raw detail of the present. Both are information about what has already happened.

None of this makes moving averages useless: a long average is a reasonable way to define "trend" or filter market regimes. Nor is order flow lag-free; your feed, platform and reaction time all add delay. The real difference is the *type* of information.

Common mistake: "RSI above 70 means price must come back down." This is a popular claim, not proven. In a strong trend, RSI can stay above 70 for a long time while price keeps rising.

Try it: On a 5-minute GC chart, find a sharp move during the New York morning. Mark the candle where the move began and the candle where a 20-period moving average first turned or crossed. Count the candles between them. Then look at the volume of the candle where the move began.

The Market Is a Two-Sided Auction

If price is the result of trades, what organises the trades? The most useful mental model comes from **Auction Market Theory**.

Where the idea comes from

In the 1980s **J. Peter Steidlmayer**, a Chicago Board of Trade trader, developed **Market Profile**, which organises trading activity by price instead of time. **James Dalton**, in *Mind Over Markets* and

Markets in Profile, turned it into a broader framework now called Auction Market Theory. Chapter 7 covers it fully; here we need only its core.

The market's job is to facilitate trade

The central idea: **the purpose of a market is to facilitate trade**. Price moves in order to find the other side.

- When price looks **cheap** to many participants, buyers become eager. They keep raising the price they are willing to pay, and price **rises until it finds enough sellers**.
- When price looks **expensive**, sellers become eager, and price **falls until it finds enough buyers**.
- When both sides are content to trade within a range, a lot of business gets done there. That range is called **value**: the area the market currently accepts as fair.

A useful image is a pendulum swinging between "too cheap, buyers step in" and "too expensive, sellers step in". Most of the time it swings around the middle, where both sides agree.

Balance and imbalance

The market alternates between two states:

- **Balance**: price rotates around an accepted value area. On a chart this looks like a range. Lots of volume builds up at similar prices.
- **Imbalance**: one side is in control, and price travels directionally to look for a new value area. On a chart this looks like a trend.

A price that the market trades through quickly, with very little volume, suggests that the price was **not accepted**: few people wanted to do business there. A price where a great deal of volume builds up suggests **acceptance**: both sides were willing.

Why this matters for volume

Auction theory ties volume directly to the idea of acceptance:

- **High volume at a price** = both sides agreed to trade there.
- **Low volume and fast movement** = the market did not accept that price.

The tools in later chapters, such as the **volume profile**, the **value area** and the **point of control (POC)**, come straight from this idea. When you read a volume profile in Chapter 5, you are reading where the auction found agreement.

Picture three days on a 30-minute GC chart. On the first day gold rotates between roughly 3990 and 4010 (about 3,960–3,983 on XAUUSD), and the volume piles up in a bell shape around 4000. That is balance. On the second day gold leaves that range and travels to 4045, with thin, stretched-out volume along the way. That is imbalance: the auction searching for a new value. On the third day it settles into a new range around 4040. The bulge of volume shows where agreement was found; the thin stretch shows prices the market passed through without interest.

Key idea: The market is not trying to go up or down. It is trying to find someone to trade with. Volume shows where it succeeded.

A framework, not a formula

Auction Market Theory is a **descriptive framework**: a way of organising what you see. It is not a tested law that produces numerical predictions. Some well-known rules attached to it have not been firmly established by careful statistical testing.

Common mistake: "The value area is 70% of volume because that is a law of the market." The 70% figure is a convention, chosen as a rough analogy to one standard deviation of a bell curve. It is a definition, not a law of nature.

Common mistake: "The 80% rule: if price enters the value area and stays for two periods, there is an 80% chance it crosses to the other side." This is a popular claim, not proven.

Bid, Ask and the Order Book

To understand who moves price, you need to see where the waiting orders sit. That is the **order book**.

Bids, asks and the spread

The order book is the list of all **limit orders** waiting to be filled, sorted by price. (A limit order is an order to buy or sell only at a specific price or better; the next section explains it fully.)

- A **bid** is a waiting order to **buy**. The **best bid** is the highest price anyone is currently willing to pay.
- An **ask**, also called an **offer**, is a waiting order to **sell**. The **best ask** is the lowest price anyone is currently willing to sell at.
- The **spread** is the difference: best ask – best bid.

In GC during normal, busy hours, the spread is usually one tick: \$0.10 per ounce. For example, best bid 4000.0, best ask 4000.1. One tick is worth \$10 on a GC contract (100 troy ounces × \$0.10) and \$1 on an MGC contract (10 ounces × \$0.10).

The DOM: the order book as a ladder

Most order-flow platforms show the order book as a vertical **price ladder**, called the **DOM** (Depth of Market). Prices run down the middle. Waiting buy quantities sit on one side and waiting sell quantities on the other.




Figure 6. Resting limit orders: buyers (bids) below, sellers (asks) above. The gap between best bid and best ask is the spread.

Here is a simplified GC ladder:

Bid size (waiting buyers)	Price	Ask size (waiting sellers)
	4000.4	31
	4000.3	22
	4000.2	18
	4000.1	12
9	4000.0	
15	3999.9	
24	3999.8	
20	3999.7	

Reading it: the best bid is 4000.0 with 9 contracts waiting; the best ask is 4000.1 with 12 contracts waiting; the spread is one tick. Deeper in the book, 22 contracts are offered at 4000.3, 24 are bid at 3999.8, and so on. The quantity resting at each price is called **depth**.

Two names you will see:

- **Level 1** data shows only the best bid and best ask (and their sizes), plus the last trade.
- **Level 2**, or **market depth**, shows several price levels on each side. CME's main market-data feed publishes aggregated depth for a number of levels on each side (typically ten for many

products).

Worked example: what a market order does to the ladder

Using the ladder above, suppose a trader sends an order to buy 20 contracts immediately at the best available prices. What happens?

- The first 12 contracts fill at 4000.1, using up every seller there.
- The remaining 8 fill at the next price up, 4000.2, leaving 10 contracts still offered there.

The trader's average price is $(12 \times 4000.1 + 8 \times 4000.2) \div 20 = 4000.14$. If they had expected to pay 4000.1 for everything, the 0.04 difference per ounce is **slippage**. On 20 GC contracts that is $0.04 \times 100 \text{ oz} \times 20 = \$80$. The new best ask is 4000.2, and the last trade printed at 4000.2.

This is the mechanism behind "push" in the previous section. Price moved because an aggressive order used up all the resting supply at one level.

Who gets filled first: price, then time

On CME Globex, most products, including gold futures, match orders by **price-time priority**, also called **FIFO** (first in, first out). Better prices are filled first; among orders at the same price, the one that arrived earlier is filled first. If you join a bid of 9 contracts at 4000.0 with your own order of 1, you stand behind those 9 in the queue.

Time & Sales: where trades are recorded

Every trade is recorded on **Time & Sales**, also called **the tape**: a running list of time, price and size. Most platforms colour each line according to where it printed relative to the bid and ask:

- A trade at the **best ask** usually means an aggressive buyer arrived and bought from a waiting seller. Traders say the buyer **lifted the offer**.
- A trade at the **best bid** usually means an aggressive seller arrived and sold to a waiting buyer. Traders say the seller **hit the bid**.

This simple rule is the foundation for classifying volume as "buying" or "selling" in Chapter 3.

What the DOM cannot tell you

The DOM is a snapshot of the current moment, not a promise.

- It shows only **visible** orders. An **iceberg order** shows a small part of its true size and refills automatically as it is filled, so the displayed quantity understates what is really there.
- Orders are **added and cancelled constantly**, often many times per second. A large quantity can vanish before price ever reaches it. Ordinary cancellation is perfectly legal and extremely common. (Deliberately placing orders you intend to cancel, in order to mislead others, is called *spoofing* and is illegal.)

Common mistake: "The price on my chart is the price I can buy at." The chart shows the last trade. If you buy immediately, you pay the ask; if you sell immediately, you receive the bid.

Common mistake: "The bid is the price I sell at, so it is my price." The bid is what *other people* are offering to pay. You can sell to them there with a market order.

Common mistake: "A big wall of orders in the DOM will hold price." This is a popular claim, not proven. Large resting orders can be cancelled or simply eaten through.

Market vs Limit Orders: Aggressor vs Passive

Now we can name the two roles that every trade involves.

The main order types

- A **market order** says: "Fill me now, at the best price available." It buys immediacy, and in return it pays the spread and accepts whatever slippage the book causes. A market buy matches against the best ask; a market sell matches against the best bid.
- A **limit order** says: "Fill me only at this price or better." If it cannot be filled right away, it **rests** in the book and waits. Resting limit orders are what create depth; they **provide liquidity**.
- A **stop order** is dormant until price reaches a trigger level. A **stop-market** order then becomes a market order. So at the moment it is triggered, it acts aggressively, and it can suffer slippage. (A *stop-limit* order becomes a limit order instead, which caps the price but may not fill at all.)

A limit order priced so that it can trade immediately, for example a buy limit at 4000.3 when the best ask is 4000.1, is called a **marketable limit order**. It behaves aggressively up to its limit price.

Aggressor and passive

The aggressor, also called the **initiator**, is the order that causes a trade to happen at the moment it arrives. CME's own market-data documentation uses essentially this definition: the aggressor is the order that, when entered into the book, immediately triggers a trade. Usually it is a market order or a marketable limit order.

The passive side, or **resting** side, is the limit order that was already sitting in the book and got filled by the aggressor.



Trade: 5 lots at 4231.0, counted as BUY volume (aggressor = buyer)

The ask now shows 6. Delta for this trade: +5.

Figure 7. Every trade has a buyer and a seller; order flow labels it by the side that crossed the spread (the aggressor).

Worked example: the best ask in GC is 4000.1 with 80 contracts resting. A trader sends a market order to buy 80. All 80 fill at 4000.1. The buyer is the aggressor; the sellers whose limit orders were resting at 4000.1 are passive. Because the level is now empty, the best ask moves to 4000.2. On Time & Sales you would see a print of 80 at 4000.1, usually coloured as a buy, and the price ladder would shift up one tick.

Every trade has one buyer and one seller

This is the single most important sentence in order flow: **every trade has exactly one buyer and one seller**. Every contract bought is a contract sold. So it is impossible for "more contracts to be bought than sold." The totals are always equal.

What differs is **which side was in a hurry**. When order-flow traders say "500 contracts of buying came in," they mean 500 contracts where the **buyer was the aggressor**. The sellers of those 500 contracts were real too; they were simply waiting passively.

Key idea: "Buying" and "selling" in order flow always mean *aggressive* buying and *aggressive* selling. The number of contracts bought and sold is always equal.

How price actually moves

Putting it together:

- **Aggressors consume liquidity.** When they use up all the resting orders at a price, the best price shifts and the market moves.
- **Passive orders provide liquidity.** Large passive orders can soak up a lot of aggressive activity without letting price move. This is called **absorption** (Chapter 9 looks at it closely).
- **Passive orders can also move price by leaving.** When resting orders are cancelled or moved, the best price shifts without much trading at all (Story B from earlier).

Aggressive does not mean smart or big

It is easy to assume that the aggressive side is the "strong" or "informed" side. That is not necessarily so. Large participants often prefer to trade **passively**: resting limit orders avoid paying the spread and leave a smaller visible footprint. Their algorithms frequently split a large order into many small passive pieces. Meanwhile, a lot of aggressive volume comes from stop orders being triggered, which is forced, not chosen.

Common mistake: "The aggressor is always the smart money." This is a popular claim, not proven.

Common mistake: "Limit orders never move the price." Cancelling or moving limit orders can shift the best bid or ask just as surely as aggressive orders can.

Pattern Trader vs Volume Trader: Same Gold Chart

The best way to see what volume adds is to watch two traders look at the same moment. We will not declare a winner. The point is to see the different questions they ask.

The scene

It is the New York session. Yesterday's high in GC was 4035.0 (about 4,005–4,008 on XAUUSD). Gold rallies into that level and prints a 5-minute candle that opens at 4031.2, spikes to 4036.0, and closes at 4032.0. The candle has a long upper wick, called a **wick** or **shadow**: the thin line above the body showing prices that traded but did not hold.

The pattern trader's read

The price-pattern trader sees:

- price at a known **resistance** level (yesterday's high);
- a **pin bar**, a candle with a long wick and small body, often read as rejection;
- RSI(14) above 70.

Their conclusion is a hypothesis: "Resistance, rejection candle and an overbought reading. A pullback is more likely than not." Their data is the shape of the candle, the price levels and the indicators. Their underlying question is: "How has this pattern behaved in the past?"

The volume trader's questions

The volume trader looks at the same candle and asks different questions, using data the pattern trader does not see:

1. **Was this candle's volume high or low *for this time of day*?** Gold's volume has a strong daily rhythm, so a candle should be compared with candles at the same clock time on previous days (Chapter 3 shows how).
2. **How much traded inside the wick?** Did the prices from 4033 to 4036 trade 60 contracts or 1,450 contracts? A wick built on almost no volume means price briefly poked up where nobody did business. A wick built on very heavy volume means a lot of business happened there and price still failed to hold.
3. **Who was aggressive?** Was the candle's delta (aggressive buying minus aggressive selling) strongly positive? If many aggressive buyers arrived at the high and price still could not rise, someone large may have been selling passively against them: possible **absorption**. If almost no aggressive buyers arrived at all, the reversal may simply reflect a lack of interest.
4. **Is the session's running delta moving with price?** Is the cumulative delta (Chapter 3) rising with the rally, or flat while price climbs?

Question	Pattern trader	Volume trader
What is the main question?	What shape is this?	Who traded here, and how hard?
Main data	Candle shape, levels, indicators	Volume vs same hour, volume in the wick, delta, cumulative delta
What would change their mind?	Price closing back above the level	Aggressive buyers returning with volume <i>and</i> price rising this time

What the difference really is

The two traders might reach the **same conclusion**. The difference lies in the **quality of the explanation** and in how precisely each can say what would prove them wrong. The volume trader can state a sharper **invalidation condition**: "If aggressive buyers return in size and this time price moves up through the high, my 'sellers are absorbing' explanation is wrong." That is testable against data; "the pin bar failed" is vaguer.

One example proves nothing

Any book can show a chart where its method worked. That is **cherry-picking**, and one chosen example proves nothing about average behaviour. The message here is "more information and better questions," not "my method beats yours." Many experienced traders combine price levels with order flow rather than choosing one.

Common mistake: "Volume traders are always ahead of pattern traders." This is a claim without evidence.

Common mistake: "Long wick plus high volume means a certain reversal." This is a popular claim, not proven.

Try it: Take any gold candle with a long wick at a prior day's high or low. Before reading the outcome, write down the four volume trader questions and answer them from the data. Only then look at what happened next, and remind yourself that it is one example.

Order Flow Myths: The Holy Grail and "Big Volume = Big Buyers"

Order-flow education is often sold with big promises. This section is a vaccination against them. For each myth, ask three questions: **Where is the evidence? Does it hold after costs? On how large a sample?**

Myth 1: "Order flow is the holy grail"

No data set is a holy grail. Well-known educators in this field openly describe order flow as neutral information that cannot rescue a method that does not work in the first place. In our own testing, order-flow signals on their own did not predict direction after costs. Order flow is information, not a system.

Myth 2: "Big volume means big players are buying"

Volume is always **both** buying and selling. High volume means **a lot of trading happened**; it says nothing about direction. A very high-volume candle might be big participants exiting, a cluster of stops being triggered, or a battle between two sides.

Myth 3: "I can see the institutions"

Exchange data has no identities: you see price, size, time and (with good data) the aggressive side. Large orders are usually split by algorithms into small, often passive pieces, while one large print may be several small orders matching at once. "Who" in order flow means *size, side and timing*, never names.

Myth 4: "Delta divergence means a reversal is coming"

A popular claim, not proven. When price and cumulative delta disagree (Chapter 3), the disagreement can last hours or days while price keeps going; some platform educators warn it may reflect large passive flows rather than weakness.

Myth 5: "Absorption always leads to a bounce"

A popular claim, not proven. Heavy aggressive volume that fails to move price can still break through; an apparent wall may be a temporary iceberg order that simply finishes.

Myth 6: "Tick volume in forex is the same thing"

Related, but not the same; the next section explains why.

"A popular claim, not proven" does not mean false. It means the idea has not been shown to work in a fair test: on enough data, out of sample, after costs. Treat such ideas as hypotheses, not facts.

Key idea: Volume is direction-free. High volume tells you "a lot happened here," never "this way next."

Try it: Find three high-volume 5-minute candles on a gold chart from the past week. For each, note what happened in the following hour. You will very likely find that the same kind of candle was followed by continuation in one case and reversal in another.

Real Volume vs Tick Volume: Why Futures, Not Spot or CFD

Many traders first meet gold through a forex broker, trading XAUUSD on MetaTrader or a similar platform. Those charts also have a "Volume" panel. It is important to understand that this is usually not the same kind of volume as on a futures chart.

Real volume

Real volume is the actual number of contracts traded, as reported by the exchange. In GC, every trade is published with its price, its size (say, 3 contracts) and its time. CME's market-data feed even identifies which side was the aggressor (Chapter 3 explains how). Because all GC trades pass through one central order book on CME Globex, the exchange's count is the *complete* count for that contract.

Tick volume

Tick volume is the number of times the quoted price changed during a candle, regardless of how much was traded. In MetaTrader and most forex or CFD charts, this count is what is labelled "Volume".

A small example shows the difference. Suppose in one minute of real trading:

- a 1-contract trade prints at 4000.1;
- a 1-contract trade prints at 4000.2;
- a 500-contract trade prints at 4000.2.

Real volume for that minute is 502 contracts. A tick-volume counter that counts price changes sees only two updates (to 4000.1 and to 4000.2). The 500-contract trade, the most important event of the minute, did not change the price, so it may not register at all. In tick volume, a 1-lot trade and a 1,000-lot trade weigh the same, or nothing.

Why spot and CFD markets have no central volume

Spot gold and CFDs are **decentralised** (over-the-counter) markets. A **CFD** (contract for difference) is a private contract between you and your broker, not an exchange-traded product. Each broker and liquidity provider streams its own prices, and no single place records every trade in the market. The "volume" on your broker's XAUUSD chart reflects only the price updates in *that broker's* feed.

Is tick volume worthless?

Not completely. Studies have found a fairly strong correlation between tick volume and real volume on higher timeframes and during busy hours. It is a reasonable rough proxy for *how active* the market is. But the relationship is unstable, and it tends to break down exactly at important moments, such as when a few large trades go through with very few price changes.

More importantly, tick volume contains **no information about the aggressive side, trade sizes, or volume at each price**, so a genuine footprint, delta or cumulative delta cannot be built from it. Any "delta" on CFD data is an estimate from price movements.

Key idea: For footprint, delta and cumulative delta you need trade-by-trade data with sizes and an aggressive side. Only a centralised exchange market such as CME gold futures provides that.

What this book uses

The reference data throughout this book is **GC (and MGC) from CME**, even if a reader ultimately trades something else. If you trade XAUUSD, you can still read GC's order flow as a guide to what is happening in the central gold market, translating levels with the approximate futures-to-spot gap.

Gold note: When you transfer a GC level to XAUUSD, subtract the current gap (recently roughly \$27–30), and remember that the gap moves. A GC level of 4035.0 would sit at approximately 4,005–4,008 on XAUUSD. Check the live gap on the day rather than using a fixed number.

Real futures data has a cost: CME charges a monthly market-data fee (lower for non-professional users), and the quality of what you see depends on your data feed, which is the subject of the next section.

Common mistake: "The Volume in MT4/MT5 is real gold volume, and the delta on my XAUUSD chart is real." The first is almost always your broker's tick volume; the second is an estimate, not exchange aggressor data.

The Tool Map: Platforms and Data Feeds

To read order flow you need three separate things, which are often chosen independently:

1. A **platform**: the software that displays and analyses the data.
2. A **data feed**: the service that delivers trade-by-trade data and order-book depth from the exchange to your platform.
3. A **broker** (in futures, an **FCM**, futures commission merchant): the firm that holds your account and sends your orders to the exchange.

The data flows like this: CME Globex → data feed → platform → you. Orders flow the other way: you → platform → broker/FCM → CME Globex.

The main platforms

The descriptions below are neutral summaries, not recommendations. Features, prices and versions change often, so always check the current documentation.

- **NinjaTrader 8**. A general trading platform whose **Order Flow+** add-on provides **Volumetric Bars**, its name for footprint charts. Delta can be calculated as **BidAsk** (each trade compared with the bid and ask) or **UpDownTick** (each trade compared with the previous trade's price). BidAsk bars on historical data need **Tick Replay** and tick history that includes bid and ask.
- **ATAS**. Focused on volume analysis: many **cluster** (footprint) styles, cumulative delta with an optional session reset, and cluster statistics.
- **Bookmap**. Built around a **heatmap** of order-book history (time across, price up, colour = resting size), with trades drawn as bubbles. Suited to watching *intent* and changing depth (Chapter 10).
- **Sierra Chart**. Fast and highly configurable; its footprint is called **Numbers Bars**, and it stores full tick-by-tick history. Steeper learning curve.

Data feeds

For CME futures, the most common feeds among retail order-flow traders are **Rithmic** and **CQG**. Sierra Chart also offers its own feed, **Denali**. When comparing feeds, ask:

- Is it **unfiltered, tick-by-tick**, or are trades aggregated or sampled? Aggregation ruins footprint accuracy.
- How many **levels of depth** does it deliver?
- Does it provide **historical tick data with bid and ask**, so your platform can rebuild past footprints?
- What does the **CME market-data subscription** cost for your user category? This is a separate exchange fee.

Same data, different lenses

All of these platforms display **the same CME data**. None of them has access to secret information. The differences are in how they display it and how they calculate things from it.

This explains a confusing experience many beginners have: two platforms showing slightly different delta for the same gold candle. The usual reasons are a different classification method (BidAsk vs UpDownTick, for instance), different feed filtering, different handling of trades that are hard to classify, and different reset or session settings. Chapter 3 explains each of these in detail.

Key idea: Platform, data feed and broker are three different layers. The platform is a lens; the data underneath is the exchange's.

Common mistake: "Platform X shows hidden market data." All of them use public exchange data.

Common mistake: "A free, delayed feed is good enough for order flow." Delayed or aggregated data distorts footprint and delta.

Chapter summary

- A chart price is the **last trade**; the chain is decisions → orders → trades (volume) → price.
- A one-tick move can be a **push** (aggressive buyers) or a **pull** (sellers cancelling). The candle looks the same; volume and delta differ.
- Explaining is not predicting. In our own testing, order-flow signals alone did not predict direction after costs.
- MA, RSI and MACD are built mostly from closes; smoothing means **lag**, and they ignore volume, path, aggressive side and location.
- **Auction Market Theory**: price rises to find sellers and falls to find buyers; high volume suggests acceptance, thin fast trading suggests rejection. A framework, not a formula.
- The **order book** holds bids and asks; the **spread** in GC is usually one tick (\$0.10). Globex mostly uses price-time (FIFO) priority.
- The **aggressor** triggers a trade; the **passive** order was resting. Bought and sold contracts are always equal; "buying" means aggressive buying.
- Pattern and volume traders ask different questions; the volume trader can often state sharper invalidation. One example proves nothing.
- Popular myths (holy grail, big volume = big buyers, seeing institutions, divergence = reversal, absorption = bounce) are unproven or wrong.

- **Real volume** (exchange contracts) is not **tick volume** (price updates); genuine footprint and delta need exchange data such as CME GC/MGC.
- Order flow needs a **platform**, a **data feed** and a **broker**; all platforms show the same CME data through different lenses.

Checklist

- I can explain the chain decisions → orders → trades → price, and why a chart price is not "the price I can buy at."
- I can describe a "push" and a "pull" one-tick move and how volume tells them apart.
- I can explain why averaging creates lag and list four things MA, RSI and MACD ignore.
- I can explain value, balance, imbalance, acceptance and rejection in auction terms.
- I can read a DOM ladder: best bid, best ask, spread, depth, and the effect of a market order on it.
- I can define aggressor and passive and explain why bought and sold contracts are always equal.
- I know the difference between real volume and tick volume, and why genuine delta needs exchange data.
- I can name the three layers of an order-flow toolkit and explain why two platforms may show different delta.

Quiz

1. Gold ticks up from 4000.1 to 4000.2 while only a handful of contracts trade. What is the most likely explanation?
2. Gold has closed at 4000.0 for four candles and then closes at 4010.0. What is the value of a 5-period simple moving average after that fifth candle, and what does the result illustrate?
3. The GC ladder shows best bid 4000.0 and best ask 4000.1, with 12 contracts offered at 4000.1 and 18 at 4000.2. A trader sends a market order to buy 20 contracts. Where does the order fill, and what is the average price?
4. A resting sell limit order at 4000.1 is filled by an incoming market buy order. Which side is the aggressor, and how should a volume trader describe this trade?
5. A 5-minute bar on a forex broker's XAUUSD chart shows "Volume: 950". Why can you not build a genuine delta from this number?

Quiz answers

1. Most likely the sellers resting at 4000.1 **cancelled or moved their orders** (a "pull"), so the best ask moved up without much aggressive buying. If aggressive buyers had pushed the price up, you would expect to see substantial volume traded at 4000.1.

2. $(4000.0 \times 4 + 4010.0) \div 5 = \mathbf{4002.0}$. Price is already \$10 higher, but the average has moved only \$2. This illustrates **lag**: averaging the past always delays the response to a change.
3. **12 contracts fill at 4000.1** and **8 at 4000.2**. Average price = $(12 \times 4000.1 + 8 \times 4000.2) \div 20 = \mathbf{4000.14}$. The 0.04 above the expected 4000.1 is slippage, about \$80 in total on 20 GC contracts. The best ask becomes 4000.2.
4. The **buyer is the aggressor**, because the market buy order triggered the trade on arrival; the seller was passive. A volume trader would count this as **aggressive buying** (a buy at the ask), while remembering that the trade, like every trade, had exactly one buyer and one seller.
5. Because that "volume" is almost certainly **tick volume**: a count of price updates in one broker's feed. It does not contain trade sizes, does not cover the whole market, and has no information about which side was aggressive. Genuine delta requires exchange trade-by-trade data with sizes and an aggressive side, such as CME gold futures data.

PART II • THE TOOLS OF VOLUME

3

Volume and Delta

In this chapter: You will learn how the raw stream of gold trades is turned into the numbers that order-flow traders read: how each trade gets labelled a "buy" or a "sell" (and where that label can be wrong), how to judge a bar's volume against the normal activity for its time of day, how a footprint bar opens a candle up price by price, how bar delta and delta percentage summarise aggression, and how cumulative volume delta (CVD) adds it all up. You will also meet the CVD "anchor trap", see why high volume never tells you direction by itself, and get a first, careful look at what it means when delta and price disagree. Everything here is a *measurement*. None of it is a signal.

Who Hit Whom? Classifying Trades as Buys or Sells

Chapter 2 established that every trade has one buyer and one seller, so "buy" and "sell" in order flow can only mean one thing: **which side was the aggressor**. A trade is called a "buy" when an impatient buyer crossed the spread and took a waiting seller's offer. It is called a "sell" when an impatient seller crossed the spread and hit a waiting buyer's bid. Every number in this chapter (footprint cells, delta, CVD) rests on getting this label right.

There are two ways to find out who was aggressive.

Method 1: the exchange tells you

CME's market-data feed, called **MDP 3.0**, publishes a **Trade Summary** message for every trade. One of its fields is **AggressorSide** (FIX tag 5797):

AggressorSide value	Meaning
1	The aggressor was the buyer
2	The aggressor was the seller
0	No aggressor

CME defines the aggressor as the order that, on entering the book, immediately triggers a trade. That is exactly the definition we used in Chapter 2.

Some trades legitimately have **no aggressor**. The most important example is the **opening auction**: when the market reopens after the daily halt or after a pause, orders accumulated during the pre-open are matched all at once at a single price. Nobody "crossed the spread", so no side is aggressive. Some **implied** trades, which arise when orders in calendar-spread markets are combined with orders in individual contract months, can also carry no aggressor flag.

Method 2: compare the trade with the quote

Many platforms and feeds do not pass the exchange flag through to the chart. Instead, they compare each trade with the **bid and ask at the moment of the trade**, a method called the **quote rule** (or the at-ask/at-bid rule):

- trade price **at or above the ask** → **aggressive buy** (the buyer lifted the offer);
- trade price **at or below the bid** → **aggressive sell** (the seller hit the bid).

In NinjaTrader this is the **BidAsk** method of calculating delta.

A worked example: ten gold prints

Here are ten consecutive GC trades from a busy morning minute, laid out the way you would see them on Time & Sales, with the bid and ask that were in force at the time. (GC prices around 4000 correspond to roughly 3,970–3,973 on XAUUSD; the gap is approximate and changes over time.)

#	Time (ET)	Price	Size	Bid	Ask	Label (quote rule)
1	09:31:02.114	4000.1	3	4000.0	4000.1	Buy
2	09:31:02.120	4000.1	12	4000.0	4000.1	Buy
3	09:31:02.388	4000.0	5	4000.0	4000.1	Sell
4	09:31:02.391	4000.0	1	4000.0	4000.1	Sell
5	09:31:02.905	4000.1	22	4000.0	4000.1	Buy
6	09:31:03.010	4000.2	8	4000.1	4000.2	Buy
7	09:31:03.011	4000.2	30	4000.1	4000.2	Buy
8	09:31:03.350	4000.1	4	4000.1	4000.2	Sell
9	09:31:03.612	4000.2	6	4000.1	4000.3	?
10	09:31:03.700	4000.3	10	4000.2	4000.3	Buy

Adding up: aggressive buys = 3 + 12 + 22 + 8 + 30 + 10 = **85** contracts; aggressive sells = 5 + 1 + 4 = **10** contracts; unclassified = **6** contracts (print 9). Total volume = 101. Net aggression (buys minus sells) = **+75**, with 6 contracts we could not label.

Print 9 is the interesting one. For a moment the spread had widened to two ticks (bid 4000.1, ask 4000.3), and the trade printed at 4000.2, in the middle. It is neither at the bid nor at the ask, so the quote rule cannot decide.

Where the quote rule goes wrong

The quote rule is simple and usually right, but it has three weak spots:

1. **Synchronisation.** Quotes and trades arrive as separate messages, milliseconds apart. If the platform pairs a trade with a quote that is slightly too early or too late, a trade at the old ask may be compared with a new bid and get the wrong label. This happens most in fast markets: data releases, the open, sudden moves.
2. **Trades inside the spread.** When the spread is wider than one tick, or changes at the instant of the trade, a print can land between bid and ask, like print 9.
3. **Sweeps.** A large aggressive order that eats through several price levels is reported as several prints at different prices. Each print is labelled separately, and the later ones can be compared against quotes that have already moved.

The practical conclusion: the delta you see on screen is a **measurement with a small error**, not absolute truth. The error is usually small, but it exists, and it is largest exactly when markets are fastest.

Key idea: "Buy" and "sell" in order flow mean *aggressive* buy and *aggressive* sell. The label comes either from the exchange's aggressor flag or from comparing the trade with the bid and ask, and both routes have edge cases.

Common mistake: "A trade at the bid means someone bought." Every trade has a buyer, but a trade at the bid means an aggressive **seller** hit the bid. The buyer was passive.

Common mistake: "My platform's delta is 100% accurate." Accuracy depends on the feed and the classification method. And some trades, such as those in the opening auction, have no aggressor at all.

When Quotes Fail: The Tick Rule and Lee-Ready

What should happen to a trade like print 9? Academic market microstructure, the study of how trading works at the level of individual orders, has worked on this question for decades, because many historical data sets have no aggressor flag at all.

The tick rule

The **tick rule** (or tick test) ignores quotes and compares each trade with the **previous trade's price**:

- higher than the previous trade (an **uptick**) → buy;
- lower (a **downtick**) → sell;
- the same price (a **zero tick**) → inherit the direction of the **last price change**. A zero tick after an uptick (a *zero-uptick*) is a buy; a zero tick after a downtick (a *zero-downtick*) is a sell.

Example sequence of GC trades:

Trade price	Change	Tick-rule label
4000.1	(start)	—
4000.2	uptick	Buy
4000.2	zero-uptick	Buy
4000.1	downtick	Sell
4000.1	zero-downtick	Sell
4000.2	uptick	Buy

The tick rule needs nothing but trade prices, which makes it easy to apply to any data. NinjaTrader's **UpDownTick** delta method follows this logic.

Its weakness shows up in quiet, ranging markets. Suppose gold trades back and forth between 4000.0 and 4000.1 for several minutes, and the last price change was an uptick. Dozens of trades at 4000.1 will all be labelled buys simply because of one old uptick, even if many of them were sellers hitting a bid that had moved up to 4000.1. The tick rule **carries errors forward**.

Lee-Ready

In 1991, Charles Lee and Mark Ready proposed combining the two approaches. The **Lee-Ready algorithm**:

1. Compare the trade with the **midpoint** of the bid and ask (the quote rule in its midpoint form):
above the midpoint → buy; below → sell.
2. If the trade is **exactly at the midpoint**, use the tick rule.

Apply it to print 9 from the previous section. The bid was 4000.1 and the ask 4000.3, so the midpoint was 4000.2, exactly where the trade printed. Step 2 applies: the previous trade (print 8) was at 4000.1, so 4000.2 is an uptick, and Lee-Ready labels print 9 a **buy**.

How accurate is it? Studies on stock markets, where researchers could check against true order data, have reported accuracy of roughly 80–90%. That figure varies with the market, the period and the data quality, and fast electronic markets tend to be harder. The lesson is not the exact number. It is that **even the best inference methods mislabel a meaningful share of trades**.

Same bar, different delta sign

The choice of method is not a technicality. It can flip the sign of a bar's delta. Here is a small constructed example of five GC trades in a choppy moment:

Trade	Price	Size	Bid / Ask at the time	Quote rule	Tick rule
A	4000.0	20	4000.0 / 4000.1	Sell	Sell (downtick from 4000.1)
B	4000.1	50	4000.1 / 4000.2	Sell (at bid)	Buy (uptick)
C	4000.1	30	4000.1 / 4000.2	Sell (at bid)	Buy (zero-uptick)
D	4000.1	15	4000.0 / 4000.1	Buy (at ask)	Buy (zero-uptick)
E	4000.0	10	4000.0 / 4000.1	Sell	Sell (downtick)

- **Quote rule (BidAsk):** buys 15, sells 20 + 50 + 30 + 10 = 110 → delta **-95**.
- **Tick rule (UpDownTick):** buys 50 + 30 + 15 = 95, sells 20 + 10 = 30 → delta **+65**.

Same five trades, opposite conclusions. The disagreement comes from trades B and C: the quotes had moved up, so sellers were hitting a higher bid, but the tick rule only saw that the price was higher than before.

An honest "unknown" bucket

There is a third option that deserves more use: **admitting you do not know**. The approach we prefer works like this:

1. Use the quote rule first.
2. If the quote cannot decide, use the tick rule **only** when the previous trade was very recent (a fraction of a second earlier) and at a **different** price, so the tick direction reflects fresh information.
3. Otherwise label the trade **UNKNOWN** and leave it there. It is never moved into buys or sells later.

The result is a delta that is slightly "smaller", because some volume is set aside, but **less contaminated** by guesses. It also lets you report how much of the volume was labelled by each method, which tells you how much to trust the number on a given day.

No method is perfect. What matters most is **consistency**, so that comparisons from day to day are meaningful, and **transparency**, so that you and anyone reading your charts know what the numbers mean.

Key idea: Trade classification is inference, not observation, unless you have the exchange's aggressor flag. Prefer a consistent method, know which one your platform uses, and treat unknowns honestly.

Common mistake: "UpDownTick and BidAsk are the same thing." They are different methods and can give different results, sometimes with opposite signs.

Common mistake: "Lee-Ready is 85% accurate everywhere." Accuracy depends on the market, the period and the data quality.

Try it: If your platform offers both BidAsk and UpDownTick delta, display both on the same 1-minute GC chart during a quiet Asian-session hour. Count how many bars show deltas with opposite signs.

Volume per Bar and Relative Volume

Before splitting volume into buys and sells, it helps to read the simplest number of all: how much traded in a bar.

Bar volume

Bar volume is the total size of all trades during a bar's time window, measured in contracts for GC. Platforms usually draw it as a histogram under the price chart.

One detail confuses many beginners: the histogram bars are usually **coloured by the candle's direction**. A green volume bar means the candle closed higher than it opened. It does **not** mean buyers were more aggressive. A green candle can have more aggressive selling than buying, as you will see when we reach delta.

Gold's daily rhythm

Gold trades nearly 23 hours a day, from Sunday 18:00 to Friday 17:00 New York time, with a daily halt from 17:00 to 18:00. Volume is very uneven across that clock. Asian hours are usually quiet, activity rises with London from around 03:00 ET, and the busiest stretch is normally the London–New York overlap, especially around US data releases (often at 08:30 ET) and the first hours of the New York session.

This means a raw number like "800 contracts in five minutes" has no fixed meaning. At 02:00 ET it may be exceptional; at 09:35 ET it may be quiet.

Relative volume (RVOL)

The fix is to compare each bar with what is normal **for that time of day**. **Relative volume (RVOL)** is:

$$\text{RVOL} = \text{bar volume} \div \text{average volume of bars at the same clock time over the previous } N \text{ days}$$

Here is an illustrative example with round numbers (not from a specific day). Suppose that over the previous 20 trading days:

- the average 5-minute GC bar at 02:00–02:05 ET traded 250 contracts;
- the average 5-minute bar at 09:35–09:40 ET traded 2,400 contracts.

Today, both bars trade 800 contracts.

Bar	Volume	Same-time average	RVOL
02:00–02:05 ET	800	250	$800 \div 250 = 3.2$
09:35–09:40 ET	800	2,400	$800 \div 2,400 \approx 0.33$

The same raw volume is three times normal in the first case and one third of normal in the second. RVOL = 2 simply means "twice the usual activity for this hour."

Simpler alternatives exist: comparing a bar with a moving average of the last N bars' volume, or ranking a bar by percentile within the day. These are easy to compute, but they do not fully remove the time-of-day effect. A moving average of the last 20 five-minute bars at 08:35 ET is still mostly made of quieter pre-data bars.

Volume filters: where to look, not what to do

A **volume filter** marks bars whose relative volume exceeds a chosen threshold, for example "more than twice normal for this hour". Treat such marks as an **attention highlight**: "something unusual happened here in terms of activity." They are not buy or sell signals. As you will see later in this chapter, high volume has no direction.

What distorts the averages

Several calendar effects can make "normal" misleading:

- **Roll periods**, when traders move from one contract month to the next. Volume collapses in the old contract and surges in the new one (Chapter 1).
- **Holidays and shortened sessions**, which shrink volume.
- **Scheduled data releases**, which make the same minute look huge every month.
- **Daylight-saving changes**. The US and Europe change clocks on different dates, so for a few weeks a year the London–New York overlap shifts by an hour.

Key idea: Volume only means something relative to its hour. Compare like with like.

Common mistake: "A green volume bar means buyers were stronger." It only means the candle closed above its open.

Common mistake: "A fixed threshold of 1,000 contracts works for every hour." Asia and New York are not comparable.

Common mistake: "An exceptionally high-volume bar is always a turning point." This is a popular claim, not proven.

Inside the Candle: Volumetric and Footprint Bars

A candle compresses everything that happened in its time window into four prices. A **footprint** reverses that compression. Different platforms use different names: NinjaTrader calls it **Volumetric Bars**, ATAS calls it **Clusters**, Sierra Chart calls it **Numbers Bars**. The idea is the same: for **every price level** inside the candle, show how much traded and which side was aggressive.

Chapter 8 is devoted to footprint charts. Here we only need to learn to read one.

Reading bid × ask

The most common layout is **bid × ask**. At each price inside the bar you see two numbers:

- the **left** number is the volume that traded **at the bid**, which means **aggressive selling**;
- the **right** number is the volume that traded **at the ask**, which means **aggressive buying**.

For example, a level reading **35 × 120** at 4001.2 means that at that price, 35 contracts were sold aggressively and 120 were bought aggressively. The total volume at the level is $35 + 120 = 155$.

Here is a complete small footprint for one illustrative 5-minute GC candle that opened at 4000.6, traded between 4000.4 and 4001.2, and closed at 4001.0:

Price	Bid (aggr. sells)	Ask (aggr. buys)	Level volume
4001.2	35	120	155
4001.1	88	142	230
4001.0	210	265	475
4000.9	160	190	350
4000.8	130	115	245
4000.7	95	70	165
4000.6	60	48	108
4000.5	72	30	102
4000.4	40	20	60
Total	890	1,000	1,890

What can we read from it?

- **Total volume** for the bar is 1,890 contracts.
- The level with the most volume is 4001.0, with 475 contracts. That level is the bar's **point of control (POC)**: the price where the most trading took place inside the candle. Here it sits near the top, not in the middle.
- The low of the bar (4000.4) traded only 60 contracts: price touched it briefly with little business done.
- Overall, 1,000 contracts were bought aggressively and 890 sold aggressively. That difference is the bar's **delta**, the subject of the next section.

The candle chart would have shown only: open 4000.6, high 4001.2, low 4000.4, close 4001.0. The footprint shows where the business was done and who initiated it.

Other ways to display the same data

Platforms offer several views of the same per-level data:

- a **volume profile inside the bar**: a small horizontal histogram of total volume at each price;
- **delta per level**: ask minus bid at each price;
- **bar statistics** under each candle: total volume, delta and more.

In the bid × ask layout, experienced readers also compare numbers **diagonally** (the bid at one price with the ask one tick higher) to look for imbalances. That technique belongs in Chapter 8; for now, the basic reading is enough.

Settings that change the picture

Footprints are sensitive to settings, so two charts of the same gold candle can look different:

1. **Ticks per level.** Each row can be one tick (\$0.10) or several ticks grouped together. Larger groupings are easier to read but hide detail. A busy gold candle with a \$5 range has 50 one-tick levels; grouped by 5 ticks, it has 10.
2. **Bar type.** Time bars (for example, 5 minutes), tick bars, volume bars and range bars all slice the same trades differently.
3. **Classification method.** BidAsk, UpDownTick or an approach with an unknown bucket, as described above.
4. **Historical data.** Building past footprints accurately requires historical tick data with the bid and ask at each trade. Without it, history is either unavailable or rebuilt with a different method than live data.

What a footprint cannot show

A footprint shows **what happened**, not what will happen. It also shows only **trades**. Orders that were placed and cancelled, which can matter a great deal, appear nowhere in it (Chapter 10 deals with the order book).

Key idea: In a bid × ask footprint, left = aggressive sells, right = aggressive buys, and their sum is the volume at that price.

Common mistake: "The left number is buyers who bought at the bid." The buyers at the bid were *passive*. The footprint counts the *aggressive* side, which at the bid is the seller.

Common mistake: "The bar's POC always becomes support or resistance." This is a popular claim, not proven.

Bar Delta: Net Aggression in One Number

A footprint holds a lot of numbers. **Bar delta** summarises them in one.

The formula

Bar delta = aggressive buy volume – aggressive sell volume

In footprint terms: the sum of the ask column minus the sum of the bid column. For the candle above, delta = 1,000 – 890 = **+110**.

A larger worked example: a 5-minute GC bar with 1,200 contracts traded at the ask and 900 at the bid has:

- $\text{delta} = 1,200 - 900 = \mathbf{+300}$;
- $\text{total volume} = 2,100$ (assuming no unclassified trades).

Delta percentage

A raw delta of +300 means different things in different bars. **Delta %** normalises it:

Delta % = delta ÷ total volume

For our bar: $300 \div 2,100 \approx \mathbf{14\%}$. Compare:

Bar	Volume	Delta	Delta %
Quiet Asian bar	500	+300	+60%
Our New York bar	2,100	+300	+14%
Data-release bar	10,000	+300	+3%

The same +300 is overwhelming one-sidedness in the first bar and almost perfect balance in the third. Delta % is what lets you compare bars of very different sizes fairly.




Figure 8. Real 1-minute gold futures candles (Oct 2, around the jobs report). Delta numbers show net aggressive buying (+, above the high) or selling (–, below the low) where it reached the minimum.

Max delta and min delta

Delta changes continuously as a bar forms. Many platforms show the **max delta** (the highest value delta reached during the bar) and the **min delta** (the lowest). A bar might close with delta +300 after spending part of its life at -400. Those two numbers sketch the path of aggression inside the bar, something the closing delta alone hides.

What positive delta tells you, and what it does not

A positive delta tells you exactly one thing: **during this bar, buyers were more in a hurry than sellers**. More contracts were bought by crossing the spread than were sold by crossing it.

It does **not** tell you:

1. **That price went up.** A bar can close lower with positive delta. Picture aggressive buyers repeatedly lifting offers while a large seller keeps refilling the offer side with passive limit orders and the bid side quietly steps lower. Plenty of aggressive buying, yet a red candle.
2. **That "smart money" was buying.** Identity is invisible, and many large participants trade passively.
3. **That the next bar will be higher.** Delta describes the bar that just finished.
4. **Anything independent of classification.** As shown above, the method can change the number, even its sign.

There is a simple way to keep this in perspective. Every bar is an interaction between **aggressors**, which delta measures, and **passive orders**, which delta does not see at all. Delta is half the story.

In our own testing, the sign of a bar's delta on its own did not predict the direction of the following bars once costs were included. Think of delta as **a measurement of aggression**, not a signal.

Key idea: Delta measures who was in a hurry, not who was "stronger", and not where price goes next.

Common mistake: "Positive delta means price will go up." Not necessarily, in this bar or the next.

Common mistake: "Delta is the number of buyers minus the number of sellers." Delta counts contracts, not people, and bought and sold contracts are always equal; delta measures only the aggressive side.

Common mistake: "A red candle with large positive delta is certain absorption and a reversal." This is a popular claim, not proven.

Try it: Take ten consecutive 5-minute GC bars and write each bar's delta and delta % underneath. Mark any green bar with negative delta and any red bar with positive delta. They are more common than most people expect.

Cumulative Volume Delta (CVD): The Running Total

Bar delta describes one bar. **Cumulative volume delta (CVD)** strings bars together.

The formula

CVD is the running sum of bar deltas from a chosen **starting point**, called the **anchor**:

$$\text{CVD}(n) = \Delta(1) + \Delta(2) + \dots + \Delta(n)$$

Worked example: four bars with deltas +300, -100, +250, -50.

Bar	Bar delta	CVD
1	+300	300
2	-100	200
3	+250	450
4	-50	400

CVD is usually drawn as a line, or as "CVD candles", in a panel under the price chart.

Reading CVD: the slope, not the number

- **Rising slope:** over this stretch, aggressive buying has outweighed aggressive selling.
- **Falling slope:** aggressive selling has dominated.
- **Flat:** aggression has been roughly balanced.
- **Compared with price:** if price and CVD move together, the move is accompanied by aggression from the same side. If they do not, there is a disagreement (the last section of this chapter).

The **absolute level** of CVD, such as "+12,400", means almost nothing on its own, because it depends entirely on where the sum started. What carries meaning is the **change** over a specific stretch of time. That is the gateway to the next section, on anchors.

CVD inherits, and accumulates, delta's limits

- **Errors add up.** Small classification errors, harmless in a single bar, accumulate over hours. A CVD line can slowly drift for reasons that have nothing to do with the market.

- **Opening or closing?** CVD does not know whether aggressive buying came from new buyers or from sellers closing short positions (**short covering**). Both look identical.
- **Passive flow is invisible.** A large participant buying patiently with limit orders all day can build a big position while CVD slopes *down*, because every one of their fills counts as an aggressive sell by someone else.

The realistic use of CVD is to **summarise the flow of aggression across a session at a glance** and compare it with the path of price. It is context, not an entry signal.

Key idea: Read the slope of CVD over a defined window. The absolute number depends on where you started counting.

Common mistake: "Positive CVD means more money came into gold today." It means more net *aggressive* buying, nothing more.

Common mistake: "CVD always turns before price." This is a popular claim, not proven.

The CVD Anchor Trap: Session Reset vs Chart Load

Because CVD is a running total, the anchor controls what you see. This is one of the most common sources of confusion in order flow, and it is easy to avoid once you understand it.

Trap 1: "since chart load"

Many platforms, by default, start the sum at the **first bar loaded on the chart**. Load five days of history today and ten days tomorrow, and the CVD line starts from a different place.

Suppose that over the five days before today, net delta was +4,800, while over the ten days before today it was -2,300. Today's own net aggression is +1,150 either way. But:

Chart setting	CVD at today's start	CVD at today's end
Since load, 5 days	+4,800	+5,950
Since load, 10 days	-2,300	-1,150
Session reset	0	+1,150

Mathematically, loading more history shifts the line by a constant on the days the two charts share. In practice, though, the panel rescales, the earliest loaded bars differ, data gaps fall in different places, and any comparison involving the start of the data changes. Screenshots stop being

reproducible and level comparisons become meaningless. ATAS, for example, states plainly that without its session delta mode, the indicator uses all the loaded data.

Session reset

The alternative is to **reset CVD to zero at the start of each session**. Days become comparable: each day's CVD answers the same question. But that raises the next question: **which session?** For gold the reasonable choices include:

Anchor	Starts at	Question it answers
CME trading day	18:00 ET the previous evening	Net aggression since the official trading day began
Asia	18:00 ET	Net aggression during the Asian session
London	about 03:00 ET	Net aggression since London became active
New York	about 08:20 ET	Net aggression since the New York open (many educators prefer this for intraday work)
New York initial balance	08:20 ET, first hour	Net aggression during New York's opening hour

Remember from Chapter 1 that the CME trading day starts at 18:00 ET the evening before, so a Monday-evening Asian session belongs to Tuesday's trading day.

Why the anchor can change a "divergence"

Here is a concrete case. During one illustrative day, GC makes a high at 4020.0 at 04:10 ET during London, and a slightly higher high at 4024.0 at 10:30 ET during New York (about 3,990–3,993 and 3,994–3,997 on XAUUSD).

- **Anchored at the CME day (18:00 ET):** CVD was +1,500 at the London high and +1,300 at the New York high. Price made a higher high, CVD a lower high: it looks like a "divergence".
- **Anchored at the New York open (08:20 ET):** the London high happened before the anchor, so there is nothing to compare it with. From 08:20 to 10:30, CVD rose by about +900 alongside price.

Both statements are true. They answer different questions: "Since the trading day began, was aggression at the second high weaker than at the first?" versus "Since New York opened, has aggressive buying accompanied the rally?" A comparison between two points *after* the same anchor gives the same answer whatever the anchor is. A comparison that straddles an anchor, or that compares levels across days, depends entirely on the anchor.

Trap 2: the incomplete anchor

Suppose your platform connects at 09:10 ET because your computer was off, and your "session CVD" is supposed to start at 18:00 ET the previous evening. Unless the platform backfills complete tick history, your CVD actually starts at 09:10: it is missing all of Asia and London. It looks exactly like a complete line. A good tool should flag an anchor as **incomplete** when data before it is missing or when there was a gap or disconnection during the session.

Trap 3: the wrong scale

A weekly or multi-day CVD buries short intraday swings under days of accumulated flow. The rule of thumb: **the scale of the anchor should match the scale of the question**. Intraday questions need intraday anchors.

Label your anchor

Whenever you show or share a CVD chart, write the anchor on it, for example "CVD anchored: NY open 08:20 ET". Without the anchor, a CVD claim cannot be checked.

Key idea: A CVD line without a stated anchor is incomplete information. Choose the anchor that matches your question, and say what it is.

Common mistake: "My CVD shows a divergence," without saying the anchor. With a different anchor the divergence may not exist.

Common mistake: "CVD is the same on every platform." Anchors and classification methods differ.

Common mistake: "Weekly CVD gives a bigger, better picture." For intraday questions it mostly adds noise.

Try it: On a 5-minute GC chart, display CVD three ways: since chart load, reset at 18:00 ET, and anchored at 08:20 ET. Pick two price highs on the same day, one before 08:20 and one after, and compare what each version says about them.

High Volume ≠ Direction

This idea appeared in Chapter 2 as a myth to discard. Now that you have the tools of this chapter, we can make it practical.

High volume means **a large number of contracts changed hands**. Every one of those contracts had a buyer and a seller, so volume is, by its nature, **direction-free**. The same tall volume bar can come from completely different situations.

Five stories behind one tall bar

1. **Initiative.** One side arrives with heavy aggression and moves price. Typical footprint: high volume, delta strongly in the direction of the move, a large range.
2. **Stop run or liquidation.** Many stop orders sit beyond an obvious level, such as the Asian-session high. When price reaches it, the stops trigger and turn into market orders. The volume is high because traders were *forced* to act, not because they chose to. Afterwards, the move may continue or may simply end.
3. **Absorption, or a two-sided battle.** Heavy aggression runs into large passive orders on the other side. Volume is high but **price barely moves**.
4. **Scheduled events.** A US data release at 08:30 ET, the GC settlement window from 13:29 to 13:30 ET, or a roll period. The volume is high because of the **calendar**.
5. **Spread and routine trading.** Part of the volume can come from rolls between contract months or spread strategies that have no net directional intent.

Here are four illustrative 5-minute GC bars, each with RVOL above 3 (round numbers, not from a specific day):

Story	Volume	Delta	Delta %	Range
Data release at 08:30 ET	6,200	+1,400	+23%	\$9.0
Break of the Asian high (stop run)	3,100	+1,100	+35%	\$4.5
Heavy volume, tiny range (possible absorption)	4,000	+900	+22%	\$0.8
Bar containing the settlement window	3,500	+60	+2%	\$1.2

All four would light up a volume filter. They are four different events.

What to read next to volume

To interpret a high-volume bar, look at least at:

- **delta and delta %:** was the activity one-sided or balanced?
- **price movement within the bar:** high volume with a large range is different from high volume with almost no range;
- **location:** did it happen at a meaningful level (yesterday's high, a prior volume cluster) or in the middle of nowhere?

- **time:** was there a data release, the settlement window, a session open?
- **what happened afterwards:** did the market accept the new prices (trade there, build volume) or reject them?

A high-volume bar says, honestly, "something mattered to participants here." Working out *what* requires the other evidence, and even then the conclusion is a probability, not a certainty.

Key idea: A tall volume bar means "pay attention here," never "go this way."

Common mistake: "Big volume means big players are buying." Volume has no direction.

Common mistake: "A climactic high-volume bar marks the end of a trend." This is a popular claim, not proven.

Common mistake: "Low volume means the move is fake." This is a popular claim, not proven. When the order book is thin, moves on low volume are entirely real.

When Delta and Price Disagree (Introduction)

We usually expect aggressive buying to lift price and aggressive selling to push it down. When that does not happen, we have a **disagreement**, often called a **divergence**. This section only introduces the idea; later chapters return to it with more tools.

Two levels of disagreement

1. Within a single bar. A green candle with negative delta, or a red candle with positive delta. For example, GC rises during a bar while aggressive selling dominates. Price went up without buyers pushing it: perhaps passive sellers stepped away, or passive buyers kept raising their bids.

2. Across a stretch of CVD. Price makes a higher high, but CVD, with a stated anchor, makes a lower high. Or price makes a lower low while CVD makes a higher low.




Figure 9. Monday Sep 28, 2026: gold fell about \$82 while cumulative delta finished around +5,504. Aggressive buying was absorbed by resting sellers.

The figure shows a real gold session (28 September) in which price fell while cumulative delta rose: aggressive buying was dominating, yet price drifted lower.

Possible explanations

Each of these is plausible, and from the delta numbers alone you usually cannot tell which one applies:

- **Absorption.** One side's aggression is being soaked up by large passive orders on the other side.
- **Passive accumulation or distribution.** A large participant is building or reducing a position with limit orders, perhaps through an algorithm that spreads the order over time (for example a TWAP, time-weighted average price, algorithm). Some platform educators point out that a session-long divergence can reflect institutional passive flow rather than weakness.
- **Liquidity withdrawal.** The other side cancels its orders, and price moves through empty space.
- **A measurement artifact.** A classification error (as in the "same bar, different delta sign" example) or an unsuitable anchor.
- **Short covering or position closing.** Aggressive buying from traders closing shorts looks the same as new buying.

A question, not a signal

Why only an introduction? Because divergence is widely sold online as a "reversal signal". **That is a popular claim, not proven.** A divergence can persist for hours while price continues in its own direction.

The useful way to treat a disagreement is as a **question generator**, not a decision maker. Ask: "Why isn't this side's aggression moving price? Who is passive on the other side?" Then look for other evidence: the order book or heatmap (Chapter 10), the location within the volume profile (Chapters 5 and 6), and how the market behaves next.

A helpful rule from ATAS's educational material: **the scale of your reaction should match the scale of the disagreement.** A disagreement in one bar does not justify a conclusion about the whole day.

Key idea: When delta and price disagree, the honest statement is "aggression on this side is not what is driving price right now; find out why."

Common mistake: "Delta divergence means a reversal." This is a popular claim, not proven.

Common mistake: "A CVD divergence always means big players are on the other side." This is a popular claim, not proven.

Common mistake: "Several divergences in a row make a reversal certain." There is no statistical evidence for this.

Chapter summary

- In order flow, "buy" and "sell" mean the **aggressive** side. CME's feed flags it directly (AggressorSide, tag 5797: 1 buy, 2 sell, 0 none); otherwise platforms use the **quote rule** (at or above ask = buy, at or below bid = sell).
- Opening-auction trades and some implied trades have **no aggressor**. Quote/trade synchronisation, trades inside the spread and sweeps cause classification errors.
- The **tick rule** compares each trade with the previous trade; **Lee-Ready** uses the midpoint first and the tick rule for midpoint trades. Studies report roughly 80–90% accuracy in equities, never 100%. Methods can flip a bar's delta sign. An honest **unknown** bucket is better than guessing.
- **Bar volume** must be judged against the same time of day: $\text{RVOL} = \text{bar volume} \div \text{same-time average}$. Volume histogram colours follow the candle, not aggression. Volume filters highlight; they do not signal.

- A **footprint** shows bid × ask volume at every price in the bar: left = aggressive sells, right = aggressive buys. The bar's **POC** is its highest-volume level. Ticks per level, bar type, classification and historical data change the picture.
- **Bar delta** = ask volume – bid volume; **delta %** = delta ÷ volume. Positive delta means buyers were in more of a hurry, not that price rose or will rise. Delta does not see passive orders.
- **CVD** is the running sum of bar deltas from an **anchor**. Read its slope over a defined window, not its absolute level. Errors accumulate; short covering and passive flow are invisible.
- The **anchor trap**: "since chart load" CVD depends on how much data you loaded; session resets make days comparable; the choice of session answers a specific question. Flag incomplete anchors, match anchor scale to question scale, and always label the anchor.
- **High volume has no direction**. Initiative, stop runs, absorption, scheduled events and spread trading can all produce the same tall bar. Read delta, range, location, time and what followed.
- When **delta and price disagree**, possible causes include absorption, passive accumulation, liquidity withdrawal, measurement artifacts and short covering. Divergence is a question, not a reversal signal.
- In our own testing, order-flow measurements such as delta on their own did not predict direction after costs. Everything in this chapter is a measurement.

Checklist

- I can label a trade as an aggressive buy or sell from its price and the bid/ask at the time.
- I know what the CME aggressor flag is and which trades have no aggressor.
- I can apply the tick rule and Lee-Ready to an ambiguous trade, and I know which method my platform uses.
- I can compute RVOL against the same time of day and explain why a fixed volume threshold misleads in gold.
- I can read a bid × ask footprint, find the bar's total volume and POC, and say which column is aggressive buying.
- I can compute bar delta and delta % and explain what positive delta does and does not mean.
- I can build a CVD from bar deltas and read its slope rather than its level.
- I can choose a CVD anchor that fits my question and I always label it.
- I can list at least four different stories behind a high-volume bar.
- I treat a delta/price disagreement as a question to investigate, not a signal.

Quiz

1. The bid is 4000.0 and the ask is 4000.1. A 12-lot trade prints at 4000.0. How is it classified under the quote rule, and who was passive?
2. Quotes are unavailable. The previous trade printed at 4000.2 and this trade prints at 4000.3. What does the tick rule say? What would it say if the following trade also printed at 4000.3?

3. A 5-minute bar at 02:15 ET trades 600 contracts. Over the previous 20 days, the average 02:15 bar traded 200 contracts. What is the RVOL, and what does it tell you about direction?
4. A footprint bar's ask column sums to 600 and its bid column to 900. What are the bar delta and delta %? Could this bar have closed higher than it opened?
5. Bar deltas, anchored at the first bar, are +200, -50 and +100. What is the CVD after bar 3? If you reloaded the chart with more history and no session reset, what would change?

Quiz answers

1. It is an **aggressive sell**: a seller crossed the spread and hit the bid at 4000.0. The **buyer was passive**; their limit order was resting at the bid.
2. 4000.3 is higher than 4000.2, an uptick, so the tick rule says **buy**. A following trade at 4000.3 is a zero-uptick, so it would also be labelled a **buy**, inheriting the last change. (In a quiet range, this inheritance is exactly how the tick rule carries errors forward.)
3. $RVOL = 600 \div 200 = 3$: three times the normal activity for that time of day. It tells you **nothing about direction**; it only says something unusual happened in terms of activity. You would need delta, range, location and what followed to interpret it.
4. $\Delta = 600 - 900 = -300$. Volume = 1,500 (assuming no unknowns), so $\Delta \% = -300 \div 1,500 = -20\%$. **Yes**, the bar could still have closed higher: delta measures only aggressive orders, and price can rise if passive sellers withdraw or passive buyers keep raising their bids.
5. $CVD = 200 - 50 + 100 = 250$. Reloading with more history and no reset would shift the **level** of the whole CVD line (it would start from the sum of all the earlier loaded bars), and the chart could look different because of rescaling and data differences. The bar deltas themselves, and the change from bar 1 to bar 3, would not change.

4

VWAP

In this chapter: You will learn what the Volume-Weighted Average Price (VWAP) actually measures, why it began life as a yardstick for how well big institutions execute their orders, and how to calculate it by hand for a few gold candles. We will work through the standard deviation bands step by step, choose a session reset for gold deliberately rather than by accident, stack daily, weekly and monthly VWAPs, anchor VWAP to events, and read price location and VWAP slope as context. We finish by separating VWAP from moving averages and listing the seven most common VWAP mistakes. Throughout, the message is the same: VWAP is a reference for context, not a signal that tells you where price goes next.

In Chapter 3 you learned to count volume and delta: how many contracts traded, and which side was the aggressor (the side that crossed the spread with a market order). Those numbers live bar by bar. This chapter introduces the first tool that summarises all of that volume into a single, slowly evolving line on the chart. Candles show *what* happened to price; volume shows *who, where and how hard*. VWAP is the simplest way to fold the "where and how hard" into one number: the average price at which the market has actually done business.

What Is VWAP and Why Do Institutions Care?

VWAP (Volume-Weighted Average Price) is the average price at which trading took place over a defined period, where each price counts in proportion to the **volume** (number of contracts) that traded there. That second half of the definition is the whole point. A simple average treats every price the same. VWAP asks a sharper question: *at what price did most of the actual business happen?*

A tiny example makes the difference obvious. Suppose that in some short window 1,000 gold futures contracts trade at \$4,200.0 and only 10 contracts trade at \$4,210.0. The simple average of those two prices is \$4,205.0, right in the middle. But almost nobody traded in the middle. VWAP weights each price by its volume:

$$\text{VWAP} = (4,200.0 \times 1,000 + 4,210.0 \times 10) \div (1,000 + 10) = 4,242,100 \div 1,010 \approx \mathbf{4,200.1}$$

So VWAP sits almost exactly where the business was done. The 10 contracts at \$4,210 barely move it, which is exactly what you would want from a measure of "the typical traded price".

Born as an execution benchmark

VWAP was not invented for retail chart reading. It became popular as a benchmark for **execution quality**, that is, how well (how cheaply or expensively) a large order was filled. A well-known 1988 paper by Berkowitz, Logue and Noser in the *Journal of Finance*, "The Total Cost of Transactions on the NYSE", used the day's VWAP as the reference against which institutional trades were measured. The logic is simple. If a pension fund buys a large block of shares throughout the day and its average purchase price is below that day's VWAP, it bought more cheaply than the average participant, and its trading desk is judged to have executed well. If it paid above VWAP, it did worse than average.

That idea is still everywhere. Many large orders today are handled by **VWAP algorithms** (often called "VWAP algos"): programs that slice a big order into many small pieces and release them in proportion to the market's typical volume pattern, so that the final average fill lands close to the day's VWAP. For a desk that has to buy 2,000 contracts without moving the market, "match VWAP" is a reasonable, measurable goal.

Gold note: Gold futures have their own built-in connection to VWAP. Under CME's settlement procedures, the daily settlement price of the active GC contract is based on the volume-weighted average price of trades in a short window, 13:29–13:30 New York time. In other words, the exchange itself uses VWAP logic to decide the official daily price. Procedures can change, so check CME's current settlement documentation before relying on the exact details.

What VWAP tells you, and what it does not

At any moment during the session, VWAP tells you where the average traded price sits *so far*, counting from its starting point. That is **context**: it tells you whether the current price is expensive or cheap relative to the business done today. If gold trades at \$4,212 and VWAP is \$4,203, current trades are happening well above the day's average transaction price. That is a fact about the present and the past. It is not a forecast.

Two features of VWAP will matter in every section that follows:

- **VWAP is cumulative.** It includes every trade from its start point (called the **reset** or **anchor**) up to now. Nothing ever drops out.
- **VWAP only has its precise meaning on real exchange volume.** Futures such as GC (100 troy ounces) and MGC (10 troy ounces) report every contract traded on one central exchange. Most CFD (contract for difference) brokers offering XAUUSD show only **tick volume**, a count of how many times the broker's own price feed changed. Tick volume is not the number of contracts that changed hands, so a "VWAP" built from it is a rough, broker-specific approximation.

Key idea: VWAP is the market's average traded price, weighted by volume, measured from a chosen start point. It began as a benchmark for grading executions, not as a prediction tool.

Common mistake: "Institutions always defend VWAP." This is a popular claim, not proven. Being a benchmark means fills are *compared* against VWAP after the fact; it does not mean anyone is obliged to buy or sell there. A related claim, "price always returns to VWAP like a magnet", is also a popular claim, not proven. On strong trend days price may never come back to it.

The VWAP Formula, Step by Step

The general formula is short:

$$\text{VWAP} = \Sigma(\text{Price} \times \text{Volume}) \div \Sigma(\text{Volume})$$

The symbol Σ (the Greek capital letter sigma) simply means "add up". So: multiply each traded price by the volume traded at that price, add all those products together, and divide by the total volume.

If you had every single trade, you would apply this directly to each one. On a candle chart, however, the platform usually knows only the open, high, low, close and volume of each candle. So it picks one representative price per candle. The most common choice is the **Typical Price (TP)**:

$$\text{TP} = (\text{High} + \text{Low} + \text{Close}) \div 3$$

With that, the calculation becomes five mechanical steps:

1. Compute the typical price of each candle.
2. Multiply the typical price by that candle's volume ($\text{TP} \times \text{V}$).
3. Keep a running total of $\text{TP} \times \text{V}$.
4. Keep a running total of volume.
5. Divide the running total of $\text{TP} \times \text{V}$ by the running total of volume.

A worked example on micro gold

Here are four one-minute MGC candles. The prices are simplified and hypothetical so the arithmetic is easy to follow; real gold prices will differ.

Candle	High	Low	Close	TP	Volume	TP × V	$\Sigma \text{ TP} \times \text{V}$	$\Sigma \text{ V}$	VWAP
1	4202	4198	4200	4200	100	420,000	420,000	100	4200.00
2	4204	4200	4202	4202	100	420,200	840,200	200	4201.00
3	4206	4202	4204	4204	700	2,942,800	3,783,000	900	4203.33
4	4203	4199	4201	4201	100	420,100	4,203,100	1,000	4203.10

Check one row yourself. Candle 3: $\text{TP} = (4206 + 4202 + 4204) \div 3 = 4204$. Multiply by 700 contracts: 2,942,800. Add to the previous running total of 840,200: 3,783,000. Running volume: $200 + 700 = 900$. Divide: $3,783,000 \div 900 = 4,203.33$.

Now compare two "averages" of these four candles:

- The simple average of the four typical prices is $(4200 + 4202 + 4204 + 4201) \div 4 = \mathbf{4201.75}$.
- The VWAP after candle 4 is **4203.10**.

Why the difference of more than a dollar? Candle 3 carried 700 of the 1,000 contracts, which is 70% of all the volume. VWAP is pulled towards where the volume was. That is what "volume-weighted" means in practice: **the heavy candle wins**.

Gold note: These are hypothetical futures prices. If you wanted to mark the same level on a spot-gold CFD chart, you would subtract the current futures-to-spot gap, which has recently been roughly \$27–30 (always approximate, and it changes over time). A GC VWAP of 4203.1 would sit at roughly 4173–4176 on XAUUSD. Read the live gap at the time; never hard-code it.

Why VWAP slows down as the day goes on

Look again at candle 4. Price dropped back to a typical price of 4201, more than two dollars below the VWAP. Yet VWAP only moved from 4203.33 to 4203.10, about 23 cents. The reason is that 100 new contracts are being averaged against 900 old ones.

You can see this precisely with an equivalent "update" form of the formula:

$$\text{New VWAP} = \text{Old VWAP} + V_{\text{new}} \times (P_{\text{new}} - \text{Old VWAP}) \div (\text{Old } \Sigma V + V_{\text{new}})$$

For candle 4: $4203.33 + 100 \times (4201 - 4203.33) \div (900 + 100) = 4203.33 - 0.233 \approx 4203.10$.

The denominator keeps growing all session. Early in the day a single busy candle can shove VWAP around; late in the day the same candle barely nudges it. This is not a flaw. It is simply what a cumulative average does. A trader who complains that "VWAP is too slow in the afternoon" is really observing that the afternoon VWAP already contains most of the day's business.

A few practical consequences:

- Dividing by the **number of candles** instead of the total volume is the most common hand-calculation error. That gives you a simple average, not VWAP.
- VWAP is expressed in price units: dollars per troy ounce for gold.
- Two platforms fed with identical data and identical settings should show practically the same VWAP. When they differ, the cause is usually a settings difference (session start, price source, data resolution), which the next two sections cover.

Try it: Candle A trades at a price of 100 with a volume of 900. Candle B trades at 110 with a volume of 100. Before calculating, guess whether VWAP is closer to 100 or 110. Then compute it: $(100 \times 900 + 110 \times 100) \div 1,000 = 101$. The simple average would be 105.

Which Price Goes Into VWAP? Typical Price vs Tick Data

The "ideal" VWAP is computed from **every individual trade**: the price of each trade multiplied by its size, summed and divided by the total number of contracts. NinjaTrader's documentation for its order-flow VWAP describes it in exactly these terms: total dollars traded divided by total contracts traded. We will call this **tick-level VWAP**, because it is built from tick-by-tick trade data, the raw record of each transaction.

Most charts, though, are built from candles, and a candle does not record *where inside its range* each contract traded. So the platform chooses a single **source price** for each candle and pretends the candle's entire volume traded there. Common choices:

Source	Formula	Notes
HLC3 (Typical Price)	$(\text{High} + \text{Low} + \text{Close}) \div 3$	Default on many platforms, including TradingView
HL2	$(\text{High} + \text{Low}) \div 2$	Midpoint of the candle's range
OHLC4	$(\text{Open} + \text{High} + \text{Low} + \text{Close}) \div 4$	Average of the four prices
Close	Close	Uses only the final price of each candle

Every one of these is an **assumption**: that the candle's volume was concentrated at one price. On a one-minute candle in a quiet market the error is tiny, because the candle's whole range may be only a few ticks. On a 30-minute candle during a news release, the error can be large.

How big can the error be?

Consider one 30-minute GC candle on a news day (hypothetical numbers): open 4201, high 4220, low 4200, close 4218, with 1,000 contracts traded. Suppose that, inside the candle, the volume was distributed like this:

Price	Contracts traded
4220	20
4218	190
4216	190
4204	290
4202	290
4200	20
Total	1,000

Most of the business happened low in the candle, around 4202–4204, before the release; then price jumped and traded less. The true, trade-by-trade VWAP of this candle is:

$$(4200 \times 20 + 4202 \times 290 + 4204 \times 290 + 4216 \times 190 + 4218 \times 190 + 4220 \times 20) \div 1,000 = \mathbf{4208.6}$$

Now compare what each candle-based source would assume:

Method	Assumed price	Error vs 4208.6
Tick-level (true)	4208.6	—
OHLC4	$(4201 + 4220 + 4200 + 4218) \div 4 = 4209.75$	+1.15
HL2	$(4220 + 4200) \div 2 = 4210.0$	+1.40
HLC3	$(4220 + 4200 + 4218) \div 3 \approx 4212.67$	+4.07
Close	4218.0	+9.40

None of the shortcuts "knows" that most contracts traded near the bottom of the candle. On this one candle the Close method is off by more than nine dollars, about 94 ticks. On one-minute candles the same news move would be split into many small candles, each with a much narrower range, and the approximation error would shrink dramatically.

Resolution: calculate fine, view coarse

A second setting is **resolution**: what data the platform uses underneath the chart you are looking at. Some platforms, NinjaTrader among them, can compute VWAP from one-tick data even while you view a five-minute chart. That is more accurate but heavier on your computer.

A practical rule for intraday gold work: calculate VWAP on the finest data you can (tick data or one-minute bars), even if you look at a five-minute or fifteen-minute chart.

Key idea: When your VWAP and someone else's VWAP disagree by a few ticks, neither is necessarily "wrong". Check three settings first: session (reset time), source price, and data resolution.

Common mistake: Believing that choosing HLC3 over Close, or vice versa, gives you an "edge". On fine data the difference is usually small. The source setting is about accuracy and consistency, not about secret advantages. Equally, calculating VWAP from daily candles and expecting intraday precision is a mismatch of tools.

When Does Gold's VWAP Reset? Globex vs RTH

A daily VWAP needs a **starting point**, the moment the running totals are set back to zero. For US stocks this is easy: the regular session opens at 09:30 New York time and that is the obvious reset.

Gold futures are different. As covered in Chapter 1, GC and MGC trade on CME Globex, the exchange's electronic platform, from **Sunday 18:00 to Friday 17:00 New York time (ET)**, with a **daily halt from 17:00 to 18:00 ET**. CME's trading day for gold begins at 18:00 ET on the *previous* evening. That leaves two common choices for the reset.

Option 1: Globex reset (18:00 ET)

The VWAP starts at 18:00 ET and includes everything: the Asian session, London, and New York.

- **Advantage:** it lines up with CME's official trading day and leaves no volume out.
- **Disadvantage:** by the New York morning, a large amount of overnight volume is already baked in. If you mainly care about how US participants are pricing gold, the overnight weight can make VWAP less representative of the "main" session.

Option 2: RTH reset (about 08:20 ET)

Many platforms define a **Regular Trading Hours (RTH)** session for COMEX metals based on the old trading-floor ("pit") hours, roughly **08:20 to 13:30 ET**. The rest of the day is then called the **Electronic Trading Hours (ETH)** or overnight session.

- **Advantage:** it focuses on the hours with heavy US participation, the usual time of major US economic releases (often 08:30 ET), and the daily settlement window (13:29–13:30 ET).
- **Disadvantage:** it ignores overnight volume, and sometimes the most important move of the day happened overnight in London.

Neither choice is "more correct". They answer different questions. A Globex-reset VWAP answers "what is the average traded price of CME's whole trading day?" An RTH-reset VWAP answers "what is the average traded price during the main US hours?" Some traders display both at once, one for the overnight and full session and one for RTH.

Our rule throughout this book is short: **pick one, label it on the chart, and stay consistent**. A VWAP without a session label cannot be compared with anything.

Gold note: Session times are quoted in New York time. When you convert to your local time zone, remember that daylight saving time changes on different dates in different countries, so the local equivalent of 18:00 ET can shift by an hour for a few weeks each spring and autumn. Always check CME's official hours before relying on them.

Common mistake: Comparing your VWAP with a screenshot from a video or a colleague without knowing which session they used. Another frequent one is assuming gold "opens" at 09:30 ET like stocks; it has been trading for many hours by then.

Try it: Put two VWAPs on the same five-minute GC chart, one reset at 18:00 ET and one at 08:20 ET. At 10:00 ET, write down the distance between them. Do this for five days and note on which days the overnight session pulled the Globex VWAP far from the RTH one.

How VWAP Standard Deviation Bands Are Calculated

Most VWAP indicators can draw **bands** above and below the VWAP line. To understand them you need one statistical idea.

The **standard deviation**, written σ (lower-case sigma), measures how far values typically sit from their average. A small σ means prices clustered tightly around VWAP; a large σ means they were spread out. The **variance** is σ squared: the average of the *squared* distances from the mean. We square the distances so that values above and below the mean do not cancel each other out, and so that far-away values count more heavily.

For VWAP, the common and logical version is **volume-weighted**. Sierra Chart documents it in this form:

$$\sigma^2 = \Sigma [V \times (TP - VWAP)^2] \div \Sigma V, \text{ and } \sigma = \sqrt{\sigma^2}$$

An equivalent shortcut, handy in a spreadsheet, is $\sigma^2 = \Sigma (V \times TP^2) \div \Sigma V - VWAP^2$.

The bands are then:

$$\text{Upper band } k = VWAP + k \times \sigma, \text{ and Lower band } k = VWAP - k \times \sigma$$

where k is usually 1 and 2, sometimes 3.

Worked example: continuing the four candles

From the previous example, VWAP = 4203.10 and total volume = 1,000.

Candle	TP	Distance from VWAP	Distance ²	Volume	V × Distance ²
1	4200	-3.10	9.61	100	961
2	4202	-1.10	1.21	100	121
3	4204	+0.90	0.81	700	567
4	4201	-2.10	4.41	100	441
Total				1,000	2,090

$$\sigma^2 = 2,090 \div 1,000 = 2.09, \text{ so } \sigma = \sqrt{2.09} \approx \text{\$1.45}.$$

The bands, rounded to gold's tick of \$0.10:

Band	Calculation	Level
+2σ	$4203.10 + 2 \times 1.4457$	4206.0
+1σ	$4203.10 + 1.4457$	4204.5
VWAP		4203.1
-1σ	$4203.10 - 1.4457$	4201.7
-2σ	$4203.10 - 2 \times 1.4457$	4200.2

Notice something about candle 1. It had only 100 contracts, the same as candles 2 and 4, yet it contributes 961 of the 2,090 total, nearly half. That is because it sat furthest from VWAP and the distance is **squared**. Prices far from the average have an outsized effect on σ .

Bands are cumulative too, and noisy early

Like VWAP, σ is computed from the reset onward. Watch how it evolved in our four-candle example:

After candle	VWAP	σ	$\pm 1\sigma$ band width
1	4200.00	0.00	0.00
2	4201.00	1.00	2.00
3	4203.33	1.33	2.67
4	4203.10	1.45	2.89

After the first candle, with only one typical price, σ is zero and the bands collapse onto VWAP. With very little data, every new candle changes σ noticeably. As the session progresses and volume accumulates, the bands usually widen and become steadier. In practice, the bands in the first fifteen to thirty minutes after a reset are unreliable simply because they are built on too little information.

Common mistake: Thinking the $\pm 1\sigma$ band is a fixed dollar distance. It changes every day and every minute. A related mistake is comparing bands from two platforms that use different formulas: some use a volume-weighted σ , some a plain standard deviation, and some (Sierra Chart among them) also offer percentage or fixed-distance bands. Read your platform's documentation before comparing.




Figure 10. VWAP (solid) with $\pm 1\sigma$ and $\pm 2\sigma$ bands on a real gold session. Bands widen as the day's prices spread out.

What VWAP Bands Do (and Don't) Tell You

Here is where many traders go wrong, so it is worth slowing down.

In statistics, if data follow a **normal distribution** (the familiar symmetric "bell curve"), about **68%** of values lie within $\pm 1\sigma$ of the mean and about **95%** within $\pm 2\sigma$. This is called the 68–95 rule. Many traders transfer it straight onto VWAP bands: "Price is at $+2\sigma$, so there is a 95% chance it reverses." That conclusion is wrong, for three separate reasons.

Reason 1: prices are not normally distributed

Market returns have **fat tails**: very large moves happen far more often than a normal distribution would predict. Benoit Mandelbrot made this point famously (see *The (Mis)Behavior of Markets*, with Richard Hudson, 2004). In gold, the tails show up most clearly around major US releases such as CPI (inflation), NFP (the monthly jobs report) and FOMC (the Federal Reserve's interest-rate decisions). On those days, "three-sigma" or "four-sigma" moves relative to the morning's bands are not rare at all.

Reason 2: σ describes the past, not the future

The band tells you, "so far today, trades have typically been this far from the average." If the market changes character, for example when a news release arrives, σ adapts only slowly, because it is cumulative. Even the statement "95% of today's volume traded inside $\pm 2\sigma$ " would be a description of trades that already happened in this session, not a probability for the next trade.

Reason 3: the bands move with price

VWAP and σ are both cumulative and both respond to new trades. On a strong trend day, price can **ride** the $+1\sigma$ or $+2\sigma$ band for hours: price climbs, new volume trades at higher prices, VWAP rises, σ widens, and the band climbs right behind price. The band is not a wall; it is attached to the thing it is measuring.

So what are the bands good for?

Bands are a **normalised ruler for stretch**. Saying "price is \$6 above VWAP" means different things on different days. On a quiet day with $\sigma = \$3$, \$6 above VWAP is 2σ , a large stretch relative to that day's activity. On a CPI day with $\sigma = \$12$, \$6 above VWAP is only 0.5σ , barely noteworthy. Expressing distance in σ units lets you compare across days and volatility regimes:

Distance in $\sigma = (\text{Price} - \text{VWAP}) \div \sigma$

Using our worked example, if price were 4205.0 with VWAP 4203.1 and σ 1.45, the distance would be $(4205.0 - 4203.1) \div 1.45 \approx 1.3\sigma$.

That is a useful *description*. Whether anything happens at that stretch is a separate question, answered by other evidence: the order-flow behaviour at that price (absorption, aggression, delta, covered in Chapters 3, 8 and 9), the structure of the day, and the volume profile (Chapter 5).

Key idea: " $+2\sigma$ " honestly means "stretched relative to today's traded distribution so far." It does not mean "95% likely to reverse."

Common mistake: " $\pm 2\sigma$ means a 95% chance of reversal" is a popular claim, not proven, and it misapplies the statistics. So is treating $\pm 1\sigma$ as an automatic buy or sell line. Treating every touch of $+2\sigma$ as a turning point on a trend day is a classic way to be repeatedly on the wrong side of a band ride.

Try it: Pick two recent GC days: one quiet, range-bound day and one major-news day. On each, note the σ value at 11:00 ET and how many times price closed a five-minute bar outside $\pm 2\sigma$. You will see how differently the same band setting behaves.

Daily, Weekly and Monthly VWAP

The VWAP formula never changes. Only the **reset point** changes:

- **Daily VWAP:** from the start of the session (Globex 18:00 ET, or RTH about 08:20 ET).
- **Weekly VWAP:** from the start of the trading week; for gold, Sunday 18:00 ET.
- **Monthly VWAP:** from the first trading session of the month.

Each one reflects a different **time horizon** of participants. The daily VWAP is the average traded price of today's business, dominated by intraday traders. The weekly and monthly VWAPs average

much longer stretches, which include the volume of slower participants: swing traders who hold for days, funds that build positions over weeks, and commercial hedgers such as miners and jewellers.



Figure 11. Daily VWAP restarts every session; weekly VWAP keeps averaging all week. Where they meet or diverge tells you whose average price is in control.

Stacking context

Putting all three on one chart gives you a layered description. Suppose gold is trading above the daily VWAP but below both the weekly and the monthly VWAP. A careful, descriptive reading is: *today, business is being done above today's average price, but on the weekly and monthly horizons price is still below the average transaction price.* That is not a signal to do anything. It is a reminder of which bigger picture today's activity sits inside.

Price vs Daily	Price vs Weekly	Price vs Monthly	Descriptive reading
Above	Above	Above	Trading above the average on all three horizons
Above	Below	Below	Today is firm, but still below the longer-term averages
Below	Above	Above	A soft day inside a firmer longer-term picture
Below	Below	Below	Trading below the average on all three horizons

Why the monthly VWAP goes flat

By late in the month, the monthly VWAP contains an enormous amount of volume, so each new bar has almost no weight. Using illustrative round numbers and the update formula from earlier:

- Daily VWAP at 08:30 ET with, say, 40,000 contracts traded so far. A five-minute bar of 2,000 contracts trades \$10 above it. The VWAP moves by $2,000 \times 10 \div 42,000 \approx \text{\$0.48}$.
- Monthly VWAP on the twentieth trading day with, say, 4,000,000 contracts so far. The same bar moves it by $2,000 \times 10 \div 4,002,000 \approx \text{\$0.005}$, less than a tenth of one tick.

So a nearly flat monthly VWAP late in the month is completely normal.

Watch out for the contract roll

Futures contracts expire, and traders move their positions from the expiring contract to the next one, a process called the **roll** or **rollover**. Gold's main active months are usually February, April, June, August, October and December, and volume migrates to the next active contract towards the end of each cycle. A weekly and especially a monthly VWAP can straddle a roll.

How this affects you depends on your price series:

- A **continuous back-adjusted** series shifts historical prices so the chart has no jump at the roll. Old prices on the chart are then not the prices that actually traded.
- A **non-adjusted** series may show a sudden jump equal to the price difference (the **spread**) between the two contract months.

Either way, a long VWAP that crosses a roll needs care. Always know which price series your weekly or monthly VWAP is calculated on.

Common mistake: "The monthly VWAP is guaranteed support or resistance" is a popular claim, not proven. Also common: comparing a monthly VWAP computed on two different contract months, and covering the chart with ten VWAPs without a clear question for each.

Anchored VWAP: Start the Clock at an Event

An **Anchored VWAP (AVWAP)** is an ordinary VWAP whose starting point *you* choose, instead of letting the calendar choose it. The formula is unchanged: $\sum(TP \times V) \div \sum V$ from the anchor candle up to now. Brian Shannon of Alphatrends popularised the approach (see his book *Maximum Trading Gains with Anchored VWAP*, 2023). His central idea is that the average traded price is most meaningful when measured from a **significant event**, not from an arbitrary clock time.

Common anchors for gold:

- **News events:** the CPI or NFP release (usually 08:30 ET) or the FOMC statement (usually 14:00 ET).
- **Significant swing highs or lows:** points where the market clearly changed direction.
- **Session or week starts:** which simply reproduce the daily or weekly VWAP.

- **Unusual days:** an exceptionally high-volume day, or a price gap after a weekend or holiday.

What an AVWAP measures

An AVWAP anchored at an important low is the average price of every contract traded from that low until now. Shannon's common interpretation is that above a rising AVWAP, the participants who have traded since that event are, on average, in profit, and below a falling AVWAP the reverse.

That interpretation is a useful simplification, but be precise about its limits. Every contract has a buyer and a seller. VWAP is the average price of **trades**, not the average cost of **open positions**. Some of those trades opened positions, some closed them, and the buyer of each contract had a seller who now holds the opposite result. "Who is in profit" is a helpful mental picture, not a measurement.

The weak spot: choosing the anchor

The main weakness of AVWAP is that the choice of anchor is **subjective**, and Shannon himself acknowledges this. The danger is **hindsight bias**: after the fact, it is easy to pick the anchor that "worked best" and then believe the tool works. If you look at a chart, try five anchors, and keep the one whose line price happened to respect, you have learned nothing about the future.

The remedy is a rule written *before* you look: for example, "always anchor at the 08:30 ET candle on CPI days" or "always anchor at the low of the prior week". Then apply that rule every time, including on the days it looks unimpressive.

Key idea: Anchored VWAP is the same formula with a deliberately chosen window. Its value depends entirely on choosing that window by a rule, not by hindsight.

Common mistake: "Price will definitely return to the AVWAP from the last low" is a popular claim, not proven. Another frequent error is anchoring to insignificant candles until the chart is a tangle of lines.

Try it: Write down one anchor rule today, for example "the 14:00 ET candle on FOMC days". Apply it to the last six FOMC days on a 15-minute GC chart without skipping any. Note honestly what you see on each.

Price vs VWAP and VWAP Slope as Context

VWAP gives you two simple, descriptive pieces of information.

1. Location

If price is **above** VWAP, current trades are happening at prices **more expensive than the average** traded price since the reset. **Below** VWAP means cheaper. When price stays above VWAP for a

long time, the market is showing **acceptance** of higher prices: it is willing to keep doing business up there. (Acceptance will be defined more carefully with volume profiles in Chapters 5 to 7.)

2. Slope

Because VWAP is cumulative, its **slope** tells you where new business is being done relative to the old average. A rising VWAP means new trades keep happening above the existing average; a falling VWAP means they keep happening below it. A **flat** VWAP means trading is rotating around one fair price, which is a sign of **balance**: buyers and sellers broadly agree on value and the market is not moving away from it.

A simple context matrix

Combining location and slope gives a simple descriptive grid:

Price location	VWAP slope	Descriptive reading
Above	Rising	Higher prices are being accepted; trending context upwards
Below	Falling	Lower prices are being accepted; trending context downwards
Crossing back and forth	Flat	Rotational, balanced session
Above	Flat or falling	Possibly a change, possibly just a short move; more evidence needed
Below	Flat or rising	Possibly a change, possibly just a short move; more evidence needed

Shannon puts it simply: above a rising VWAP, buyers are in control until proven otherwise. We treat this as **context**. Its job is to help you recognise whether the market is in a trending state or a balanced one, so that you know which tools and which questions are appropriate. It is not an entry rule.

Add the σ ruler from earlier. "Price is above VWAP by 0.3σ " and "price is above VWAP by 2.5σ " are very different stories, even though both are "above VWAP".

Be cautious with slope early in the session. In the first minutes after a reset, VWAP is built on very little volume and its slope can swing wildly. Also remember that the slope depends on the session you chose: an RTH VWAP and a Globex VWAP on the same morning can slope in different directions.

Common mistake: "Above VWAP, only buy; below VWAP, only sell." This is a popular claim, not proven: a simplification that cannot replace context. Reading slope in the first ten minutes after the reset is another frequent error.

Try it: Find three recent GC days on a five-minute chart with an RTH VWAP: one that trended up, one that trended down, and one that rotated. Count how many times price crossed VWAP on each day and describe each day using the matrix above.

VWAP vs Moving Average: Not the Same Line

VWAP and a moving average are both "averages" and can look similar on a chart. They answer different questions.

A **moving average (MA)** is the average of the last N closing prices. A **Simple Moving Average (SMA)**, such as the SMA 20, gives each of the last 20 candles equal weight. An **Exponential Moving Average (EMA)** gives more weight to recent candles and less to older ones. There are three structural differences between these and VWAP.

1. The window

A moving average uses a **rolling** window: the SMA 20 always averages the latest 20 candles; as a new candle arrives, the oldest one drops out. VWAP uses a **cumulative** window: everything from the reset or anchor up to now, with nothing ever dropping out.

2. The weighting

The standard SMA weights all candles equally; the EMA weights recent candles more. Neither looks at volume in its standard form. VWAP weights by **volume**: a candle with 5,000 contracts has fifty times the influence of a candle with 100 contracts.

3. The economic meaning

VWAP is a real quantity: the average price at which contracts actually changed hands since the reset, the very number a fund's trading desk uses to grade its executions. A moving average is a statistical smoothing tool. The "20" or "50" in its name is a convention among traders, not a property of the market.

Feature	Moving average (SMA/EMA)	VWAP
Window	Rolling: last N candles	Cumulative: from reset or anchor
Weighting	Equal (SMA) or recency (EMA)	By volume
Uses volume?	No (standard versions)	Yes
Resets?	No; carries yesterday's data	Yes, at each session (or anchor)
Responsiveness	Constant, set by N	Fast early in the session, slow late
Economic meaning	Statistical smoother	Average traded price; execution benchmark

The practical consequence: a 20-period EMA reacts to a sharp 11:00 ET move at roughly the same speed it would at 09:00 ET. VWAP reacts strongly at 09:00 (little accumulated volume) and only slightly at 12:00 (lots of accumulated volume). Neither is "better". VWAP is the right tool for the question "where is price relative to today's business?"

There is also a hybrid called **Rolling VWAP**: a VWAP calculated over a rolling window of the last N candles. It weights by volume like VWAP but forgets old data like a moving average.

Common mistake: "VWAP is just an EMA in a different colour." It is not; the window, weighting and meaning are all different. Expecting VWAP to react at noon as quickly as it did at 9 a.m. is another frequent misunderstanding. And "the 200 MA or VWAP is a magic line" is a popular claim, not proven.

7 Common VWAP Mistakes

Here is the chapter condensed into seven mistakes, each with the section where it was explained.

- 1. Not knowing the reset session.** A VWAP reset at 18:00 ET and one reset at 08:20 ET are two different numbers. Decide and label your session first (see "When Does Gold's VWAP Reset?").
- 2. Using data without real volume.** On most CFD platforms, XAUUSD "volume" is tick volume: the number of price changes in that broker's own feed, not centralised exchange volume. A VWAP built on it is approximate and depends on the broker. GC and MGC futures report real exchange volume. If you trade XAUUSD, a sounder approach is to read VWAP on the futures and translate the level using the current futures-to-spot gap (recently roughly \$27–30, always approximate).
- 3. Treating VWAP as a magnet.** "Price always returns to VWAP" is a popular claim, not proven. On trend days it may never return.
- 4. Treating bands as probabilities.** $\pm 2\sigma$ is not a 95% chance of reversal (see "What VWAP Bands Do (and Don't) Tell You").
- 5. Trusting VWAP and bands in the first minutes.** With little cumulative volume, every candle shifts VWAP and σ noticeably.
- 6. Ignoring the contract roll** on weekly and monthly VWAPs (see "Daily, Weekly and Monthly VWAP").
- 7. Replacing context with a signal.** VWAP is a reference. Price reaching VWAP means only that price has reached today's average traded price. What happens there, whether aggressive orders are absorbed, whether delta shifts, whether the structure holds, has to be read with order-flow tools.

There is an eighth mistake that affects learning more than trading: **selection bias in examples**. If you only ever look at, or are only ever shown, days when price bounced neatly off VWAP, you will build a false picture. For every clean example you study, find one where the same idea failed, using

the same settings. This book tries to follow that rule, and you should demand it from any educational material.

Your VWAP setup checklist

Check	Question to answer
Session	Globex (18:00 ET) or RTH (about 08:20 ET)? Is it labelled?
Source	Tick data, HLC3, HL2, OHLC4 or Close?
Resolution	What data is VWAP calculated on underneath the chart?
Band formula	Volume-weighted σ , plain σ , or something else?
Contract	Which contract or continuous series? Any roll inside the window?
Volume	Is this real exchange volume (futures) or tick volume (CFD)?

Common mistake: "VWAP alone is a complete strategy" and "institutions hide orders at VWAP" are both popular claims, not proven. Comparing a CFD gold VWAP directly with a futures VWAP is also a mismatch: different instruments, different volume, and a price gap between them.

Key idea: VWAP is more information about where business has been done. It is not a magic edge. Our own testing of order-flow information in general found that signals on their own did not reliably predict direction once trading costs were included, which is why this book treats every tool as context to be combined, not as a trigger.

Chapter summary

- VWAP is the volume-weighted average traded price from a chosen start point: $\Sigma(\text{Price} \times \text{Volume}) \div \Sigma \text{Volume}$.
- It began as an institutional execution benchmark (Berkowitz, Logue and Noser, 1988), and CME uses a short VWAP window (13:29–13:30 ET) in settling the active GC contract.
- On candle charts VWAP uses a source price per candle, most often the typical price $(H + L + C) \div 3$; this is an approximation of the true tick-level VWAP, and the error grows with candle size.
- VWAP is cumulative: it is responsive early in the session and slow later, by design.

- For gold you must choose a reset: Globex (18:00 ET, the full CME day) or RTH (about 08:20 ET, the main US hours). Label it and stay consistent.
- Standard deviation bands are $VWAP \pm k \times \sigma$, with σ usually volume-weighted. They are cumulative, noisy early, and implemented differently across platforms.
- The 68–95 rule does not turn bands into reversal probabilities: prices have fat tails, σ describes the past, and bands move with price. Use bands as a ruler of stretch.
- Daily, weekly and monthly VWAPs share one formula with different resets; together they layer context. Long VWAPs can be distorted by contract rolls.
- Anchored VWAP starts the clock at an event; choose anchors by a written rule to avoid hindsight bias.
- Location (above or below) and slope (rising, flat, falling) describe trending versus balanced context, not entries.
- VWAP differs from moving averages in window, weighting and economic meaning.

Checklist

- I can define VWAP in one sentence and explain why it is a benchmark rather than a forecast.
- I can calculate VWAP by hand for a few candles using typical price and running totals.
- I know my platform's VWAP source price and data resolution.
- I have chosen a session reset for gold (Globex or RTH) and labelled it on my chart.
- I can calculate a volume-weighted σ and the $\pm 1\sigma$ and $\pm 2\sigma$ bands for a small example.
- I describe price distance from VWAP in σ units, not just dollars.
- I do not treat $\pm 2\sigma$ as a 95% reversal probability.
- I know which contract or continuous series my weekly and monthly VWAPs use, and when the last roll was.
- If I use anchored VWAP, my anchor rule is written down before I look at the chart.
- I use VWAP on real exchange volume (GC/MGC) and translate levels to XAUUSD with the live, approximate gap.

Quiz

1. Candle A trades at a price of 100 with a volume of 900. Candle B trades at 110 with a volume of 100. What is the VWAP of the two candles, and what is their simple average?
2. Why does a session VWAP move much less in response to a busy candle at 12:00 ET than to the same candle at 08:40 ET?
3. VWAP is 4203.1 and the volume-weighted σ is 1.45. Where are the -2σ band and the $+1\sigma$ band (rounded to the tick)?
4. Price is at $+2\sigma$ above VWAP. Which statement is honest: (a) there is a 95% chance price reverses; (b) price is stretched relative to today's traded distribution so far; (c) sellers must

step in here?

5. Name two reasons why a VWAP drawn on a spot-gold CFD chart is less reliable than one drawn on GC futures.

Quiz answers

1. $VWAP = (100 \times 900 + 110 \times 100) \div 1,000 = 101,000 \div 1,000 = \mathbf{101}$. The simple average is $(100 + 110) \div 2 = \mathbf{105}$. VWAP sits near the price where most of the volume traded.
2. Because VWAP is **cumulative**. At 12:00 ET the running total of volume is much larger, so the new candle's volume is a much smaller fraction of the whole and its weight in the average is smaller. The update formula shows this: the change equals $V_{\text{new}} \times (P_{\text{new}} - VWAP) \div (\text{total volume including the new candle})$.
3. $-2\sigma = 4203.1 - 2 \times 1.45 \approx \mathbf{4200.2}$; $+1\sigma = 4203.1 + 1.45 \approx \mathbf{4204.5}$ (approximately 4170–4177 on XAUUSD with a gap of roughly \$27–30).
4. **(b)**. The band describes stretch relative to the trades so far today. Price distributions have fat tails, σ describes the past, and bands move with price, so (a) misapplies the 68–95 rule; (c) is an unsupported claim.
5. Any two of: CFD "volume" is usually broker **tick volume** (price-change counts), not centralised exchange volume; it is specific to one broker's feed, so two brokers give different VWAPs; the CFD price differs from the futures price by a moving gap, so CFD and futures VWAPs cannot be compared directly.

5

Volume Profile I: Building and Reading a Profile

In this chapter: You will learn to turn the volume bars at the bottom of a chart sideways and read volume *at price* instead of volume over time. You will find the Point of Control (POC), build the 70% Value Area by hand with both the single-row and the two-row method, and learn where the 70% figure came from and why it is a convention rather than a law. We then read the edges of value (VAH and VAL), separate high-volume nodes from low-volume nodes, distinguish developing profiles from final ones (and the look-ahead trap that hides between them), and choose the right profile window: session, weekly, monthly, composite, fixed range and visible range. As always, the profile is a map of where the market has done business, not a forecast of where it will go.

Chapter 4 folded a whole session's volume into a single line, VWAP. That is useful, but a single average hides a lot. Two days can have the same VWAP while one spent the whole session in a narrow range and the other spent half the day at one price and half at another. The volume profile keeps that detail. Instead of one number, it gives you the full distribution: how much business was done at *every* price.

What Is a Volume Profile?

The ordinary volume bars under a price chart answer the question: **when** did trading happen, and how much? Each bar belongs to a time period. A **volume profile** answers a different question: **at which prices** did trading happen, and how much at each?

To build one, you take all the trades in a chosen period, sort them by price, and add up the volume at each price level. Then you draw a horizontal bar for each price, its length proportional to the volume traded there. The result is a **histogram on the price axis**, a bar chart turned on its side. A helpful picture: imagine the volume bars under the chart sliding sideways and stacking up next to the prices where they traded.

A small profile built by hand

Here is a hypothetical half-day of GC trading, split into three half-hour periods. Each cell shows how many contracts traded at that price during that period.

Price	09:00–09:30	09:30–10:00	10:00–10:30	Profile total
4206	0	20	60	80
4205	10	60	90	160
4204	40	120	60	220
4203	90	70	20	180
4202	50	10	0	60
Total	190	280	230	700

The bottom row is what ordinary volume bars would show: 190, then 280, then 230 contracts per half-hour. The right-hand column is the volume profile. Drawn as a sideways histogram, with each block representing 20 contracts:

Price	Volume	Profile
4206	80	
4205	160	
4204	220	
4203	180	
4202	60	

Both views use exactly the same 700 contracts. The time view tells you the middle half-hour was busiest; the price view tells you \$4,204 was where the most business was done, and that the market traded little at the extremes.

Why volume at price is meaningful

In **auction market theory** (introduced in Chapter 2 and developed fully in Chapter 7), the market is a continuous auction searching for the price that facilitates the most trade. Prices where a lot of volume traded are prices that both buyers and sellers were willing to accept: there was **two-sided agreement**. Prices with little volume are prices the market passed through quickly or rejected. The profile is a record of this **volume memory** laid out along the price axis.

Two settings that change the picture

Row size. Each horizontal bar can represent one tick or a group of ticks. Gold's tick is \$0.10, so a profile can have one row per tick or, for example, one row per \$1.00 (10 ticks). Larger rows produce a smoother but less detailed profile, and they can even move the POC (the busiest row, defined in the next section).

Here is the same 700-contract profile grouped into \$2.00 rows, in two different ways:

Grid A (\$2 rows starting at 4202)	Volume	Grid B (\$2 rows starting at 4201)	Volume
4206–4207	80	4205–4206	240
4204–4205	380	4203–4204	400
4202–4203	240	4201–4202	60

Same trades, same row size, but simply shifting where the grid starts moves the busiest row from 4204–4205 to 4203–4204. Whenever you compare two profiles, make sure they use the same row size and alignment.

Data source. The most accurate profile is built from **tick data**: each trade added at the exact price where it happened. When a profile is built from candles instead, the platform does not know where inside each candle the volume traded, so it has to assume something, typically that the candle's volume was spread evenly from its low to its high. A five-minute candle with a range from 4201 to 4205 and 500 contracts might be drawn as 100 contracts at each of five \$1 rows, even if in reality 300 of them traded at 4202. On fine data this approximation is mild; on large candles it can distort the shape noticeably.

Key idea: Ordinary volume is volume over time. A volume profile is volume at price. It shows where the market has spent its time and money, not where price will go next.

Common mistake: "The profile shows where institutions bought." This is a popular claim, not proven. A profile shows total volume at each price. Every contract has a buyer and a seller, so a heavy row is just as much "where institutions sold" as "where they bought". Other frequent mistakes: comparing profiles built with different row sizes, and building profiles on data without real exchange volume, such as CFD tick volume.

Gold note: As with VWAP, build your profile on futures (GC or MGC), where volume is real exchange volume, and translate the key levels to XAUUSD using the current futures-to-spot gap, recently roughly \$27–30 and always approximate.

POC: The Point of Control

The **Point of Control (POC)** is the price, or row, in the profile where the **most volume** traded: the longest bar of the histogram. In auction language, the POC is where the auction found the most agreement, the price at which the largest number of contracts changed hands. In our hand-built profile above, the POC is 4204 with 220 contracts.

Ties and tiny margins

What if two rows have exactly the same volume? Platforms use different **tie-break** rules. CQG's Market Profile documentation, for example, picks the row closer to the middle of the range. These details are small, but they explain why two platforms sometimes disagree on the POC by a tick.

More important than exact ties are *near* ties. Because the POC is just "the single winning row", a broad, flat profile can have two or three rows within a handful of contracts of each other, and the POC can jump between them with a few late trades. As you saw in the previous section, changing the row size or grid can also move it. So think of the POC as a **zone**, not a laser-thin line, especially with small row sizes.

POC is the mode; VWAP is the mean

POC and VWAP both describe the "centre" of the day's trading, but they are different statistics:

- **VWAP** is the **mean** (average) of all traded prices, weighted by volume.
- **POC** is the **mode** of the distribution: the single most-traded price.

On a symmetric, balanced day they sit close together. On a **skewed** day, where most volume is in one area with a long thin tail on one side, they separate: the mean is pulled towards the tail, while the mode stays where the volume is heaviest.

A hypothetical example. The morning rotates quietly around 4200–4202, then a fast afternoon rally trades thinly up to 4208:

Price	Volume
4208	50
4207	50
4206	50
4205	50
4204	50
4203	100
4202	250
4201	400
4200	200
Total	1,200

The POC is clearly **4201**. Now compute VWAP. To keep the numbers small, measure each price as an offset from 4200: $\Sigma(\text{offset} \times \text{volume}) = 1 \times 400 + 2 \times 250 + 3 \times 100 + 4 \times 50 + 5 \times 50 + 6 \times 50 + 7 \times 50 + 8 \times 50 = 2,700$. Divide by 1,200: 2.25. So VWAP = **4202.25**.

The VWAP sits more than a dollar above the POC because the thin upper tail pulls the average upward, while the busiest price stays at 4201. The distance between VWAP and POC is itself descriptive information: it tells you how lopsided the day's distribution is.

Gold note: On a hypothetical futures POC of 4201, the XAUUSD equivalent would be roughly 4171–4174, using an approximate gap of \$27–30. Always use the live gap.

What the POC is, and is not

The POCs of previous days are widely tracked as reference levels (you will use them in Chapter 6). That is reasonable: they mark where the market most agreed. But a reference level is a place to pay attention, not a guarantee of a reaction.

Common mistake: "The POC is always support or resistance" and "the POC is a magnet for price" are both popular claims, not proven. Treating POC and VWAP as the same thing is another error; on skewed days they can be several dollars apart.

Building the Value Area by Hand

The **Value Area (VA)** is the range of prices, built around the POC, in which roughly **70% of the profile's volume** traded. It is the "core" of the day's business. The algorithm is simple enough to do with pencil and paper:

1. Add up the total volume of the profile. The target is 70% of that total.
2. Start at the POC. It is the first row of the Value Area.
3. Look at the rows immediately above and below the current Value Area. Add the side with the **larger** volume.
4. Repeat until the Value Area's volume is **equal to or greater than** the target.

There are two common versions of step 3:

- **Single-row method:** compare one row above with one row below. TradingView describes its calculation this way.
- **Two-row method:** compare the *sum* of the next two rows above with the *sum* of the next two rows below, and add both rows of the winning side. This is the classic Market Profile approach.

The results are usually close, but they can differ by a row or two.

Worked example: an 11-row gold profile

A hypothetical GC session with \$1 rows:

Price	Volume
4210	20
4209	40
4208	80
4207	120
4206	200
4205	260 (POC)
4204	180
4203	150
4202	90
4201	40
4200	20
Total	1,200

Step 1: total = 1,200; target = 70% × 1,200 = **840**.

Single-row method:

Step	Candidate above	Candidate below	Added	VA volume	VA range
Start	—	—	4205 (260)	260	4205
1	4206 (200)	4204 (180)	Above: 4206	460	4205–4206
2	4207 (120)	4204 (180)	Below: 4204	640	4204–4206
3	4207 (120)	4203 (150)	Below: 4203	790	4203–4206
4	4207 (120)	4202 (90)	Above: 4207	910	4203–4207

After step 4 the Value Area holds 910 contracts, which is at least 840, so we stop.

Two-row method:

Step	Pair above	Pair below	Added	VA volume	VA range
Start	—	—	4205 (260)	260	4205
1	4206 + 4207 = 320	4204 + 4203 = 330	Below pair	590	4203–4205
2	4206 + 4207 = 320	4202 + 4201 = 130	Above pair	910	4203–4207

Both methods give the same answer here: **Value Area = 4203 to 4207**, containing 910 contracts, which is $910 \div 1,200 \approx 75.8\%$ of the volume.

Three things to notice:

- The Value Area usually ends up **a little above 70%**, because rows are added whole. You cannot add "part of" a row.
- The Value Area is **not necessarily symmetric** around the POC. It grows towards whichever side has more volume. Here it extends two rows on each side, but in the two-row method the lower side was added first because it was heavier.
- The highest row in the Value Area is the **Value Area High (VAH)**, here 4207, and the lowest is the **Value Area Low (VAL)**, here 4203. We examine them in their own section shortly.




Figure 12. Building the value area by hand: start at the POC, then keep adding the larger neighbouring row until 70% of the volume is inside.

Gold note: These are hypothetical futures prices. A Value Area of 4203–4207 on GC would correspond to roughly 4173–4180 on XAUUSD with a gap of about \$27–30, approximate and to be checked live.

When the two methods disagree

To see that the method matters, try a second, smaller hypothetical profile:

Price	Volume
4217	20
4216	40
4215	50
4214	90
4213	100 (POC)
4212	60
4211	30
4210	10
Total	400

Target = $70\% \times 400 = 280$.

Single-row: start 100. 4214 (90) vs 4212 (60): add above, 190. 4215 (50) vs 4212 (60): add below, 250. 4215 (50) vs 4211 (30): add above, **300**. Stop. VA = **4212–4215**, 300 contracts, 75.0%.

Two-row: start 100. Above pair 4214 + 4215 = 140 vs below pair 4212 + 4211 = 90: add above, 240. Above pair 4216 + 4217 = 60 vs below pair 4212 + 4211 = 90: add below, **330**. Stop. VA = **4211–4215**, 330 contracts, 82.5%.

The two methods disagree on the VAL by one full dollar (ten ticks). Neither is wrong; they are two conventions. This is one reason two platforms can show different VAH and VAL values for the same day.

Edge cases

- **One side runs out.** If the Value Area reaches the top or bottom of the profile, keep adding from the other side only.
- **Ties between sides.** If the candidates above and below are equal, platforms use their own rule (for example, adding both, or preferring one side). Check your platform's documentation.
- **Row size.** Because the Value Area is built from rows, a different row size can shift VAH and VAL.

Common mistake: Thinking the Value Area means "70% of the price range". It is 70% of the **volume**, not of the range. Expecting the Value Area to be symmetric around the POC, or expecting every platform to give identical VAH and VAL values, are the other two classic errors.

Try it: A profile has total volume of 2,000 contracts. At what Value Area volume do you stop expanding? (Answer: when it reaches at least 1,400.) Then take any real GC session, export the volume at each \$1 row, and build the Value Area by hand with both methods. Compare your result with your platform's.

Why 70%? The Origin of the Value Area

The Value Area comes from **Market Profile**, a way of organising market activity developed by **J. Peter Steidlmayer**, a trader at the **Chicago Board of Trade (CBOT)**, in the early 1980s and made public around 1985 (see Steidlmayer and Koy, *Markets and Market Logic*, 1986). His insight was that if you plot price on the vertical axis and activity on the horizontal axis, a typical day tends to look like a **bell curve** turned on its side: wide in the middle, narrow at both ends.

In a **normal distribution**, about **68%** of the data lie within **± 1 standard deviation** of the mean (the exact figure is 68.27%). Steidlmayer, and later **Jim Dalton** in *Mind Over Markets* (with Eric Jones and Robert Dalton), called this central portion **value**: the range where most of the day's business took place, the prices both sides considered fair. The 68% was rounded to **70%**.

Three honest points

- 1. Seventy per cent is a convention.** There is no empirical proof that 70% is "better" than 65% or 80%. Most platforms, TradingView included, let you change the percentage.
- 2. Real profiles are often not normal.** They are frequently skewed, double-humped (a **double distribution**, two separate areas of acceptance on one day) or stretched out on trend days. You will meet these shapes in Chapter 6. On such days the Value Area is not " $\pm 1\sigma$ " in any precise statistical sense; it is simply the high-volume core.
- 3. The real value is standardisation.** Because almost everyone uses 70%, VAH and VAL are comparable between traders. When someone says "yesterday's VAH", others know what is meant. That is communication value, not predictive value.

Value Area vs VWAP bands

Both ideas trace back to "one standard deviation", but they are built differently:

Feature	VWAP $\pm 1\sigma$ band (Chapter 4)	Value Area
Centre	VWAP (mean)	POC (mode)
Method	Variance formula	Counting row volumes outward from the POC
Symmetry	Always symmetric around VWAP	Usually asymmetric
Depends on	All prices, distance squared	Volume of rows adjacent to the growing area

Common mistake: "Seventy per cent is a magic, optimal number" is a popular claim, not proven. "The Value Area is always exactly one standard deviation" is only roughly true for a normal distribution. And changing the percentage until past results look better is **overfitting**: tuning a setting to fit history rather than to describe the market (see Chapter 12).

VAH and VAL: The Edges of Value

The **Value Area High (VAH)** is the highest price of the Value Area; the **Value Area Low (VAL)** is the lowest. Inside them, the market had its greatest agreement. Outside them are prices the market accepted less.




Figure 13. A real daily volume profile built from every trade: the POC (busiest price) and the value area (VAL to VAH, about 70% of the volume).

The edge asks a question

In auction logic, when price reaches the edge of value, a question is being asked: **will the market accept prices outside value, or reject them?** There are three descriptive scenarios:

1. **Rejection.** Price moves outside VAH or VAL, but little volume builds there and it returns quickly into the Value Area. The interpretation: prices outside value are not yet considered fair.
2. **Acceptance.** Price moves outside and **time and volume** build at the new prices. The interpretation: value may be shifting. This is the beginning of **value migration**, covered in Chapter 6.
3. **Probing or indecision.** Price reaches the edge and tests it again and again without a clear outcome.

Defining "acceptance" before you look

There is no universal, quantitative definition of acceptance. The Market Profile literature uses heuristics (rules of thumb) such as "at least two 30-minute periods trading outside the edge" or "volume building at the new prices". These are practical conventions, not laws. Whatever definition you use, write it down before you study a chart, so you cannot unconsciously change it to fit what you see.

Using previous edges

The VAH and VAL of a **previous**, completed profile (yesterday's, or last week's) are common reference levels for the current session, a theme Chapter 6 develops. But like every level, an edge is only a place to **pay close attention**. What completes the story is the order-flow behaviour at that price: whether aggressive orders are absorbed, whether delta shifts, whether new volume builds. Those tools come in Chapters 8 to 10.

Gold note: Using the hypothetical Value Area of 4203–4207 from the previous section, a trader watching XAUUSD would mark these edges at roughly 4173–4180, recalculating the gap on the day since it is never fixed.

Common mistake: "VAH means sell and VAL means buy, always." This is a popular claim, not proven. Two other common errors: treating the edges of a still-developing profile as fixed levels (see "Developing vs Final Profiles"), and forgetting that VAH and VAL depend on row size and calculation method.

Try it: Find two GC days where price went above the prior day's VAH. On one, did it return within roughly fifteen minutes with little volume built outside? On the other, did it stay and build a new high-volume area? Label each "rejection" or "acceptance" according to a definition you wrote down first.

HVN and LVN: Where Price Lingers vs Where It Rushes

Look at almost any profile and you will see peaks and valleys.

- A **High Volume Node (HVN)** is a local peak: a price area where noticeably more volume traded than at its neighbours.
- A **Low Volume Node (LVN)** is a local valley: an area where very little volume traded compared with its surroundings.

The POC is simply the largest HVN, but a profile can have several HVNs.

Why they form

- **HVNs** form where the market **rotated** for a while, with both sides trading actively. That is **acceptance**.
- **LVNs** form where the market **moved through quickly**, perhaps on an aggressive push or a news release, or where it **rejected** prices. Little time, little volume. Some traders call this an **inefficiency**: a price range the auction passed through without much two-sided trade.

A worked example: a two-day composite

A hypothetical profile combining two GC sessions, with \$1 rows:

Price	Volume	Note
4234	30	
4233	60	
4232	140	
4231	210	Upper HVN
4230	180	
4229	90	
4228	35	LVN zone
4227	20	LVN zone
4226	40	
4225	120	
4224	240	
4223	260	Lower HVN (POC)
4222	200	
4221	90	
4220	40	

The picture is two "hills" with a narrow valley between them. Perhaps the first day rotated around 4223, a news release then lifted price quickly through 4226–4229, and the second day rotated around 4231. The valley at about 4227–4228 is the LVN: the market barely traded there.




Figure 14. High-volume nodes (HVN) are prices where the market built acceptance and tends to slow down; low-volume nodes (LVN) are thin prices it tends to cross quickly.

Descriptive tendencies when price returns

Auction theory describes typical behaviour when price comes back to these areas:

- **Returning to an HVN:** price is often described as *slowing down and rotating*, because there was agreement there before.
- **Returning to an LVN:** two behaviours are commonly described. Either the LVN is rejected again, or, if price passes into it, it travels relatively **quickly** to the next HVN, because there is "nothing to get stuck on".

Both descriptions are a popular claim, not proven. They are descriptive tendencies from auction theory, not verified statistics. Treat them as questions to watch, not as rules.

Detecting HVNs and LVNs in practice

There is no single standard for "how high" or "how low" counts. Software typically uses a threshold such as a percentage of the POC's volume, or a local peak or valley with some minimum separation. Different settings produce different nodes. This is why it is more honest to think of HVNs and LVNs as **zones**, a few ticks or a dollar or two wide, rather than single-tick lines. A very small row size can also produce dozens of tiny "LVNs" that are just noise.

Gold note: In gold, LVNs often mark the path of fast moves around major US releases such as CPI, NFP and FOMC. HVNs often mark the quieter, rotational hours, for example the Asian session or the lull in the middle of the US session.

Common mistake: "Price always shoots through an LVN like lightning" and "an HVN is always support or resistance" are both popular claims, not proven. Finding dozens of LVNs with a one-tick row size is mistaking noise for structure.

Developing vs Final Profiles

While a session is still in progress, its profile is **developing**: every new trade adds volume to some row, so the **developing POC (dPOC)**, **developing VAH (dVAH)** and **developing VAL (dVAL)** can move. When the session ends, the profile becomes **final** (also called **frozen**), and its levels no longer change.

Watching the dPOC move

Here is a hypothetical GC RTH session, shown as snapshots of the cumulative profile:

Price	09:00	10:30	12:00	13:30 (final)
4207	0	10	30	45
4206	10	40	90	130
4205	30	100	170	210
4204	70	120	150	160
4203	90	110	120	125
4202	40	50	55	60
dPOC	4203	4204	4205	4205

Each column includes all volume from the session start up to that time, so every number can only grow. The dPOC steps from 4203 to 4204 to 4205 as the busiest price shifts higher during the morning. Many platforms draw the dPOC as a stepped line across the chart.

Why this matters

1. The dPOC path carries information. A dPOC that keeps stepping higher means the high-volume area, where value is being built, is moving up: higher prices are being accepted. A dPOC that stays put while price swings above and below it describes a market balanced around one fair price.

2. The look-ahead trap. When you review past days, you see the final profile. But at 09:30 on the day above, the final POC of 4205 **did not exist yet**; the dPOC was around 4203. If you look back and say "at 09:30 price bounced perfectly off the POC at 4205", you are using a level that only formed later, possibly *because* of the trading around it. This is **look-ahead bias**: using information that was not available at the moment you are analysing. It is one of the most common errors in trading education, and Chapter 12 deals with it in depth.

3. Today's levels are unstable. A dVAH at 09:00 can be several dollars away from the final VAH. That is why fixed reference levels usually come from **previous final** profiles (yesterday, last week), while today's developing profile is used to read *how value is being built*.

The rule this book follows: whenever a level is shown, say whether it was developing or final, and if final, whether it was known at the moment being discussed.

Common mistake: Drawing a fixed horizontal line at the morning's dVAH and expecting it to stay put. Naming the day's profile shape, or "today's value", before the session has ended is the same error in another form.

Try it: Use your platform's replay mode, or simply take screenshots, to capture a GC session's profile at 09:00, 10:30, 12:00 and 13:30 ET. Write down the dPOC at each time. Then look at the final profile and note any "perfect reactions" at the final POC that happened before that POC existed.

Session, Weekly, Monthly and Composite Profiles

A profile always belongs to a **time window**. Change the window and you change the profile. The common choices:

Session (daily) profile

One session. For gold, the same choice from Chapter 4 applies: the full **Globex** session (18:00–17:00 ET) or **RTH** (about 08:20–13:30 ET). Many Market Profile analysts keep the **overnight (ON)** profile separate from the **RTH** profile, so they can compare what US hours did with what happened in Asia and London.

Weekly and monthly profiles

A **weekly profile** includes every trade of the week, starting Sunday 18:00 ET for gold; weekly value is where the week's volume concentrated. A **monthly profile** covers a month and gives a longer view; its HVNs and LVNs describe the larger structure.

Composite profiles

A **composite profile** merges any sessions you choose: the last 5, 10 or 20 days, or "everything since the last FOMC". Jim Dalton uses composites to see **longer-term balance areas**. When the market rotates within a range for several days, the composite of those days often forms a clear bell-shaped distribution, even if no single day did.

Profiles are nested

Profiles fit inside one another like nesting dolls. A price that is an LVN on one day's profile may sit in the middle of an HVN on the weekly profile, because the market came back and traded there

heavily on other days. A day that looks lopsided on its own can be one piece of a perfectly balanced week. Shapes and levels only make sense **within their own window**.

So ask the question first, then choose the window:

Question	Suitable window
What is today's intraday context?	Session profile (Globex or RTH, labelled)
How did US hours differ from overnight?	Separate ON and RTH profiles
What is the multi-day balance area?	5–10 day composite, or weekly
What is the large structure?	Monthly, or a long composite

The roll warning from Chapter 4 applies here too. A monthly profile or a long composite can merge two contract months; on a continuous chart, check whether prices were back-adjusted and where the roll fell.

Common mistake: Comparing the Globex-session POC of one day with the RTH POC of another. Building a very long composite on continuous data without checking for rolls. And "monthly levels are always stronger than daily ones" is a popular claim, not proven.

Fixed Range and Visible Range Profiles

Besides calendar-based profiles (session, week, month), most platforms offer two "free-form" profiles.

Fixed Range Volume Profile (FRVP)

With a **Fixed Range Volume Profile**, you choose the **start and end points** by hand, and the profile covers only the trades in that window. Typical uses:

- the profile of a **specific move**, such as from the low before a CPI release to the high after it;
- the profile of a **balance area**, several days of sideways trading;
- a profile "from the FOMC statement until now".

This is the profile equivalent of anchored VWAP from Chapter 4: a window chosen around an event.

Visible Range Volume Profile (VRVP)

A **Visible Range Volume Profile** covers every candle **currently visible on your screen**. Zoom out or scroll, and the window changes, and so do the POC and Value Area.

Why visible range is a poor reference

A visible-range profile depends on something accidental: the size of your monitor and how far you happen to be zoomed in. Two people looking at the same chart can see two different POCs. That is

fine for a quick glance at where volume concentrated, but it cannot be reproduced, so it should not be used for a reference level you plan to rely on or present to others.

The hindsight trap in fixed range

Fixed range has the same weakness as anchored VWAP: the window can be chosen after seeing the outcome. If you redraw a fixed range three or four times until its POC lands exactly where price turned, you have proved nothing. Choose the window by a rule written beforehand (an event, a defined swing, a fixed number of days) and state the rule whenever you share the result.

Feature	Fixed range	Visible range
Who sets the window?	You, with start and end points	Your screen and zoom level
Reproducible?	Yes, if the rule is stated	No; changes as you scroll or zoom
Main risk	Hindsight in choosing the window	Accidental, unrepeatable levels
Good for	Profiling a move, a balance area, post-event trading	A quick look at where volume sits

Common mistake: Reporting visible-range levels without stating the window. Redrawing a fixed range repeatedly until it looks "perfect". And "a fixed range finds the hidden institutional level" is a popular claim, not proven.

Key idea: A volume profile is more information about where the market has done business. It is not a magic edge. In our own testing of order-flow information, signals on their own did not reliably predict direction once costs were included. Profiles earn their place as context, a map that tells you where to pay attention, combined with everything else in this book.

Try it: On a five-minute GC chart, draw a fixed range from the session low before the most recent NFP release to the end of that session, using the rule "from the last swing low before 08:30 ET". Then switch to a visible-range profile and zoom out twice. Write down the POC each time and notice how much the visible-range POC moves.

Chapter summary

- A volume profile is a histogram of volume at each price over a chosen period: volume at price rather than volume over time.
- Heavy volume at a price indicates two-sided agreement; light volume indicates fast travel or rejection.
- Row size, grid alignment and data source (tick versus candle approximation) all change the profile's shape and can move the POC.

- The POC is the most-traded row, the mode of the distribution; VWAP is the mean. On skewed days they separate, and their distance describes the skew. Treat the POC as a zone.
- The Value Area is built outward from the POC, adding the heavier side until at least 70% of the volume is included. The single-row and two-row methods can give slightly different VAH and VAL values.
- The 70% figure comes from Steidlmayer's Market Profile and the 68% of a normal distribution within $\pm 1\sigma$; it is a convention whose value is shared language, not prediction.
- VAH and VAL are the edges of value, where the market "asks" whether outside prices will be accepted or rejected. Define acceptance before you look.
- HVNs mark acceptance and rotation; LVNs mark fast travel or rejection. Behaviour on a return is descriptive, not proven; treat nodes as zones.
- Developing profiles change all session; final profiles are fixed. Using a final level as if it were known earlier is look-ahead bias.
- Session, weekly, monthly and composite profiles are nested windows; choose the window to fit the question, and watch for contract rolls.
- Fixed range profiles need a pre-written window rule; visible range profiles are not reproducible reference levels.

Checklist

- I can explain the difference between volume over time and volume at price.
- I know my profile's row size, data source and session window, and I label them.
- I can find the POC and explain why it is the mode, not the mean.
- I can build a Value Area by hand with both the single-row and two-row methods.
- I know that 70% is a convention and I do not tune it to fit past results.
- I have written a definition of "acceptance" before studying edges.
- I treat POC, HVNs and LVNs as zones, not single-tick lines.
- I label every level as developing or final, and I check whether it existed at the time I am analysing.
- I choose the profile window (session, ON/RTH, weekly, composite) based on the question I am asking.
- I do not use visible-range levels as references, and I choose fixed ranges by a written rule.
- I build profiles on futures volume and translate key levels to XAUUSD with the live, approximate gap.

Quiz

1. What does a volume profile show that ordinary volume bars under a chart do not?

2. Statistically, the POC is which measure of the distribution: the mean, the median or the mode? Why can it be far from VWAP on a skewed day?
3. A profile has a total volume of 1,200 contracts. The POC row (4205) has 260. The next rows above are 4206: 200 and 4207: 120; the next rows below are 4204: 180, 4203: 150 and 4202: 90. Using the single-row method, what is the Value Area and how much volume does it contain?
4. Reviewing yesterday's chart, you notice that price "perfectly bounced" off the final POC at 10:00 a.m. What is the main risk in drawing a conclusion from this?
5. Why is a visible range volume profile a poor choice for a reference level?

Quiz answers

1. It shows **how much volume traded at each price** over the chosen period, laid out on the price axis, rather than how much traded in each time period. It reveals where the market agreed (heavy volume) and where it moved quickly or rejected prices (light volume).
2. The POC is the **mode**: the single most-traded price. VWAP is the volume-weighted **mean**. On a skewed day, a long thin tail of trading pulls the mean towards it while the busiest price stays where the bulk of volume traded, so the two separate.
3. Target = 840. Start at 260; add 4206 (200 > 180) for 460; add 4204 (180 > 120) for 640; add 4203 (150 > 120) for 790; add 4207 (120 > 90) for **910**. Value Area = **4203–4207** with 910 contracts, about **75.8%** of the volume (roughly 4173–4180 on XAUUSD with an approximate \$27–30 gap).
4. **Look-ahead bias**. The final POC may not have existed at 10:00 a.m.; the developing POC could have been elsewhere, and the final POC may have formed there partly because of the trading around that time. You are using information that was not available at the moment.
5. Because it **changes whenever you zoom or scroll**. Its window is set by your screen and zoom level, not by a rule, so two people looking at the same chart can see different POCs and Value Areas. It is useful for a quick glance, not for a reproducible level.

6

Volume Profile II: Shapes, Yesterday's Map and Value Migration

In this chapter: In Chapter 5 you learned how a volume profile is built and how to find its point of control (POC), value area (VAH and VAL), and high- and low-volume nodes. This chapter teaches you to *read* a finished profile. You will learn the common profile shapes (D, P, b, thin trend and double distribution) and the auction story each one is usually said to tell, with the limits of those stories spelled out. You will learn how yesterday's final profile becomes today's map, what "naked" levels are and why the claim that price "always returns" to them is folklore, how to classify value migration from one day to the next, how to describe where gold opens relative to yesterday's value (including an honest look at the famous "80% rule"), how volume profile differs from the older TPO (Market Profile), and finally the setup and reading mistakes that trip up most beginners, ending with a pre-session checklist for gold.

A quick word on prices before we begin. Every price in this chapter is an illustrative gold futures (GC) price. If you trade gold through a CFD broker, the instrument you see is usually called XAUUSD, and its price sits below the futures price by a gap that has recently been roughly \$27–30 (the gap changes with interest rates, time to contract expiry and your broker's own pricing). Where a price matters, we give an approximate XAUUSD equivalent beside it. Treat those equivalents as rough, never exact: always measure the live gap on your own platform.

The Balanced Day: The D-Shaped (Normal) Profile

What it looks like

Picture a finished daily profile turned on its side so the price axis runs vertically. On a balanced day, the horizontal bars are longest in the middle of the day's range and get gradually shorter towards the top and the bottom. The outline looks like the capital letter **D**, or like a bell curve standing on its edge. This is called a **D-shaped profile**, and in Market Profile literature it is also called a **normal profile** (the name comes from the "normal" bell-shaped distribution in statistics).

Three features usually go together:

- **One main bulge.** There is a single high-volume area, roughly in the middle of the range.
- **The POC sits near the middle.** The point of control (the single price with the most volume) is close to the midpoint of the day's high and low.
- **A roughly symmetric value area.** The value area (the band holding about 70% of the day's volume) extends a similar distance above and below the POC.




Figure 15. Common profile shapes and the auction story each one tells.

The auction story

A profile is a record of where business was done. A D-shape tells a story of a **two-sided auction**. During the session, price moved up, but at higher prices it did not find enough new buyers to keep going, so it rotated back down. It moved down, but at lower prices it did not find enough new sellers, so it rotated back up. After several of these swings, most of the trading piled up around one central price that both sides were willing to accept.

In the language of auction theory (which Chapter 7 covers in full), this condition is called **balance**. Neither side had dominant control, and the market found a price that was "fair" for that window of time. The important words are *for that window of time*. Balance describes the past session; it is not a promise about tomorrow.

Think of a pendulum swinging between the edges of value, losing energy at each edge because there is less willingness to trade out there, and spending most of its time in the thick middle.

Companion clues

D-days rarely appear on their own; other indicators usually tell the same story:

- **The daily VWAP is flat.** VWAP (the volume-weighted average price from Chapter 4) barely slopes, and it tends to sit close to the POC. When the average price paid all day is close to the most-traded price, nobody pushed price far away for long.
- **Price crosses VWAP and the POC many times.** On a 5-minute chart you will see price wander back and forth across both levels.
- **The range is often modest.** A D-day's high-to-low range is usually smaller than on days with major economic news.

Example: on a quiet Monday, GC trades between 4,012.0 and 4,031.0 (XAUUSD roughly 3,984–4,003), the POC forms at 4,021.5 (XAUUSD $\approx$ 3,993.5), the value area runs from 4,016.0 to 4,026.5, and VWAP finishes at 4,021.8. Value is almost centred and VWAP sits on the POC: a textbook D.

What a D-shape does not tell you

A D-shape does not say that tomorrow will also be balanced. Balance always ends eventually, usually because something new arrives (a data release, a central-bank surprise, a geopolitical headline) and pushes one side into control. That shift from balance to directional movement is called **imbalance**.

There is a popular view in auction theory that long periods of balance tend to be followed by directional moves ("balance leads to imbalance"). As a general description of how markets alternate, that is reasonable. But the profile does not tell you *when* the break will come or in *which direction*. Two related claims circulate widely:

- "After a D-day, a trend day always follows." This is a popular claim, not proven.
- "On a D-day, always fade the edges" (that is, always trade against moves to the value-area edges). This is also a popular claim, not proven. It describes what balanced days look like in hindsight, not a rule that holds in advance.

The time-frame warning

A profile's shape depends on the window it covers. The same day can look like a D on a profile built only from regular trading hours (RTH, roughly 08:20–13:30 New York time for gold) and look quite different on a profile built from the full 23-hour Globex session. Several D-days can also merge into one larger D on a multi-day composite (Chapter 5). Whenever you name a shape, name the window too: "RTH profile, D-shaped", not just "it's a D".

Common mistake: Naming the shape before the session has ended. A developing profile at 10:00 a.m. can still turn into a P, a b or a trend profile by the close. Shape names belong to finished profiles.

Key idea: A D-shaped profile describes a two-sided, rotational auction in which the market found a temporary fair price. It tells you what happened, not what will happen next.

P and b Profiles: Short-Covering and Long-Liquidation Stories

The two shapes

A **P-shaped profile** has its volume bulge near the **top** of the day's range and a thin "tail" of low volume below it, so the outline resembles the letter P. Reading it as an auction record: in the lower part of the range, price moved up **quickly, on little volume**, and then, near the top, it **slowed down and built volume**.

A **b-shaped profile** is the mirror image: the bulge is near the **bottom** of the range, with a thin tail above. Price fell quickly through the upper part of the range and then built volume near the lows.

The thin part of each shape is the fast part of the day. The thick part is where the market slowed down and accepted prices (did lots of business there).

The classic interpretation

Market Profile literature, especially the books of James Dalton and his co-authors, offers a well-known reading of these shapes:

- A **P-shape** is often linked to **short covering**. Traders who had sold short earlier (bet on a fall) buy back to close their positions. Their buying lifts price quickly. Once they have finished covering, there are not many *new* buyers arriving, so the market stalls and rotates near the top.
- A **b-shape** is often linked to **long liquidation**. Traders who had bought earlier sell to close their positions. Their selling pushes price down quickly. Once the liquidation is done, there are not many *new* sellers, so the market rotates near the lows.

The logic is about fuel. A move driven by *old positions being closed* may have less fuel to continue than a move driven by *new money entering* (what auction theory calls *initiative* activity). When the closing is done, the push is done.

Honest limits

This story is plausible, but you should hold it loosely, for three reasons.

1. **A volume profile cannot see whether trades opened or closed positions.** A contract bought to close a short looks exactly the same in volume as a contract bought to open a new long. To get even a rough idea, you would need to look at changes in **open interest** (the total number of futures contracts outstanding, which the exchange publishes once a day). Even then, open interest is a blunt, end-of-day number and cannot tell you what happened in a particular hour.
2. **P and b can simply be continuation.** In a steady uptrend, a P can mean nothing more than "price moved up, then consolidated at its new level", which is exactly what healthy trends do.
3. **Shape depends on the window.** As with D-days, an RTH P-shape may look different on a Globex profile or a multi-day composite.

So treat P and b as a **hypothesis** about the day's story, then test that hypothesis with order flow: did delta and cumulative delta (Chapter 3) confirm aggressive buying during the fast leg? Did heavy selling get absorbed in the bulge? Later chapters give you the tools to ask those questions.

A worked gold example

Imagine an 08:30 New York data release that comes in weaker than expected. GC jumps from 4,018.0 to 4,034.0 (XAUUSD roughly 3,990 to 4,006) within fifteen minutes, with only a few contracts traded at each price on the way up. For the rest of the session, price rotates between 4,030.0 and 4,040.0 and builds heavy volume there, with the POC at 4,035.5 (XAUUSD $\approx$ 4,007.5). The finished RTH profile is a P: a thin tail from 4,018 to 4,030 and a fat bulge above.

Now ask the honest questions. Was the fast move covering or new buying? The profile cannot answer. If cumulative delta rose strongly during the jump and then flattened in the bulge, that fits "aggressive buyers lifted price, then two-sided trade took over", which suits *both* the covering story and a continuation story. On some days like this the move continues; on others it fades. The shape alone does not decide.

Common mistake: "P means bullish, b means bearish." This is a popular claim, not proven, and in some material you will even see it stated the other way round (P as a sign of exhausted buying). Neither rule holds reliably. Likewise, "a P is always short covering" is a popular claim, not proven.

Key idea: The thin part of a profile shows where price moved fast; the thick part shows where it was accepted. Who was trading, and whether they were opening or closing positions, is something the profile cannot see.

Trend Days and Double Distributions

The thin, elongated trend profile

On some days the profile has no dominant bulge at all. Instead it is **long and thin**: volume is spread thinly across a wide range of prices, with no single area where the market settled. This is the **thin** or **elongated** profile, typical of a **trend day**.

The auction story is a **one-sided auction**. The market kept probing new prices in one direction, and at each new price it failed to find enough opposite-side business to stop it. No fair price was established during the day, because the price kept moving on before it could become accepted.

Companion clues:

- **VWAP slopes steadily** in the direction of the move.
- **Price spends most of the day on one side of VWAP**, sometimes riding the outer VWAP bands (Chapter 4).
- **The value area is wide and not very meaningful**. On a trend day, "the band containing 70% of volume" is mostly a stretch of price that was traded through, not a zone of agreement.

Because nothing was accepted, tools built on "returning to the average" tend to struggle on trend days. That is a structural observation, not a statistic: if the market never returns to its average, a method that waits for a return keeps waiting.

The double distribution

A **double distribution** is a profile with **two separate bulges** (two high-volume nodes, HVNs) joined by a thin neck of low volume (a low-volume node, LVN). The story: the market was balanced in one area, then something knocked it out of balance (on gold, often an 08:30 New York data release

such as CPI), price moved quickly to a new area, and there it found a **new balance**. One day, two small auctions.

Each of the two distributions has its own POC. The POC of the *whole day* may simply be the POC of the larger one, which hides the other completely. If you report only the day's POC and value area, you might draw a value area that sits right across the thin neck, an area where very little trading actually happened. That is misleading.

Many charting platforms offer a **split profile** function, which lets you break a session into separate profiles so each distribution gets its own POC and value area. On a double-distribution day, splitting the profile usually gives a much more honest picture.

The LVN in the middle

In profile literature, the thin neck between two distributions is the boundary the market crossed in a hurry. When price comes back to it on later days, traders often watch it as a **decision area**: either price is turned away and stays in its current distribution, or it passes through and travels back towards the other one. That is a popular claim, not proven. It is a descriptive tendency some traders report, not a rule.

Worked example

Suppose, before an 08:30 CPI release, GC balances between 4,040.0 and 4,048.0 with a local POC at 4,044.0. The number comes in hot, and in four minutes GC falls to 4,022.0, passing 4,026–4,038 on very light volume. For the rest of the session it balances between 4,015.0 and 4,026.0 with a local POC at 4,020.5. The whole-day POC lands at 4,020.5, because the afternoon distribution was bigger, and the morning POC at 4,044.0 (XAUUSD $\approx$ 4,016) is invisible unless you split the profile. The neck at roughly 4,026–4,038 (XAUUSD $\approx$ 3,998–4,010) is the LVN between the two auctions.

Shape	Outline	Auction story	Typical companions
D (normal)	One central bulge	Two-sided rotation; balance	Flat VWAP near POC
P	Bulge at top, thin tail below	Fast rise, then acceptance higher	Interpretation: covering or continuation
b	Bulge at bottom, thin tail above	Fast fall, then acceptance lower	Interpretation: liquidation or continuation
Thin trend	Long, narrow, no bulge	One-sided auction; nothing accepted	Sloped VWAP, price on one side
Double distribution	Two bulges, LVN between	Balance $\rightarrow$ imbalance $\rightarrow$ new balance	Split the profile

Common mistake: Fading (trading against) a trend day just because "price is outside the value area". On a trend day, the value area means little and price outside it is normal.

Common mistake: Assuming the LVN between two distributions "must be filled". That is a popular claim, not proven.

Key idea: A thin profile means no agreement was reached; a double distribution means the day contained two separate agreements. Split the profile so each gets its own POC and value area.

Yesterday's Profile as Today's Map

Why use yesterday?

Today's profile is still being built. Its POC and value area move as the day goes on (Chapter 5 called this the *developing* profile). If you use today's developing levels as references, you are aiming at a moving target. Yesterday's **final** profile, by contrast, is fixed: it will not change. That is why it is the most common reference map for the current session.

The standard levels

Before the session, traders usually mark these levels from the prior session's final profile:

- **pPOC (prior POC):** yesterday's most-traded price.
- **pVAH and pVAL:** the upper and lower edges of yesterday's value area.
- **pHigh and pLow:** yesterday's high and low, the edges of yesterday's range.
- **pVWAP:** yesterday's final VWAP (optional but useful).
- **The most prominent HVNs and LVNs** from yesterday's profile.



Figure 16. Yesterday's value area and POC become today's reference map: watch how price behaves when it returns to them.

Why these levels matter

These levels summarise *yesterday's agreement*: where the market did most of its business and where it stopped. Today the market must, in effect, answer one question: **Is yesterday's value still valid, or is the market looking for new value?**

- If price keeps rotating inside yesterday's value area, the market is still accepting yesterday's prices.
- If price moves outside yesterday's value area and is accepted there (it spends time and builds volume), the market is looking for new value.

Those two possibilities are exactly what the next sections (value migration and the open relative to value) describe in more detail.

What does "yesterday" mean for gold?

Gold trades almost 23 hours a day, so "yesterday" needs a precise definition, and it must match the session setting you chose for your profile (Chapter 4 discussed session settings for VWAP; the same logic applies here):

- **If you use Globex sessions**, "yesterday" means the session that ran from 18:00 New York time the evening before last until 17:00 New York time yesterday.
- **If you use RTH sessions**, "yesterday" means roughly 08:20–13:30 New York time yesterday. The overnight hours then form a **separate** overnight profile, which supplies its own reference levels: the **ONH** (overnight high) and **ONL** (overnight low).

Do not mix the two. A Globex profile from yesterday combined with RTH logic today produces levels that mean nothing consistent.

A simple pre-session routine

1. Load yesterday's final profile (in the session type you always use).
2. Draw pPOC, pVAH, pVAL, pHigh and pLow, extending them across today's chart.
3. Note where current price is: inside yesterday's value area, outside value but inside yesterday's range, or outside yesterday's range.
4. Add weekly levels (Chapter 5) as a second layer, if they are nearby.

Worked example: yesterday's RTH profile had pHigh 4,052.0, pVAH 4,045.5, pPOC 4,041.0, pVAL 4,036.5 and pLow 4,029.0. At 08:00 New York time this morning, GC trades at 4,048.0. The table below sets out the map, with approximate XAUUSD equivalents assuming a gap of about \$28.

Level	GC	XAUUSD (approx.)
pHigh	4,052.0	≈ 4,024
pVAH	4,045.5	≈ 4,017.5
pPOC	4,041.0	≈ 4,013
pVAL	4,036.5	≈ 4,008.5
pLow	4,029.0	≈ 4,001
Current	4,048.0	≈ 4,020

Current price sits above yesterday's value area but below yesterday's high: outside value, inside range. That single observation already frames the morning: will the market accept these higher prices, or rotate back into yesterday's value?

This is a **map**, not a forecast. It tells you where to pay close attention, not what will happen there.

Common mistake: Drawing too many levels. With dozens of lines on the chart, every wiggle in price can be "explained" by some line, and nothing can ever be tested.

Common mistake: "Today's price will definitely test the pPOC." This is a popular claim, not proven.

Try it: For five consecutive days, before the session begins, draw only the five standard levels from the prior day's final profile, and write one sentence about where price is relative to them. After the close, note which levels price reached and what happened there. Do not add any other lines for those five days.

Naked (Virgin) POC, VAH/VAL and VWAP

Definitions

A **naked POC** (also called a **virgin POC**, often abbreviated nPOC) is the final POC of a past session that price has **not traded at again** since that session ended. "Naked" simply means "not yet revisited". The same idea applies to other levels:

- **Naked VAH / naked VAL:** a past session's value-area edge that has not been revisited.
- **Naked VWAP:** a past session's final VWAP that has not been revisited.

When price later trades at the level, it is no longer naked; traders say it has been "filled" or "covered".



Figure 17. A naked (virgin) POC: the busiest price of Oct 02 that price has not traded back to since. Traders watch it as a magnet and a reaction level; it is a reference, not a promise.

The popular story

The usual narrative goes like this: the POC was where the most business was done; if the market left that price quickly and never came back, there is "unfinished business" there; therefore price tends to return and test it. You will often hear this stated as a law: "price always returns to the naked POC."

That statement is a popular claim, not proven. It is folklore, and it has some serious problems.

A reality check

1. A random walk also comes back. Imagine a price series that moves completely at random, with no memory and no magnets. Such a series still revisits nearby prices very often, sooner or later. So

when a naked POC that was \$3 away eventually gets touched, that alone proves nothing about magnetism. The right question is: *are naked POCs touched more often than random levels at the same distance and age?* Answering that requires a proper statistical test comparing naked levels with randomly placed control levels. We have not presented such a test here, and most material you will read on the topic has not either.

2. Selection bias. Naked POCs that do get filled make memorable chart examples. Naked POCs that stay untouched for months (which is common during strong gold trends) quietly disappear from the conversation. If you only ever see the successes, the idea will always look better than it is.

3. Touch is not reaction. Even when price does reach a naked level, nothing special may happen there. Price can slide straight through it. "It was touched" and "it mattered" are two different claims.

Honest use

Keep naked levels on your map as **reference levels**: places where, *if* price arrives, you will watch the order flow carefully. Always write down two facts next to each one:

- **Its distance** from current price.
- **Its age** (how many sessions ago it was formed).

For example: "nPOC from four sessions ago, 4,064.0 (XAUUSD $\approx$ 4,036), \$12 above current price." A two-week-old naked POC \$60 away is a very different object from a one-day-old naked POC \$4 away, and writing the numbers forces you to notice that.

Keep the list short. A trader who keeps dozens of old naked levels on the chart will find that *every* move ends near one of them, which is a sign the method has stopped measuring anything.

Common mistake: "A naked POC is always filled." A popular claim, not proven. "A naked POC is the strongest support or resistance." Also a popular claim, not proven.

Key idea: Naked levels are worth keeping on a map, but their famous "magnet" effect has not been shown to beat simple chance. Write down distance and age, and judge the level by what order flow does when price arrives, not by the fact that it arrived.

Value Migration: Is Value Moving Higher, Lower or Nowhere?

The idea

Price moves around all day. **Value**, the area where most trading actually happened, moves more slowly and more meaningfully. **Value migration** means comparing the position of one session's value area with the previous session's. James Dalton treats this comparison as one of the most important pieces of evidence about direction in Market Profile analysis: price can spike up and down, but the relocation of value shows where business was really being done.




Figure 18. Value migration: when each day's value area and POC step higher or lower, the market is accepting new prices; overlapping boxes mean balance.

The categories

Call today's value area VA(today) and yesterday's VA(yesterday). The standard categories are:

Category	Definition	Description
Higher value	Today's VAL is above yesterday's VAH (no overlap)	Clear acceptance of higher prices
Lower value	Today's VAH is below yesterday's VAL (no overlap)	Clear acceptance of lower prices
Overlapping-higher	Areas overlap, but both of today's edges are higher	Leaning higher, with less conviction
Overlapping-lower	Areas overlap, but both of today's edges are lower	Leaning lower, with less conviction
Inside value	Today's VA lies completely inside yesterday's	Contraction, balance, waiting
Outside value	Today's VA completely contains yesterday's	Expansion; both sides active; often volatile indecision
Unchanged	Essentially the same area	No change in where business is done

The no-overlap categories (higher and lower value) carry the strongest descriptive message. The overlapping categories show a lean with less conviction. Inside days show contraction, which often

accompanies waiting before an important event. Outside days show expansion: both sides were active and the market explored widely without settling.

Worked numbers

Yesterday: VAL 4,030.0, VAH 4,042.0.

- Today: VAL 4,044.5, VAH 4,056.0. Today's VAL is above yesterday's VAH. **Higher value.**
- Today: VAL 4,035.0, VAH 4,047.0. They overlap; both edges are higher. **Overlapping-higher.**
- Today: VAL 4,032.5, VAH 4,039.5. Entirely inside yesterday's area. **Inside value.**
- Today: VAL 4,024.0, VAH 4,050.0. Contains yesterday's area entirely. **Outside value.**

(XAUUSD equivalents simply shift every number down by roughly \$28; the *relationships* between the areas are what matter, and they are identical on both instruments.)

Price is not value

Here is why this matters. Imagine a day on which gold makes a **new high**, a dollar above yesterday's high, yet the day's value area is **overlapping-lower**. Price poked higher briefly, but most of the trading happened at lower prices than yesterday. The headline says "gold hits new high"; the auction record says "the market did not accept higher prices." Compare that with a day that has no new high at all but shows **higher value**: the market moved its whole centre of business upward.

Limits

1. **Sensitivity to settings.** The category depends on your session definition (Globex or RTH) and your value-area calculation method. Change either and a day can change category.
2. **"Unchanged" needs a threshold.** How close counts as "essentially the same"? Any answer is a convention, and you must choose one and keep it.
3. **Migration in progress is unstable.** During the day you can only see the *developing* value area, which keeps moving until the close. A classification made at 10:00 a.m. is provisional.
4. **Description, not destiny.** Value migration describes what happened. It does not guarantee that the migration will continue; migrations can fail, and value can return to where it was.

Common mistake: "Higher value today means higher tomorrow." A popular claim, not proven.

Common mistake: Comparing value areas built from different sessions, for example a Globex area for one day against an RTH area for the next.

Key idea: Price shows where the market *visited*; value shows where it *did business*. Classifying value migration day by day gives you a disciplined description of direction that a single new high or low cannot.

Where Did Gold Open Relative to Value?

Three opening situations

In *Mind Over Markets*, Dalton classifies the **open** relative to the previous day's value and range into three situations:

1. **Open within value** (inside yesterday's value area). The market is, for now, agreeing with yesterday's value. The context is more balanced, and rotational trading is more likely than on the other two types.
2. **Open outside value but within range** (above the VAH or below the VAL, but still inside yesterday's high-low range). A mild imbalance. The key question is whether the market will accept prices outside value or return into it.
3. **Open outside range** (above yesterday's high or below yesterday's low). A stronger imbalance. There is more potential for directional movement, or for a quick failure and return if the outside prices are not accepted.

Using our earlier map (pHigh 4,052.0, pVAH 4,045.5, pVAL 4,036.5, pLow 4,029.0), an open at 4,040.0 is within value; an open at 4,049.0 is outside value but within range; an open at 4,056.0 (XAUUSD $\approx$ 4,028) is outside range.

Gold's problem: which open?

GC trades around 23 hours a day. Its only true "opening" in the technical sense comes after the daily halt, at 18:00 New York time. But an 18:00 open happens in the quietest part of the day and says little about how the main session will develop.

Profile analysts on gold therefore usually use the **RTH open**, around **08:20 New York time** (the old COMEX trading-floor opening time), and compare it with **yesterday's RTH value and range**, while treating the overnight hours separately through ONH and ONL. Whichever convention you choose, state it explicitly and use it consistently. Chapter 7 returns to this question in detail.

The "80% rule", honestly

You will meet the **80% rule** sooner or later. It is usually stated like this: *if the market opens outside yesterday's value area, then moves back inside it and stays inside for two consecutive 30-minute periods, there is an 80% probability that it will travel across the whole value area to the opposite edge.*

What you should know about it:

- **Its origin.** It comes from Market Profile commentary of the 1980s and 1990s associated with the Dalton school, and it was originally defined on the **TPO value area** (built from time, explained later in this chapter), **not** on the volume value area.
- **Its status.** The 80% figure is a historical claim. It is a popular claim, not proven. Some informal public tests on stock-index futures have reported success rates well below 80%, and we are not aware of any rigorous test on gold futures.

- **How we treat it in this book.** As a historical idea to test, never as a statistic to rely on.

To examine it, define every element before looking at any chart (session, value-area method, what counts as "back inside" and "reaching the other side") and compare the result with a sensible baseline. Chapter 12 explains how.

Honest use

The position of the open relative to value **frames the day's scenarios**. It tells you which questions are most relevant this morning. It does not tell you which way price will go.

Common mistake: "The 80% rule means an 80% win rate." A popular claim, not proven, and in any case a probability of price *reaching a level* is not the same thing as the outcome of a trade, which also depends on costs, timing and where you would have been wrong.

Common mistake: Applying the 80% rule to a volume value area with arbitrary time periods and expecting the historical claim to carry over. Or using the 18:00 Globex open with RTH logic.

Volume Profile vs TPO (Market Profile)

Two ways of counting

The **Market Profile**, developed by J. Peter Steidlmayer at the Chicago Board of Trade in the 1980s, did not originally count volume at all. It counted **time**.

Its unit is the **TPO**, short for **Time Price Opportunity**. The session is divided into fixed periods, usually **30 minutes** long, and each period is given a letter (A for the first, B for the second, and so on). Whenever price trades at a given level during a period, that period's letter is printed at that price. Stack all the letters against the price axis and you get a histogram, just like a volume profile, except that the length of each row means **how many periods the market spent at that price**, not how many contracts traded there.

A **volume profile** builds the same kind of histogram using **volume** instead.

Key differences

- **Different POCs.** The **TPO POC** is the price where the market spent the most *time*. The **volume POC** (sometimes written VPOC) is the price with the most *volume*. They are often close, but not always. If a news release produces an enormous burst of trading at one price within a few minutes, the volume POC may land there, while the TPO POC, which only counts how many 30-minute periods visited each price, may sit somewhere else entirely.
- **Same 70% logic, different unit.** Both methods build a value area around their POC containing about 70% of the total, but one counts TPOs and the other counts contracts.

- **Time-based concepts.** Some important Market Profile ideas are defined by time, and they are not directly visible on a volume profile:
- **Initial Balance (IB):** the range of the first hour, i.e. the first two 30-minute periods.
- **Single prints:** prices where only one period's letter appears, meaning price moved through them quickly.
- **Tails:** single prints at the very top or bottom of the day's profile.

The nearest volume-profile equivalent of a run of single prints is an **LVN**. Chapter 7 develops all of these concepts.

Which is better?

Neither. A TPO profile shows **acceptance in time**; a volume profile shows **acceptance in volume**. Many analysts look at both side by side. Because gold futures give us genuine, exchange-reported volume, this book leads with the volume profile. But you need the TPO vocabulary to understand Dalton's literature, much of which is written in time-based terms.

A small example. Suppose gold spends five of the session's eleven 30-minute periods trading around 4,030.0, but at 08:30 an inflation release causes a burst of 6,000 contracts within three minutes at 4,041.0, after which price drifts back. The TPO POC may be 4,030.0 (most time), while the volume POC may be 4,041.0 (most contracts). Both are "correct"; they answer different questions.

Common mistake: "Market Profile and Volume Profile are the same thing." One measures time, the other volume.

Common mistake: Applying TPO-based rules (like the 80% rule) to a volume profile without checking whether they still hold. And the claim that one of the two methods has proven superiority is a popular claim, not proven.

Common Volume Profile Mistakes and a Pre-Session Checklist for Gold

This final section collects the errors that cause the most confusion, then turns everything from Chapters 5 and 6 into a short routine.

Configuration mistakes

1. **An unspecified session.** Is your profile built on Globex or RTH? Every level (POC, VAH, VAL, high, low) changes depending on the answer. If you do not know which session your platform uses, find out before you use any level.
2. **Inconsistent row size.** The row size (how many ticks are grouped into each horizontal bar) affects where the POC and value area land. If you change it from day to day, you cannot

compare days. Pick one and keep it.

3. **Bar data instead of tick data, without knowing it.** Some platforms build profiles from candles rather than individual trades, spreading each candle's volume across its range. That is an approximation. Know which one your platform uses.
4. **Ignoring the contract roll.** Gold futures expire, and traders move ("roll") from one contract month to the next. Composite and monthly profiles that span a roll can mix two different contracts with different prices. Check how your platform handles this.
5. **Using a visible-range profile as a reference.** A profile that covers "whatever is on the screen" changes every time you scroll or zoom. It cannot serve as a fixed reference level.

Reading mistakes

1. **Look-ahead.** Using a *final* profile to judge decisions that would have been made *during* the day, when only the developing profile existed (Chapter 5). This is the most common error in educational chart examples, and it makes every level look far better than it really was.
2. **Treating every level as a wall.** Profile levels are places to pay attention, not barriers that price must respect.
3. **Treating shape interpretation as law.** "P means bullish" is an interpretation, not a rule.
4. **Repeating folklore without a label.** "Naked POCs always fill", "the 80% rule", "LVNs are always crossed quickly": each of these is a popular claim, not proven.
5. **Too many levels.** With thirty lines on the chart, any reaction has an explanation and nothing is ever tested. Fewer lines give you a clearer, more honest read.

Two more claims deserve special mention. "Volume profile shows where institutions bought" is a popular claim, not proven: the profile shows where volume traded, not who traded it. "Volume profile is a trading edge" is also a popular claim, not proven. It gives you more information about the auction, which is valuable but is not profit; in our own testing, order-flow signals on their own did not reliably predict direction once trading costs were included.

A pre-session checklist for gold

Use this before the session. Fill in the blanks the same way every day.

- **Session:** ____ (Globex 18:00–17:00 New York time, or RTH about 08:20–13:30 New York time). Kept constant.
- **Row size:** ____ ; **data:** tick or bar-based.
- **Yesterday's final profile:** pPOC ____ / pVAH ____ / pVAL ____ / pHigh ____ / pLow ____ (with XAUUSD equivalents if you trade the CFD).
- **Overnight (if you use RTH):** ONH ____ / ONL ____.
- **Value migration,** yesterday versus the day before: ____ (higher, lower, overlapping-higher, overlapping-lower, inside, outside, unchanged).
- **Weekly and naked levels nearby:** at most two or three, each with its distance and age.
- **VWAP:** daily and weekly, with bands.

- **News calendar for gold** (CPI, the jobs report, Federal Reserve decisions): exact times in New York time.
- **Where current price sits:** inside value / outside value but inside range / outside range.

This checklist builds a **map**. What happens at each level, whether price is accepted or rejected, whether aggressive buyers or sellers are absorbed, has to be read with order flow, which is the subject of the chapters ahead.

Try it: Run the checklist on a real gold morning, starting from an empty chart. Stop when you have six to eight labelled levels at most. Take a screenshot, then take another after the close, and compare. Do this for two weeks before you add anything else to your routine.

Gold note: Gold's most important scheduled U.S. data usually arrives at 08:30 New York time, ten minutes after the RTH open. A profile built from the RTH session therefore often contains a news spike in its first hour. Keep that in mind when you read shapes and value areas: a single burst of news-driven volume can move the POC and change the shape of the whole day.

Chapter summary

- A **D-shaped (normal) profile** has one central bulge, a POC near the middle and a roughly symmetric value area. It describes a balanced, two-sided auction, usually with a flat VWAP near the POC. It says nothing about tomorrow.
- **P and b profiles** show a fast, thin leg and a fat area of acceptance. They are often read as short covering (P) and long liquidation (b), but the profile cannot see whether positions were opened or closed, and either shape can simply be continuation.
- A **thin trend profile** shows a one-sided auction with no accepted price. A **double distribution** shows two separate balances joined by an LVN; split the profile so each has its own POC and value area.
- **Yesterday's final profile** is today's fixed map: pPOC, pVAH, pVAL, pHigh, pLow, plus optional pVWAP and key nodes. "Yesterday" must match your session setting.
- **Naked levels** are past levels that have not been revisited. The claim that price always returns to them is folklore; random price paths also revisit nearby levels, and selection bias inflates the examples you see.
- **Value migration** compares today's value area with yesterday's: higher, lower, overlapping-higher, overlapping-lower, inside, outside or unchanged. Price and value can tell different stories.
- **The open relative to value** (inside value, outside value but inside range, outside range) frames the day's scenarios. The "80% rule" is a historical TPO-based claim, not a proven statistic.
- **TPO profiles count time; volume profiles count contracts.** Their POCs can differ. TPO gives us concepts such as Initial Balance, single prints and tails.

- Consistent settings, no look-ahead, labelled folklore and a short list of levels are the foundation of honest profile reading.

Checklist

- I always say which session (Globex or RTH) and which window a profile covers before naming its shape.
- I name shapes only on finished profiles, never on developing ones.
- I treat P and b stories as hypotheses to check with order flow.
- On double-distribution days, I split the profile.
- Before each session I draw pPOC, pVAH, pVAL, pHigh and pLow from yesterday's final profile, with XAUUSD equivalents if I trade the CFD.
- I note every naked level's distance and age, and keep no more than two or three.
- I classify value migration with the same session and value-area method every day.
- I write down where gold opened relative to yesterday's value and range, using the 08:20 New York RTH open (or my own stated convention).
- I never repeat the 80% rule, naked-POC magnetism or LVN-fill claims without labelling them as unproven.
- I keep my chart to six to eight labelled levels.

Quiz

1. A finished RTH profile has one central bulge, a POC near the middle of the range, and a flat VWAP sitting close to the POC. What kind of auction does this describe?
2. Why can a P-shaped profile not prove that short covering took place?
3. Today's VAL is above yesterday's VAH. Which value-migration category is this, and what does it describe?
4. A naked POC \$2 away from current price is touched later in the session. What does that prove about naked POCs, and why?
5. How do the TPO POC and the volume POC differ, and why might they land at different prices on a news day?

Quiz answers

1. A balanced, two-sided auction. Price rotated around a temporary fair price, with neither side in control during that session. It does not predict tomorrow's behaviour.
2. Because a volume profile records only how much traded at each price. It cannot tell whether a trade opened a new position or closed an existing one. At best, daily changes in open interest give a rough, end-of-day hint. A P can also simply reflect continuation in an uptrend.

3. Higher value. There is no overlap, and the whole area of business has moved above yesterday's, which is a clear description of acceptance at higher prices. It is a description, not a guarantee of further migration.
4. Very little on its own. Even a random price path frequently revisits nearby levels, so a touch at a short distance is expected by chance. To show a real effect, you would need to compare touch rates with random control levels at the same distance and age. And a touch is not the same as a meaningful reaction.
5. The TPO POC is the price where the market spent the most time (the most 30-minute periods); the volume POC is the price where the most contracts traded. A news release can concentrate heavy volume at one price in a few minutes. That moves the volume POC there, while the time-based TPO POC stays where the market spent most of the session.

PART III • THE AUCTION

7

Auction Market Theory

In this chapter: Auction Market Theory (AMT) is the framework that ties your volume measurements together into one picture of what the market is *doing*. You will learn why markets exist, how to tell balance from imbalance, how to judge acceptance and rejection with time and volume, how to read a TPO chart, how to define gold's Initial Balance and read range extension, the four opening types, the open's position relative to yesterday, the main day types, excess, poor highs and lows, single prints, gold's session map and reference levels, ending with a pre-New York checklist. AMT describes context; it does not predict prices.

All times in this chapter are New York time (ET) unless marked "UK". Prices are illustrative gold futures (GC) prices; where a price matters, an approximate XAUUSD (CFD) equivalent is given using a gap of about \$28. The real gap moves around (recently roughly \$27–30), so always measure it live on your own platform.

Why Markets Exist: Facilitating Trade

The core idea

Auction Market Theory grew out of the work of J. Peter Steidlmayer at the Chicago Board of Trade in the 1980s, and was later developed and popularised by James Dalton. Its central idea is simple and slightly surprising: **the purpose of a market is to facilitate trade**. The market's job is not to discover some fixed "correct" price. Its job is to find the prices at which the most business can be done.

The market works like a two-sided auction. Price moves up to find out whether sellers will become active or buyers will stop. Price moves down to test the same thing from the other side. Every move is an experiment.

Think of an auctioneer who does not know what a painting is worth. If many hands go up, they raise the price; if none do, they lower it. The auction settles where both sides keep trading. A futures market does the same, continuously.

Price advertises; time and volume answer

A useful way to remember this: **price advertises, and time and volume are the market's answer to the advertisement**.

- If price moves up and **trading increases**, the new price is acceptable to both sides. The market can stay there and do business.
- If price moves up and **trading nearly stops**, nobody is willing to transact at that price. It has gone too far, and the market moves back to find business elsewhere.

That is why AMT always cares about two things together: **price** (where) and **time plus volume** (how much business was done there). A high-volume area is a zone of temporary agreement. A low-

volume area means either disagreement (one side refused to trade) or speed (price passed through too quickly for business to build up).

Notice what high volume does *not* mean. It does not mean "a strong move in one direction". Heavy volume at a price usually means agreement in that area, not direction.

Two kinds of participant

AMT divides participants by their time horizon.

- The **Other Timeframe (OTF)** participant has a longer horizon than a single day: funds, commercial hedgers (such as miners or jewellers protecting against price changes), and traders who hold positions for days or weeks. When OTF participants become active, they tend to **move price away** from the current value area, because they are willing to pay up (or sell down) to get their business done.
- The **Day Timeframe** participant is a short-term trader who mostly opens and closes positions within the session. Day-timeframe activity tends to **rotate price inside** the current range, buying low and selling high within it.

No trade carries a label saying "OTF" or "day trader"; these are categories of *behaviour*, inferred from price, time and volume. Still, when price leaves the established range and is accepted outside it, a longer-horizon participant is usually at work.

A worked gold example

Over five sessions on a 30-minute chart, imagine two days that both push to new highs:

- **Day A:** GC spikes from 4,050.0 to 4,062.0 (XAUUSD $\approx$ 4,022 to 4,034) in the late morning. The volume profile shows almost nothing traded above 4,058.0. Within forty minutes price is back below 4,052.0. The market advertised higher prices, and nobody answered.
- **Day B:** GC rises from 4,050.0 to 4,062.0 over the morning, and then spends the afternoon trading between 4,058.0 and 4,066.0, building a thick volume node around 4,062.0. The market advertised higher prices, and both sides answered.

Both printed a new high; only Day B *accepted* it.

What AMT is not

AMT is a **descriptive framework**, not a forecasting formula. It tells you what the market is doing now: searching for new value, or rotating within current value. It does not tell you where the next price will be. Many popular statements associated with AMT ("price always comes back to fill X") have never been put through a proper statistical test.

Common mistake: "The market wants to reach the right price." The market wants trade; a "right price" is not a fixed thing.

Common mistake: "AMT tells me when to enter." It tells you what context you are in. That is valuable, but it is not an entry rule.

Key idea: Price moves to find trade. Where trade happens (time and volume), the market has accepted the price; where trade does not happen, it has rejected it. AMT is the common language behind both Market Profile and Volume Profile.

Balance vs Imbalance: Brackets and Trends

The two states

AMT says a market is always in one of two states: **balance** or **imbalance**.




Figure 19. Markets alternate between balance (two-sided rotation in a range) and imbalance (one side in control, price moving to find new value).

Balance means buyers and sellers broadly agree on a range of prices. Price **rotates** (swings back and forth) between an upper and lower boundary. Day after day, the value areas **overlap**, and a multi-day profile takes on a bell shape. The price zone where this happens is called a **bracket**.

Inside a bracket, traders typically act **responsively**: they sell near the top and buy near the bottom, expecting price to return towards the middle. Responsive activity means trading *against* the latest move, at the edges of value.

Imbalance means one side, usually the other timeframe, is more active and pushes price out of the bracket to search for a new place to trade. The signs are:

- Value areas **migrate** day after day, consistently higher or consistently lower (the value migration of Chapter 6).
- Profiles become **elongated**.
- Low-volume areas (LVNs, or single prints on a TPO chart) are left behind price.

In imbalance, **initiative** activity dominates: buying above value or selling below value. Initiative activity trades *with* the move, outside value.

The cycle

Markets alternate: **balance** → **breakout** → **imbalance (trend)** → **new balance at a different level**, and then the cycle repeats. Neither state is "better". Problems arise when a trader uses balance tools (buying the bottom of the range, selling the top) during imbalance, or uses imbalance tools (chasing breakouts) during balance.

That makes one question the most important of the day: **Are we in balance or imbalance right now?**

Worked example

Over twelve sessions, suppose gold's daily value areas look like this:

Days	Daily value areas (GC)	Reading
1–5	All overlap within 4,010–4,045	Balance; bracket ≈ 4,010–4,045 (XAUUSD ≈ 3,982–4,017)
6	Price closes above 4,045; value overlapping-higher	Test of the bracket top
7–9	Value areas 4,050–4,062, 4,061–4,075, 4,072–4,086	Imbalance; value migrating higher
10–12	Value areas overlapping around 4,078–4,090	New balance forming

Limits

A bracket's boundaries are a matter of judgement. Two analysts may draw two different brackets on the same chart. And a move out of a bracket is only a **test** until acceptance outside it is established (the next section explains how). Many breakouts fail and return inside: they are **failed breakouts**, and they are common.

Common mistake: "Every break of the high is the start of a trend." Many breakouts return inside the bracket.

Common mistake: "Balance means the market is safe." Breaks out of balance can be very fast.

Common mistake: Drawing a bracket from one candle's high and low instead of several days of overlapping value.

Key idea: Balance is rotation inside a bracket with overlapping value; imbalance is a directional search with migrating value. Responsive traders dominate balance, initiative traders dominate imbalance. Leaving a bracket without acceptance is just a test.

Acceptance vs Rejection: Time and Volume at Price

Definitions

Acceptance means the market spends **time** at a new price and builds **volume** there. Both sides are trading, and that price is turning into value.

Rejection means price reaches a level, no meaningful trading develops, and price returns quickly. Rejection usually leaves behind a tail (on a TPO chart) or a thin, low-volume area (on a volume profile).

A compact summary to remember:

Key idea: Time + volume = acceptance = new value. No time and no volume = rejection = excess.

Why it matters

When price crosses a reference level (say yesterday's value-area high, or the Initial Balance high), the only thing that separates a genuine move from a failed test is what happens **after** the cross. If price stays beyond the level, if subsequent bars close beyond it, and if volume builds there, the market is probably establishing new value. If price quickly returns, the level was defended.

How much time? How much volume?

There is no universal answer. Dalton, working with Market Profile, usually talks in terms of 30-minute TPO periods; for example, one full period beyond a level is often treated as initial evidence of acceptance. On a minute chart you need to **write your own rule in advance**, such as:

- "Three consecutive 1-minute bars close beyond the level", or
- "One 30-minute bar closes beyond the level", or
- "Price stays beyond the level for 20 minutes and a visible volume node forms there."

Then keep that rule fixed. Changing the definition after seeing the chart is self-deception.

Worked example

Yesterday's high (PDH) is 4,048.0 (XAUUSD $\approx$ 4,020). Consider two mornings, using the rule "acceptance = at least 30 minutes beyond the level with a visible volume node":

1. **Accepted:** GC crosses 4,048.0 at 09:40, trades between 4,049.0 and 4,056.0 until 10:30, and builds a node around 4,052.0. Acceptance above PDH.
2. **Rejected:** GC trades to 4,049.5 at 09:40, prints a few contracts there, and is back at 4,045.0 within four minutes. Nothing was built above the level. Rejection at PDH.

Limits

Acceptance is backward-looking. By the time it is confirmed, part of the move is already over. And acceptance can itself fail later: value accepted at 10:30 can be abandoned at 12:00.

Common mistake: "A wick above the level is a definite rejection." Without volume and context, a wick is just a wick.

Common mistake: Changing the time or volume threshold after seeing the result. That is hindsight, not analysis.

Common mistake: Ignoring time frame. What looks like acceptance on a 1-minute chart may still be rejection on a 30-minute chart.

TPO and Market Profile Basics vs Volume Profile

Chapter 6 introduced the difference between time-based and volume-based profiles. Here we look at how a TPO chart is actually built and read, because much of the AMT vocabulary in the rest of this chapter comes from it.

Building a TPO chart

The Market Profile was introduced by Steidlmayer at the CBOT and made publicly available in the mid-1980s. The session is cut into fixed periods, usually **30 minutes**, and each period gets a letter: A, B, C and so on. During each period, the period's letter is printed next to every price that traded. Each letter at a price is one **TPO (Time Price Opportunity)**.

Picture four columns of letters. Period A traded from 4,030 to 4,036, so A is printed at every price in that range. Period B traded from 4,033 to 4,040, so B is printed at those prices. C traded 4,034–4,038, D traded 4,035–4,039. Now push all the letters to the left so they stack in rows. The price rows from 4,035 to 4,036 have four letters (A, B, C, D); the rows at 4,030–4,032 have only A. The result is a "time histogram": the longer a price's row, the more time the market spent there.

From this chart come familiar concepts:

- **TPO POC:** the price with the most TPOs (the longest row).
- **Value area:** usually the range around the POC containing about 70% of all TPOs.
- **Initial Balance:** the range of periods A and B together, that is, the first hour.
- **Single prints:** prices that have only one letter.

Time versus volume

A volume profile asks "how many contracts traded here?" instead of "how many periods visited here?". The two usually look similar but can differ. A news release can produce a huge burst of volume at one price in three minutes, moving the volume POC there, while the market spent hardly any *time* at that price, so the TPO POC stays elsewhere.

Why this book leads with volume on gold

In a 23-hour market, TPO counts a quiet Asian half-hour exactly the same as a hectic half-hour after a New York data release, over-weighting the low-participation hours. A volume profile shows the real weight of participation, though it is sensitive to sudden volume spikes. So volume profile is our main tool, and TPO supplies the language: Initial Balance, periods, single prints and poor highs.

When the TPO POC and the volume POC are far apart, ask why. The answer usually tells you something about the day, typically that a short burst of news-driven trading carried a lot of volume.

Common mistake: Comparing TPO charts from two platforms that use different session templates (full Globex versus RTH). They are measuring different days.

Common mistake: "The 70% value area is a law of nature." It is a convention, inspired by the bell-shaped normal distribution in statistics, not a natural constant.

Initial Balance on Gold: Which Clock?

The classic definition

In classic Market Profile, the **Initial Balance (IB)** is the price range of the **first two 30-minute periods** of the session, that is, the first hour. Its top is the **IB High** and its bottom is the **IB Low**.

The idea behind it: during the first hour, day-timeframe traders establish an opening range. Whether the market later breaks out of that range (range extension, the next section) tells you whether the other timeframe has become active.




Figure 20. The initial balance is the first hour's range. A move outside it (range extension) shows that a longer-term participant has shown up.

Gold's problem: when does the session start?

For a stock-index future, the answer is easy: the stock market opens at 09:30 New York time. Gold has no such single moment. There are two common conventions:

1. The Globex open at 18:00 New York time. This is the true start of the CME trading day.

- *Advantage:* no part of the day is excluded.
- *Disadvantage:* the first Globex hour is usually thin, with little participation. Its IB is often very narrow and gets broken almost every day, so it carries little information.

2. 08:20 New York time. This was the opening time of the old COMEX trading floor (the "pit"), and a surge in volume is still visible around that time. It marks the start of gold's regular trading hours (RTH, roughly 08:20–13:30).

- *Advantage:* the IB is built when U.S. participants are active and London is still open, so it reflects real participation.
- *Disadvantage:* the overnight hours (Asia and London) are left out of the IB, so you must track the overnight high and low (ONH and ONL) separately.

Other conventions exist too: some traders use 09:30 to match the stock market, and some use a 90-minute IB.

In this book, gold's IB is 08:20–09:20 New York time. Which convention you choose matters less than **staying consistent**, and when you read someone else's statistics or lessons about the IB, ask which definition they used.

Worked example

On one day, the first Globex hour (18:00–19:00) trades between 4,031.0 and 4,033.5, a range of just \$2.50. By morning, gold has moved through that range many times; its high and low mean almost nothing. The 08:20–09:20 hour, by contrast, trades between 4,036.0 and 4,046.0 (XAUUSD $\approx$ 4,008–4,018), a \$10 range built on heavy participation. When price eventually breaks above 4,046.0 at 10:15 and stays there, that break carries information. The Globex IB's "break" carried none.

A daylight-saving detail

The United States and the United Kingdom do not change their clocks on the same dates. For a few weeks each year, in March and in October/November, the gap between London and New York shifts by one hour. An IB defined in New York time stays fixed in New York time, but its relationship to the London session changes during those weeks. Keep this in mind when comparing days around those dates.

Common mistake: Copying "IB = 09:30–10:30" from stock-index lessons and applying it to gold.

Common mistake: "The IB is broken on X% of days." Popular claims of this kind are not proven for your market and your definition: the figure depends entirely on how the IB is defined and on the market. Do not accept any such figure without testing it on your own data.

Gold note: Many platforms define session templates in Chicago time (CT), because CME is based in Chicago. 08:20 New York time is 07:20 Chicago time. If your IB looks wrong, check the time zone of your session template first.

Range Extension: When the Other Timeframe Shows Up

Definition

Range extension (RE) means price trades above the IB High or below the IB Low **after the IB period has ended**.

The Market Profile reasoning goes like this. Day-timeframe traders tend to rotate within the IB. For price to leave the IB with conviction, larger or longer-horizon participants (the OTF) usually have to get involved. So range extension can be seen as a possible footprint of the OTF. That is why the CBOT's Market Profile study material and Dalton's books note that the *absence* of range extension usually means short-term traders are in control.

Measuring in IB multiples

To compare days fairly, use the IB's own range as the unit. If the IB is \$8 wide and price extends to \$12 above the IB Low, the day's range is **1.5 × IB**. IB multiples make a quiet day and a busy day

comparable.

Worked numbers: IB Low 4,036.0, IB High 4,044.0 (an \$8 IB). Price later reaches 4,052.0 (XAUUSD $\approx$ 4,024). The range from IB Low is now \$16, or $2 \times$ IB.

These multiples are useful for **comparing days**. They are **not price targets**.

Three important situations

1. **One-sided RE with acceptance.** Price leaves the IB on one side, stays out, and builds volume. The structure starts to resemble a normal variation day or a trend day (see day types below).
2. **RE that quickly returns inside the IB.** This is a **failed extension**. Sometimes the market then travels to the other side of the IB.
3. **Two-sided RE.** Price extends above *and* below the IB. This produces a **neutral** structure: neither side has complete control.

The gold caveat: news inside the IB

Major U.S. economic releases arrive at 08:30, which falls **inside** an IB starting at 08:20. On those days, a single news spike can make the IB very wide, and any later range extension means less. Always ask: was today's IB a normal IB or a news IB?

Remember also a purely mechanical effect: a narrow IB is easier to extend than a wide one, because price has less distance to travel to leave it.

Common mistake: "The IB almost always extends to $2\times$ " or " $1.5 \times$ IB is the target." These are popular claims, not proven.

Common mistake: Automatically buying a break of the IB High without checking for acceptance and context.

Key idea: Range extension is a possible sign that longer-horizon participants are active. Measure it in IB multiples, check for acceptance, and remember that news inside the IB changes what it means.

Opening Types: How the Session Starts

Dalton's four types

Dalton classifies the start of the session into four types, ordered from the most to the least **conviction** (how strongly and clearly the market chooses a direction):

1. **Open-Drive (OD).** Price moves strongly in one direction right from the open and does **not** return to the opening price. Interpretation: the other timeframe made its decision before the

open.

2. **Open-Test-Drive (OTD)**. Price first tests a nearby reference level in one direction (for example yesterday's high or the overnight high), finds no business there, and then drives strongly in the opposite direction.
3. **Open-Rejection-Reverse (ORR)**. Price moves in one direction, meets opposition, and reverses back through the opening price to the other side. Conviction is lower, and the day's direction is not yet clear.
4. **Open-Auction (OA)**. Price rotates on both sides of the open with no clear direction. This usually signals balance and a rotational day, especially when the open is inside yesterday's value.

The conviction ladder is: **OD > OTD > ORR > OA**.

Picturing them

Draw the opening price as a horizontal line and sketch the first 30–60 minutes:

- **OD**: a nearly straight line away from the open.
- **OTD**: a small hook in one direction, then a straight line the other way, never returning to the open.
- **ORR**: a move one way, then back across the open and beyond it.
- **OA**: a zigzag around the open.

Worked example

Gold opens the RTH session at 4,040.0. The overnight high (ONH) is 4,044.0 (XAUUSD $\approx$ 4,016).

- Price rises to 4,043.8 by 08:28, stalls with almost no trading near the ONH, then falls steadily to 4,028.0 by 09:15 without returning to 4,040.0. That is an **Open-Test-Drive**: a test of the overnight high, no business found, then a drive lower.
- If instead price had fallen from 4,040.0 straight to 4,028.0 without any initial test upward, it would be an **Open-Drive**.
- If price had risen to 4,044.5, then fallen back through 4,040.0 to 4,034.0 and kept swinging, that would be closer to an **Open-Rejection-Reverse**.

How to use opening types

An opening type gives you an **initial hypothesis** about the day. OD and OTD make a directional, elongated day more likely; OA inside value makes rotation between known levels more likely. That hypothesis must then be confirmed or rejected by the rest of the day: range extension, acceptance, and so on.

Gold's caveat

Which moment is "the open"? This book uses 08:20 New York time. But gold has been trading since 18:00 the evening before, and the 08:20 open is really a continuation of a move that may have started in London. Opening-type labels are also partly **subjective**: one person's OD is another's

OTD. A practical fix is to write a quantitative criterion, for example: "OD = in the first 30 minutes, price never traded more than N ticks on the opposite side of the opening price", choosing N in advance and keeping it.

Common mistake: "OD means a guaranteed trend day." It only raises the likelihood; open-drives can fail.

Common mistake: Labelling the opening type after seeing the whole day (hindsight labelling).

Common mistake: Ignoring where the open sits relative to yesterday's value and range, the subject of the next section.

The Open vs Prior Value and Range

Chapter 6 introduced Dalton's three opening situations from the volume-profile side. Here we place them inside the AMT framework and turn each into scenarios.

Three situations, three sets of scenarios

Before the open, you have two things from yesterday: its **value area** (VAH to VAL, about 70% of volume) and its **range** (high to low). Where the open falls relative to these is your first clue about balance or imbalance.

1. Open inside yesterday's value. The market still accepts yesterday's prices. The base case is rotation and responsive behaviour; a day similar to yesterday is more likely. To leave value, the market has to find a reason.

2. Open outside value but inside range. A mild imbalance. Two main scenarios:

- Price returns into yesterday's value and is accepted there, which means the prices outside value have been **rejected**; or
- New value is built outside yesterday's value area.

3. Open outside yesterday's range. The strongest possible sign of imbalance. If price stays outside the range (acceptance), directional movement is more likely. If it returns inside the range, the move outside has failed, and traders who positioned in the direction of the opening move are left in a weaker position.

Worked example with scenarios

Yesterday's RTH: high 4,052.0, VAH 4,045.5, POC 4,041.0, VAL 4,036.5, low 4,029.0. Today at 08:20, GC is at 4,049.0 (XAUUSD $\approx$ 4,021): outside value, inside range.

- *Scenario A:* price falls back below 4,045.5 and is accepted inside yesterday's value. The higher prices were rejected; yesterday's POC at 4,041.0 becomes the next reference.

- *Scenario B:* price holds above 4,045.5, builds volume there, and later accepts above yesterday's high at 4,052.0. The market is building new value higher.

You do not choose between them in advance; you write them down so that you recognise each one as it unfolds.

Gold is not a stock gap

In equities, a "gap" is the jump between yesterday's close and today's open after a night with no trading. Gold's 08:20 "open" comes after about fourteen hours of overnight trading, so a gap in the stock-market sense does not exist. What we compare is **the price at 08:20 relative to the value and range of yesterday's reference session**. You must also decide whether yesterday's value comes from the whole Globex session or from RTH only; that choice changes the answer, so make it once and keep it.

The 80% rule, briefly

As Chapter 6 explained, the "80% rule" claims that when the market opens outside value, returns inside it and stays there for two consecutive 30-minute periods, it has an 80% chance of crossing to the other side of value. This is a popular claim, not proven. No transparent, repeatable public test exists for gold, and the result is sensitive to how value and the session are defined. Only your own carefully designed backtest could examine it. Until then, treat it as folklore.

Common mistake: "An open outside the range means a trend day." Many opens outside the range return inside.

Common mistake: Computing value from RTH on one day and from Globex on the next.

Key idea: The further the open is from yesterday's value, the greater the chance of imbalance, but never a certainty. A return into value with acceptance there means the outside prices were rejected.

Day Types in Brief: Normal, Variation, Trend, Double Distribution, Neutral

The classification

Market Profile classifies finished days mainly by two things: **the size of the IB** and **how much range extension occurred, and on which side**.

- **Normal day.** A wide IB that covers nearly the whole day's range; little or no extension. The market found both sides in the first hour and rotated for the rest of the day.

- **Normal variation day.** A moderate IB, then significant extension to **one** side (in common descriptions, up to about twice the IB). The other timeframe entered on one side, but without trend-day intensity.
- **Trend day.** A relatively narrow IB, then persistent extension to one side. Price barely retraces, leaves single prints behind, and closes near its extreme.
- **Double distribution trend day.** A quiet start with a narrow IB and a first small distribution; then a fast move and a second distribution at a different level, with a zone of single prints between the two.
- **Neutral day.** Extension on **both** sides of the IB. If the day closes near one of its extremes, it is a **neutral extreme** day (one side won by the close). If it closes in the middle, it is a **neutral center** day (uncertainty).

Dalton also describes other types, such as the **non-trend day**: a narrow IB with no extension, often seen ahead of major news.

Day type	IB	Extension	Close
Normal	Wide	Little or none	Anywhere
Normal variation	Moderate	One side, moderate	Often towards the extension side
Trend	Narrow	One side, persistent	Near the extreme
Double distribution	Narrow	One fast leg, two distributions	Varies
Neutral center	Any	Both sides	Middle
Neutral extreme	Any	Both sides	Near one extreme

Worked example

IB 08:20–09:20: 4,036.0–4,044.0 (an \$8 IB). Consider three ways the day might finish:

- Day range 4,033.0–4,046.0 by the close: barely extended. Closer to a **normal** day.
- Price extends to 4,056.0 (XAUUSD $\approx$ 4,028), about $2.5 \times$ IB from the IB Low, never returns to the IB, and closes at 4,055.0. Closer to a **trend** day.
- Price extends to 4,049.0 in the morning, then to 4,031.0 in the afternoon, and closes at 4,040.0. Extension on both sides with a central close: **neutral center**.

What day types are really for

Day types are mostly a **post-session review** tool. During the day you can only say "so far, the structure looks like X" and update that view each period. Labelling is subjective and the boundaries (normal variation versus trend, for example) are fuzzy: a different definition could label the same day differently.

About gold

You may come across claims that a very high share of gold's RTH sessions are trend days. That is a popular claim, not proven: without a precise definition and a test on real data, it should not be repeated. The general idea behind it, that gold may produce directional days more often than stock indices do, is worth investigating, but it is an idea to test, not a fact.

Common mistake: Trying to "predict" the day type in the first ten minutes.

Common mistake: Quoting precise frequencies of day types without stating the definition and the data. These are popular claims, not proven. And using thresholds designed for stock-index futures on gold without testing them.

Excess, Poor Highs and Lows, and Unfinished Auctions

How auctions end

When an auction in one direction is **complete**, it usually ends with **excess**. Price goes a little too far, meets strong opposition, and comes back quickly. On a TPO chart, excess appears as a **tail**: several prices at the extreme of the day, each with only one letter (in Dalton's common description, at least two single TPOs). On a volume profile, excess shows up as volume dropping off sharply towards zero at the edge of the profile.

The interpretation: "At that price, the opposite side found the market so attractive that they stepped in hard. The auction in that direction is finished."

Poor highs and poor lows

The opposite case is a **poor high** or **poor low**: an extreme with **no tail**. The profile is "flat" at the edge. For example, several periods traded at the exact same high price without price going higher and being rejected. This suggests the auction in that direction did not reach a conclusion: buyers became tired, rather than being pushed back by strong sellers. This is also called an **unfinished auction**.

Worked illustration: three 30-minute periods (D, E and F) all touch 4,061.0 exactly, and none trades at 4,061.1 or above. The high has three TPOs at the same price and no tail. That is a textbook poor high.

Why poor extremes are marked

Poor highs and lows are kept as **reference levels for the future**, based on the idea that the market will return at some point to "finish the auction". That is a popular claim, not proven: "a poor high always gets retested" is folklore. Even educational sources that teach the concept stress that a revisit is not guaranteed; in long trends, a poor extreme may never be revisited. Their real value is

knowing that the level is "untidy", and that resting orders (such as stop orders placed just beyond an obvious high) may have accumulated nearby.

A note on vocabulary

The term "unfinished auction" is also used on footprint charts (Chapter 9) for a related but different idea at the level of a **single bar**. Keep the two apart: here we are talking about the extremes of a whole session's profile.

Common mistake: "Every poor high gets filled or tested." A popular claim, not proven.

Common mistake: Treating every long candle wick as excess. Excess is defined on the profile (tails of single prints, volume dropping sharply at the edge), not by a candle's shape.

Common mistake: Forgetting the session definition. A poor high on an RTH profile may have a perfectly good tail on the full Globex profile.

Key idea: Excess (a tail) marks a decisive end to an auction and is usually a strong reference level. A poor extreme (a flat edge) marks an unfinished auction: a weaker, untidy level worth marking but never worth assuming will be revisited.

Single Prints and Tails

Definition

On a TPO chart, a **single print** is a price where only **one** letter appears: only one 30-minute period traded there. When several single prints are stacked one above another, the market passed through that zone quickly in one period and never came back. One side was so dominant that the other side had no opportunity to trade.

Two important places

1. **At the day's extreme:** a run of single prints is a **tail**, the same thing as the excess described above. It marks the rejection of price and the end of the auction in that direction.
2. **In the middle of the profile:** a **single-print section**, typically seen on trend days and double-distribution days. It acts as a "bridge" between two value areas. Its volume-profile equivalent is a **low-volume node (LVN)**.

Why they matter

These are zones where the market did not build any agreement. If price returns to them later, two scenarios are plausible:

- The zone is crossed quickly again, because still nobody wants to trade there; or
- The market finds support or resistance there, because the side that dominated originally steps in to defend it.

That is why single prints are marked as reference levels, not as signals. The claim that "single prints always get filled" is a popular claim, not proven. Whether they are filled depends on context, and no reliable fixed percentage has been established.

Technical limits

A single print depends **completely** on the period length. With 30-minute periods, a zone might show single prints; with 15-minute periods, the same zone might have two letters. And in gold's 23-hour market, a full Globex profile may already have filled in single prints that appear on the RTH profile. As with every profile tool, fix your settings and keep them constant.

Worked illustration: on a double-distribution day, period C drops gold from 4,046.0 to 4,030.0 (XAUUSD $\approx$ 4,018 to 4,002) in thirty minutes, and no later period trades between 4,032.0 and 4,042.0. Those ten dollars are a single-print bridge between the morning and afternoon distributions, and the volume profile of the same day shows an LVN in the same place.

Common mistake: Comparing single prints across platforms with different period lengths.

Common mistake: Equating a single print with every LVN, without considering how and when it was created.

Gold's Session Map: Asia, London, New York

Why split the day?

Gold trades almost 23 hours a day, but participation is far from uniform. To turn a continuous price line into an **auction map**, we divide the day into windows and record the high and low of each one.

Exact window definitions differ between platforms and traders, so the most important rule is to write yours down. The definitions used in this book are:

Window	Definition	Character
Overnight (ON)	From the Globex open (18:00 New York) to 13:00 UK (about 08:00 New York)	Everything before New York
Asia	00:00–08:00 UK	Usually low volume; moves often rotational
London	08:00–13:00 UK	Participation rises; London's physical gold market is active
New York	13:00–21:00 UK	Contains the 08:30 data, the 08:20 RTH open, the 13:30 settlement
RTH	08:20–13:30 New York	Gold's main hours; IB = 08:20–09:20

London is the centre of the physical gold market; the twice-daily **LBMA gold price** benchmark auctions take place at about 10:30 and 15:00 London time. The New York window holds three landmark times: the 08:20 RTH open, the 08:30 U.S. data, and the 13:29–13:30 window from which CME calculates gold's daily settlement price.

Each session runs its own small auction

Each window can hold its own small auction: Asia often builds a narrow bracket, London tests or breaks it, and New York decides which move is accepted. Each window's high and low mark where its auction stopped, making them reference levels for the next session.

Worked illustration: Asia trades 4,030.0–4,036.0 (XAUUSD $\approx$ 4,002–4,008) on light volume. London breaks above to 4,043.0 and returns to 4,038.0. By 08:00 New York time, the overnight range is 4,030.0–4,043.0. New York's first question is simple: will the market accept London's higher prices, or go back into the Asian range?

Technical details

- **Daylight saving.** As noted for the IB, UK and U.S. clocks change on different dates. Defining each window in its own local time (London windows in London time, New York windows in New York time) keeps "London" aligned with London's real business hours throughout the year.
- **Holidays and half-days.** On holidays and shortened sessions, windows are cut short. Use the high and low of an incomplete window with caution, if at all.

Common mistake: Assuming every platform defines the "Asia session" the same way.

Common mistake: "London always makes the high or low of the day." A popular claim, not proven.

Common mistake: Treating Asia's high and low as strong levels without noting how little volume built them.

Reference Levels: PDH/PDL, ONH/ONL and How to Use Them

What makes a reference level?

A **reference level** is a price where the auction has already done something important: it stopped there, was rejected there, or built value there. The most important ones for gold:

- **PDH / PDL:** the prior day's high and low (always using the same session definition).
- **ONH / ONL:** the overnight high and low, up to the New York session.
- **Asia and London highs and lows:** the boundaries of the overnight's smaller auctions.
- **Prior VAH, VAL and POC**, and any **naked POC** (a POC not yet revisited), from Chapters 5 and 6.
- **Poor highs and lows, and single prints** (earlier in this chapter).
- **IB High and IB Low**, once the IB has formed.

How to use them

A level is a **place to ask a question**, not a place for an automatic order. When price reaches a level, the question is: *What is the market doing here?*

- If price crosses the level and builds time and volume beyond it (**acceptance**), the level has been broken, and attention moves to the next reference level.
- If price comes back quickly (**rejection**), the level has been defended.

This is exactly where footprint charts (Chapters 8 and 9) help: they show who was aggressive at the moment of the touch, and whether that effort achieved anything.

For each level on your map, write two scenarios in advance:

Level (GC)	XAUUSD (approx.)	If accepted beyond	If rejected
ONH 4,043.0	≈ 4,015	Overnight highs taken; watch PDH 4,048.0	Back into the overnight range; watch prior POC
PDH 4,048.0	≈ 4,020	Outside yesterday's range; watch for new value	Yesterday's range holds; watch prior VAH
Prior VAL 4,031.0	≈ 4,003	Below yesterday's value; watch PDL	Value defended; rotation back towards POC

Two practical rules

1. **Fewer is better.** A chart with 25 lines makes every price look "near a level", and therefore tells you nothing.
2. **Merge nearby levels into zones.** If the ONH is at 4,047.5 and the PDH at 4,048.0, they are effectively one zone: 4,047.5–4,048.0.

An important limit

"Price reacted at the level" almost always looks true, because price **always** reacts somewhere: it never moves in a straight line. Without comparing behaviour at your levels with behaviour at random prices (a **baseline**), you cannot say that a level "works". Chapter 12 shows how to make such comparisons.

Common mistake: Drawing dozens of lines and "finding" a reaction at each one. This is confirmation bias: seeing what confirms what you already believed.

Common mistake: "A naked POC almost always gets touched." A popular claim, not proven.

Common mistake: Using PDH and PDL calculated from different sessions on different days.

Putting AMT Together: A Pre-New York Checklist for Gold

Why a checklist?

AMT becomes useful when it turns into a **fixed routine**. The goal of this checklist is not to predict direction. It is to know, before the main session opens, **what it would mean if X happened**. It takes about ten minutes before 08:20 New York time and moves from the big picture down to the detail.

The nine steps

1. **Big picture (several days).** Have recent value areas been overlapping (balance, a bracket) or migrating (imbalance)? Where are the bracket's edges?
2. **Yesterday.** What was the day type? Did it close near an extreme or in the middle? Did its extremes show excess, or a poor high or poor low?
3. **Overnight.** Mark the ONH and ONL, and the Asia and London highs and lows. Did the overnight session stay inside yesterday's value or leave it? Was overnight volume heavier or lighter than usual?
4. **Where are we opening?** Where is price around 08:20 relative to yesterday's value and range? (Inside value / outside value but inside range / outside range.)
5. **Calendar.** Is there U.S. data at 08:30 or 10:00? If so, the IB may become a "news IB".

6. **Levels.** At most six to eight levels or zones, each with an approximate XAUUSD equivalent if you follow the CFD.
7. **If-then scenarios.** For example: "If price is accepted above the ONH, yesterday's value has been rejected; watch the PDH. If the ONH is rejected and price returns into value, the rotation scenario towards yesterday's POC is in play." Always write at least two scenarios.
8. **After the IB.** What was the opening type? Was the IB wide or narrow compared with its recent average (for example, the average of the last 20 sessions)? Update your day-type hypothesis.
9. **After the session.** Compare your sheet with what actually happened. Where was the scenario right, and where wrong? This journal becomes the data from which you can later build your own real statistics.

A worked context sheet

Item	Entry
Multi-day	Value overlapping for 4 days, bracket $\approx$ 4,020–4,055 GC (XAUUSD $\approx$ 3,992–4,027)
Yesterday	Normal variation up; closed near the high; poor high at 4,052.0
Overnight	ONH 4,050.5, ONL 4,038.0; stayed inside yesterday's value; light volume
Open (08:20)	4,046.0: inside yesterday's value (VAH 4,047.0)
Calendar	08:30 data release: possible news IB
Levels	PDH/poor high zone 4,050.5–4,052.0; prior VAH 4,047.0; prior POC 4,041.5; ONL 4,038.0; bracket top 4,055
If-then A	Accepted above 4,052.0 → outside yesterday's range and the bracket top is near; watch for new value
If-then B	Rejected at 4,050.5–4,052.0 and back inside value → rotation scenario towards prior POC 4,041.5
After IB	(filled in at 09:20)
Review	(filled in after 13:30)

Notice what the sheet does not contain: no prediction, no entry, no target.

Limits

This sheet does not remove risk and does not guarantee an edge. Its value lies in **consistency** and in reducing emotional decisions: when something happens, you have already thought about what it means.

Common mistake: Writing only one scenario. That is directional bias dressed up as analysis.

Common mistake: Moving levels during the session to "fit" the move.

Common mistake: Believing that a completed checklist means a trade must follow. Many days, the honest conclusion is that nothing on the map was resolved clearly.

Try it: Fill in the context sheet before 08:20 New York time for ten consecutive sessions. Take a screenshot of your chart at 08:10 and another at 13:30. After the tenth day, count how often each of your written scenarios played out, without changing anything you wrote in advance.

Chapter summary

- **Markets exist to facilitate trade.** Price advertises; time and volume show whether the market accepts the advertisement. The other timeframe moves price to new value; the day timeframe rotates it within current value.
- **Balance** is rotation inside a bracket with overlapping value and responsive trading. **Imbalance** is a directional search with migrating value and initiative trading. Markets cycle between them.
- **Acceptance** is time plus volume at a new price; **rejection** is a quick return on little volume. Write your acceptance rule before you look at the chart, and remember that confirmation always comes late.
- A **TPO chart** prints one letter per 30-minute period at each traded price, measuring time. It supplies key vocabulary (IB, single prints, tails), but on 23-hour gold it over-weights quiet hours, so this book leads with volume.
- Gold's **Initial Balance** in this book is 08:20–09:20 New York time. Consistency of definition matters more than the choice itself.
- **Range extension** beyond the IB hints at other-timeframe activity. Measure it in IB multiples, check for acceptance, and watch for 08:30 news inside the IB.
- **Opening types** (Open-Drive, Open-Test-Drive, Open-Rejection-Reverse, Open-Auction) and the **open relative to prior value and range** give initial hypotheses, not signals. The 80% rule is folklore until tested.
- **Day types** (normal, normal variation, trend, double distribution, neutral) are mainly review tools. Labels depend on definitions.
- **Excess** (tails) marks a completed auction; **poor highs and lows** mark unfinished ones. **Single prints** mark fast, one-sided movement. None of them is guaranteed to be revisited.
- **Gold's session map** (overnight, Asia, London, New York, RTH) gives each window its own high and low, which become **reference levels** along with PDH/PDL, ONH/ONL, prior value, naked POCs and poor extremes.

- A **pre-New York checklist** turns AMT into a routine of context, levels and if-then scenarios, followed by an honest review.

Checklist

- I can say whether the market is currently in balance or imbalance, and why.
- I have a written acceptance rule (time and volume) that I do not change after seeing the chart.
- I know my platform's session template and time zone, and my IB is 08:20–09:20 New York time (or another written convention I keep).
- I measure range extension in IB multiples (never as targets) and note whether 08:30 data fell inside the IB.
- I record the opening type with a written criterion, not by feel after the fact.
- I classify the 08:20 price relative to yesterday's value and range, and label day types only after the session.
- I mark poor highs, poor lows and single prints as references, never as certainties.
- I keep six to eight levels or zones at most, merging levels that sit close together, with XAUUSD equivalents where needed.
- For each level I have written an acceptance and a rejection scenario, and I review the sheet after the close.

Quiz

1. According to Auction Market Theory, what is the main purpose of a market, and what do time and volume tell you about a price?
2. Daily value areas have overlapped for five consecutive days. Which market state is this most likely to be, and which type of trading activity usually dominates it?
3. Price breaks above yesterday's high, spends 40 minutes there and builds a volume node. How would you describe this, and why?
4. Price tests the overnight high just after 08:20, finds almost no trading there, then drives lower without ever returning to the opening price. Which opening type is this?
5. The day's high shows three TPO periods at exactly the same price and no tail. What is this called, and what is the honest way to use it?

Quiz answers

1. To facilitate trade: to find the prices at which the most business can be done. Time and volume are the market's response to a price: lots of time and volume means the price is accepted (agreement), while little of either means it is rejected or was passed through quickly.
2. Balance, with price rotating inside a bracket. Responsive activity usually dominates: selling near the top of the range and buying near the bottom.

3. Acceptance above the prior day's high. Price spent time beyond the level and built volume there, which is the definition of acceptance (time plus volume). It suggests the market may be building new value higher, though acceptance can still fail later.
4. Open-Test-Drive: a test of a nearby reference level (the overnight high) that found no business, followed by a strong drive in the opposite direction without returning to the open.
5. A poor high, an unfinished auction. Mark it as a reference level for future sessions, knowing it is an "untidy" level where resting orders may have gathered. Do not assume it will be revisited: "poor highs always get retested" is a popular claim, not proven.

PART IV • INSIDE THE BAR

8

Footprint Charts: Inside the Bar

In this chapter: You will open up a single candle and look at every trade inside it. You will learn what a footprint chart is built from, how to read one cell ("bid × ask"), how to calculate row delta and bar delta by hand, which display layouts answer which questions, where the volume sits inside a bar (the bar POC), why bar delta and candle colour sometimes disagree, how diagonal imbalances and stacked imbalances are calculated, how two settings quietly change what you "see", and how to choose a sensible footprint set-up for gold futures. Just as important, you will learn what a footprint cannot show you, so you never mistake a detailed description for a prediction.

What Is a Footprint Chart?

A candle is a summary. It gives you four numbers — open, high, low and close — and nothing else. That is enough to tell you *what* happened to price over a period of time, but it says nothing about *how* it happened. Inside one busy one-minute candle on gold futures (GC), hundreds of separate trades may have taken place, at a dozen or more different prices, some started by impatient buyers and some by impatient sellers. The candle throws all of that away and keeps only four prices.

A **footprint chart** keeps it. It "opens" the candle and shows, for **every price level inside the bar**, how many contracts traded there and, of those, how many were traded by an **aggressive buyer** (a market buy order that lifted the ask) and how many by an **aggressive seller** (a market sell order that hit the bid). You learned the idea of the aggressor in earlier chapters: every trade has two sides, a passive side that was resting in the order book with a limit order, and an aggressive side that arrived with a market order and crossed the spread to trade immediately. The footprint is simply that aggressor information, sorted by price, for each bar.




Figure 21. One footprint cell: volume that sold into the bid (left) and bought from the ask (right) at a single price.

Different platforms use different names for the same idea. You will see it called *Numbers Bars* (Sierra Chart), *Volumetric Bars* (NinjaTrader), *Cluster* or *Footprint* charts (ATAS), and other names elsewhere. The labels differ; the underlying data is the same.

The data a footprint needs

A true footprint needs **tick data**: a record of every individual trade, with its price, its size and — crucially — whether it executed at the bid or at the ask (or, equivalently, which side was the aggressor). On CME futures such as GC and MGC, that information can be derived from the

exchange's market data feed, because every trade happens in one central order book and the exchange reports it.

If a platform only has candle data (open, high, low, close and total volume per bar), it cannot build a real footprint. Some tools "estimate" a footprint by spreading a bar's volume across its range or guessing buy and sell volume from whether the candle closed up or down. Whatever that is, it is not a footprint. It is a drawing that looks like one.

Gold note: This is one reason this book reads order flow on gold *futures* rather than on the XAUUSD CFD. A CFD broker runs its own price feed and its own internal book; there is no single central record of who hit the bid and who lifted the ask across the whole market. The footprint is a futures tool. If you trade XAUUSD, read the footprint on GC and translate the prices: XAUUSD has typically sat roughly \$27–30 below the front GC contract, so a GC footprint level at 4210.6 corresponds to roughly 4181–4184 on XAUUSD. Always treat that gap as approximate and check it live, because it drifts.

What a footprint shows

A footprint shows **executed trades**: how much traded, at which price, and which side was the aggressor. That is three pieces of information a candle does not give you — who was pushing, where they pushed, and how hard.

What a footprint does not show

This list is just as important:

- **Resting limit orders that never traded.** If a seller had 300 contracts sitting on the offer at 4211.0 and nobody bought them, the footprint shows nothing at 4211.0. The order book (the DOM or a liquidity heatmap, covered in Chapter 10) shows resting orders; the footprint does not.
- **Cancelled orders.** An order placed and pulled before it traded leaves no footprint.
- **Intent.** A market buy might be a new trader opening a long position, or a short seller closing a losing position, or a hedger, or an algorithm rebalancing. The footprint records the aggressive buy; it cannot tell you why it happened.
- **Identity.** It does not show who traded. There is no "smart money" column.

Why it is useful anyway

The footprint lets you ask the questions from Chapter 7 (Auction Market Theory) at the scale of a single bar. When price touched yesterday's high, who was aggressive? Did all that effort — the volume and delta — produce a result in price? At the very top of the bar, did the auction finish, or were both sides still trading? These are good questions, and a candle cannot answer them.

Key idea: A footprint is a *description* of aggression inside a bar. It is more information than a candle, not a crystal ball. Seeing more detail does not automatically mean predicting better.

Common mistake: "The footprint shows what smart money is doing." It shows executed trades sorted by price and aggressor side. It does not show who the traders were, what they intended, or the orders that sat in the book without trading.

Reading One Cell: Bid × Ask

The most common footprint display puts **two numbers on each row**, where each row is one price level. The format is **Bid × Ask**.

- **The left number (Bid)** is the number of contracts that traded at this price **on the bid**. Someone sent a market sell order and it hit a buyer's resting limit order at this price. In plain words: *aggressive selling at this price*.
- **The right number (Ask)** is the number of contracts that traded at this price **on the ask**. Someone sent a market buy order and it lifted a seller's resting limit order at this price. In plain words: *aggressive buying at this price*.

Here is a single row from a hypothetical one-minute GC bar:

Price (GC)	Bid (sold into bid)	Ask (bought from ask)	Row volume	Row delta
4210.6	22	81	103	+59

Read it aloud: "At 4210.6, 22 contracts traded with an aggressive seller and 81 contracts traded with an aggressive buyer." The **row volume** is $22 + 81 = 103$ contracts. The **row delta** is $\text{Ask} - \text{Bid} = 81 - 22 = +59$.

From rows to the whole bar

A full bar is just a column of these rows. Here is a small, complete, hypothetical bar so you can practise the arithmetic:

Price (GC)	Bid	Ask	Row volume	Row delta
4210.9	0	6	6	+6
4210.8	14	31	45	+17
4210.7	28	40	68	+12
4210.6	22	81	103	+59
4210.5	47	35	82	-12
4210.4	19	8	27	-11
Total	130	201	331	+71

- **Bar volume** = the sum of every number in the bar = $130 + 201 = 331$ contracts.
- **Bar delta** = sum of all Ask minus sum of all Bid = $201 - 130 = +71$. You get the same answer by adding the row deltas.

Every buy has a seller

This is the single most important idea for reading footprints correctly, and it is the one most often misunderstood. **Every trade has a buyer and a seller.** For every contract bought, one contract was sold. So in the row above, the "81" does not mean 81 contracts of buying versus 22 of selling in the sense of "more buyers than sellers". It means 81 contracts in which the *buyer was the one in a hurry* and the seller was passive (resting a limit order), plus 22 contracts in which the *seller was the one in a hurry* and the buyer was passive.

So positive delta means **"aggressive buyers were more active than aggressive sellers"**, not "there was more buying than selling". Total buying and total selling are always equal.

Key idea: Left = traded on the bid = aggressive seller. Right = traded on the ask = aggressive buyer. Row delta = Ask – Bid. Bar delta = $\Sigma \text{Ask} - \Sigma \text{Bid}$. Delta measures *who was aggressive*, not "more buying".

A technical caution: classification is not perfect

Platforms must decide, for each trade, whether it happened on the bid or the ask. They normally use the exchange's aggressor flag where it is available, or compare the trade price to the best bid and best ask at that instant. In very fast moments — a major data release, for example — quotes change many times within a millisecond, and some trades can be classified on the wrong side. Treat footprint numbers as "very nearly right", not as perfect to the contract, especially in news spikes.

Common mistake: "Bid means buyers, so a big bid number means lots of buying." The opposite is true. The bid column counts aggressive *selling* into resting bids. Some platforms also let you swap the column order, so always check which side is which on your own software before you read anything.

Try it: A row reads 40×15 (bid $\times$ ask). What is the row volume and the row delta? (Answer: volume 55, delta $15 - 40 = -25$. Aggressive sellers were more active at that price.)

Footprint Layouts: Bid $\times$ Ask, Delta, Volume, Profile

The footprint data is always the same — volume at each price, split by aggressor side — but platforms can display it in several ways. Each layout answers one question quickly and hides something else. Knowing what your current layout hides is part of reading it honestly.

1. Bid × Ask

Two numbers per row, as above. This is the most complete view: you can derive every other layout from it. It is required for spotting diagonal imbalances (later in this chapter) and finished or unfinished auctions at the bar extremes (Chapter 9). Its weakness is that it is busy and slow to read, particularly when a bar has thirty or forty rows.

Question it answers best: "Who was aggressive, and at exactly which prices?"

2. Delta per row

Only one number per row: Ask – Bid, usually coloured green for positive and red for negative. It shows at a glance where aggression was concentrated.

Question it answers best: "Where was aggression concentrated?"

What it hides: total volume. A row delta of +10 looks identical whether the row was 15 × 5 (20 contracts) or 995 × 1,005 (2,000 contracts). The first is a small one-sided trickle; the second is an enormous two-sided battle that happened to end almost even.

3. Volume per row

Only the total traded at each price. This shows where activity and agreement were — the bar's own point of control.

Question it answers best: "Where did most of the trading happen?"

What it hides: direction. You cannot tell whether a heavy row was driven by buyers, sellers, or both.

4. Profile (histogram)

The volume of each row drawn as a horizontal bar beside the candle — a miniature volume profile for each bar — sometimes split into bid and ask colours.

Question it answers best: "What shape is the volume inside this bar? Is it heavy at the top, the bottom or the middle?"

What it hides: exact numbers, unless you hover or add labels.

Here is the same hypothetical bar from the previous section in three of these views side by side:

Price	Bid × Ask	Delta	Volume
4210.9	0 × 6	+6	6
4210.8	14 × 31	+17	45
4210.7	28 × 40	+12	68
4210.6	22 × 81	+59	103
4210.5	47 × 35	−12	82
4210.4	19 × 8	−11	27

Notice how the Delta column makes 4210.6 jump out, while the Volume column shows that 4210.5 — a negative-delta row — was the second busiest price in the bar. Neither column alone tells the full story; the Bid × Ask column contains both.

Common combinations include "volume and delta" in two columns, or Bid × Ask numbers with the cell background shaded by volume. Most major platforms offer all of these views under their own names.

Key idea: One data set, several views. Learn on Bid × Ask because it hides nothing. Once you read it fluently, switch to Delta or Profile for speed — but always know what your current view has hidden.

Common mistake: Deciding from colours alone without looking at the numbers. A bright green cell on a row with 12 contracts and a pale one on a row with 1,200 contracts are not comparable. A more colourful or "professional-looking" layout does not contain more information; it contains the same information arranged differently.

The Bar POC: Where the Volume Sits Inside the Bar

In Chapter 5 you met the **point of control (POC)** of a volume profile: the price with the most volume, the place of greatest agreement. Every footprint bar has one too. The **bar POC** is simply the row with the **highest total volume** in that bar. Many platforms outline it with a box.

When you find it, ask: *where did the main agreement in this bar happen?*

Three common positions

POC in the middle of the bar. Volume is roughly symmetric around the centre. The bar was mostly two-sided rotation: price moved up and down through a central price where most business got done.

POC near the high of an up bar. Lots of trading happened at the top of the bar. There are two possible readings, and the POC alone cannot tell them apart:

1. Buyers kept transacting at higher prices and the market *accepted* those prices — trading volume built up there because both sides were willing to do business at the new, higher level.
2. Many aggressive buyers arrived at the top, and large passive sellers met them and soaked up their orders. That is the idea of *absorption*, which Chapter 9 examines carefully.

The next bars and the location (is this at a reference level?) decide between the two readings.

POC near the low of an up bar. The main volume was at the bottom of the bar, and the top of the bar was covered on thin volume. The move up happened, but it did not build agreement at the higher prices — price passed through them quickly without much trade.

Here is a small example of that third case:

Price (GC)	Bid × Ask	Row volume
4212.4	0 × 3	3
4212.3	2 × 9	11
4212.2	6 × 14	20
4212.1	31 × 44	75
4212.0	58 × 66	124 ← bar POC
4211.9	40 × 37	77

The bar is green (suppose it opened at 4211.9 and closed at 4212.3), but most of its business was done at 4212.0, near the bottom. The top three rows together hold only 34 contracts.

POC migration

Look at bar POCs in sequence. If the POC of each successive bar is higher than the last, short-term value is moving up — a miniature version of the value migration you studied in Chapter 6. If several bars in a row print their POC at roughly the same price, the market has built a small, short-term area of agreement there.

Bar (1-min)	Bar POC (GC)	Change
09:41	4208.6	—
09:42	4209.1	+5 ticks
09:43	4209.4	+3 ticks
09:44	4209.4	0
09:45	4209.9	+5 ticks

Over five minutes the bar POC rose by 13 ticks (\$1.30 on the price; about 4208.6 → 4209.9 on GC, or very roughly 4179–4182 on XAUUSD). That is a description of where agreement moved, nothing more.

Limits of the bar POC

- **Row size and bar type change it.** If you merge several ticks into one row (see the last section of this chapter), or switch from time bars to volume bars, the POC can move.
- **Thin bars make it nearly random.** In a one-minute bar with only 15 contracts, the "POC" may be the row with 4 contracts versus a neighbour with 3. That is not information.
- **Near-ties make it jump.** When two rows have almost equal volume, the POC can flip between them from one bar to the next. Don't read meaning into that flip.

Common mistake: "The bar POC always acts as support or resistance." This is a popular claim, not proven. Sometimes price reacts there, often it does not. Treat it as context about where agreement formed, and test any claim like this on your own data before relying on it.

Try it: On your platform, find ten consecutive one-minute GC bars during the New York morning and note each bar's POC. Then do the same for ten bars in the middle of the Asian session. In which session do the POCs look stable, and in which do they jump around between rows with almost the same volume?

Bar Delta vs the Close: When They Disagree

The simple expectation is: a green bar (close above open) has positive delta, and a red bar has negative delta. Most of the time that is what you see. But not always. Sometimes a bar has **positive delta and closes red**, or negative delta and closes green. When that happens, the bar is telling you: *one side was more aggressive, and it did not get the result.*

How can that be? Because **delta measures aggression, but limit orders move price too**. Three mechanical explanations cover most cases.

Reason 1: passive orders absorbed the aggression

Aggressive buyers keep lifting the ask, but a large passive seller keeps refilling their offer at the same price — sometimes as an *iceberg* order that shows only a small part of its size at a time. Delta climbs; price does not. The buyers spent effort; the seller absorbed it.

Reason 2: liquidity was pulled

Price can also move without much aggression at all. If buyers who were resting limit orders on the bid cancel them, the bid side thins out, and even a small amount of aggressive selling can push price down several ticks. Delta is only slightly negative, but the move is large. The footprint sees the small aggressive selling; it does not see the bids that disappeared.

Reason 3: timing inside the bar

Bar delta is a total for the whole bar. Aggressive buying may dominate the first forty seconds of a one-minute bar, pushing price up, and then a burst of selling in the last twenty seconds may push price back below the open. The total delta is still positive, but the close is lower.

Here is a hypothetical bar that shows the disagreement:

Price (GC)	Bid × Ask	Row delta
4215.3	4 × 38	+34
4215.2	21 × 96	+75
4215.1	33 × 58	+25
4215.0	45 × 41	−4
4214.9	28 × 12	−16
Total	131 × 245	+114

Suppose this bar opened at 4215.1 and closed at 4214.9. It is a red bar with bar delta of +114. A lot of aggressive buying happened at 4215.1–4215.3, yet price finished two ticks below the open.

Delta high and delta low inside the bar

Because of the timing problem, many platforms also show the **maximum and minimum delta reached during the bar** (sometimes called delta high/low or max/min delta). If a bar closes with delta +80 but its maximum delta during the bar was +300, buyers were very strong early on and then gave most of it back. That tells a fuller story than the closing number.

What the disagreement is good for

This disagreement is the raw material for two ideas in Chapter 9: **effort versus result** and **trapped traders**. On its own, however, it is not a signal. Bars where delta and candle colour disagree are not rare in ordinary conditions, especially in thin bars where a handful of trades decides the sign. Before you interpret them, count how often they occur on your own GC data during a normal session. They earn attention only in **context**: at a meaningful reference level, on meaningful volume.

Key idea: Delta tells you who was aggressive. The close tells you where the last trade happened. When they disagree, aggression did not get its result — and the reason is usually passive orders, pulled liquidity, or timing within the bar.

Common mistake: "Positive delta means price should go up." Delta is not a direct predictor of price. Also avoid comparing raw delta numbers between the Asian session and New York without accounting for the very different volume: +60 in a quiet Asian minute and +60 in the busiest New York minute are not the same event.

Diagonal vs Horizontal Imbalance

An **imbalance** is a row where one side was much more aggressive than the other — for example, three times as much aggressive buying as aggressive selling. That sounds simple until you ask the obvious question: *which two numbers do we compare?* There are two answers, and they can disagree completely.

Horizontal comparison

The intuitive method compares the Ask at a price with the Bid **at the same price**. For a row reading 10×60 , horizontal comparison says: $60 \div 10 = 6$, a six-to-one buy imbalance.

It is simple, but mechanically it compares two different moments. Here is why. Gold futures usually trade with a one-tick spread. Suppose at some instant the best bid is 4210.4 and the best ask is 4210.5. At **that** moment, an aggressive buyer trades at 4210.5 (the ask) and an aggressive seller trades at 4210.4 (the bid). For a trade to happen *on the bid* at 4210.5, the whole market must have been one tick higher — bid 4210.5, ask 4210.6. So the Bid and Ask numbers in the same row were created at **different times**, under different quotes.

Diagonal comparison

The diagonal method compares the two sides that were facing each other **at the same moment**: the Ask at price p against the Bid at *one tick below* ($p - 1$ tick). Several major footprint platforms make this the default or the main option.

- **Buy imbalance at price p :** $\text{Ask}(p) \div \text{Bid}(p - 1 \text{ tick})$ is at least the chosen ratio.
- **Sell imbalance at price p :** $\text{Bid}(p) \div \text{Ask}(p + 1 \text{ tick})$ is at least the chosen ratio.



Figure 22. Footprint imbalances compare the ask at one price with the bid one tick below (diagonal). Several in a row form a stacked imbalance.

Same data, two different answers

Take these two rows:

Price (GC)	Bid	Ask
4210.5	10	60
4210.4	40	15

- **Horizontal** at 4210.5: Ask 60 ÷ Bid 10 = **6:1**. "A huge buy imbalance."
- **Diagonal** at 4210.5: Ask(4210.5) = 60 ÷ Bid(4210.4) = 40 = **1.5:1**. Not an imbalance at a 3:1 ratio.

Same numbers, opposite conclusions. The diagonal answer is the one that reflects what buyers and sellers were actually doing against each other at the same instant.

A full worked example

Here is a hypothetical seven-row bar. We use a common textbook ratio of **3:1** and, for now, ignore minimum volume (next section).

Price (GC)	Bid	Ask	Diagonal buy check: $\text{Ask}(p) \div \text{Bid}(p-1)$	Diagonal sell check: $\text{Bid}(p) \div \text{Ask}(p+1)$
4211.1	0	7	$7 \div 9 = 0.8$	— (no row above)
4211.0	9	46	$46 \div 12 = \mathbf{3.8 \text{ buy}}$	$9 \div 7 = 1.3$
4210.9	12	52	$52 \div 15 = \mathbf{3.5 \text{ buy}}$	$12 \div 46 = 0.3$
4210.8	15	33	$33 \div 25 = 1.3$	$15 \div 52 = 0.3$
4210.7	25	29	$29 \div 70 = 0.4$	$25 \div 33 = 0.8$
4210.6	70	21	$21 \div 18 = 1.2$	$70 \div 29 = 2.4$
4210.5	18	6	— (no row below)	$18 \div 21 = 0.9$

Two buy imbalances (4211.0 and 4210.9). No sell imbalance: 4210.6 shows heavy selling at 70 contracts, but against 29 aggressive buyers one tick above it the ratio is only 2.4.

Now look at what horizontal comparison would have flagged: 4211.0 ($46 \div 9 = 5.1$), 4210.9 ($52 \div 12 = 4.3$), 4210.8 ($33 \div 15 = 2.2$, no), 4210.6 as a sell ($70 \div 21 = 3.3$, yes), 4210.5 as a sell ($18 \div 6 = 3.0$, yes), and 4211.1 ($7 \div 0$ — infinite). Two diagonal imbalances become as many as five horizontal ones: a different count, and a different story.

The edge rows have no partner

Notice the dashes. At the **lowest row** of a bar there is no row below, so a diagonal buy comparison has nothing to compare against. At the **highest row** there is no row above, so a diagonal sell comparison has nothing to compare against. These extremes need separate treatment — which is exactly what the "finished versus unfinished auction" idea in Chapter 9 does.

What an imbalance does and does not tell you

An imbalance says: "at this row, one side was clearly more aggressive." That is all. It does not tell you where price will go next, and it does not tell you whether that aggression succeeded or was absorbed.

Key idea: Horizontal compares $\text{Ask}(p)$ with $\text{Bid}(p)$ — two different moments. Diagonal compares $\text{Ask}(p)$ with $\text{Bid}(p - 1)$ for buying, and $\text{Bid}(p)$ with $\text{Ask}(p + 1)$ for selling — the two sides that actually met. Always know which one your platform uses.

Common mistake: Using horizontal imbalance without realising it, because it was a platform's default. A second mistake: "Every imbalance means big money arrived." A row can become an imbalance with a few dozen small contracts. A third: comparing imbalance counts between two platforms that calculate them differently.

Imbalance Settings: Ratio and Minimum Volume

Every footprint defines an imbalance with **two numbers**, and those two numbers decide how many imbalances you see. Change them and the "signals" on your screen can double or vanish, even though the market did exactly the same thing.

Setting 1: the ratio

The **ratio** is how many times larger the aggressive side must be than the diagonal opposite side. Common values range from about **1.5:1 to 4:1** (often written as 150% to 400%). Some educational material describes 2:1 as an imbalance and 4:1 as a strong imbalance; **3:1 (300%)** is a very widespread default.

A lower ratio means more imbalances and more noise. A higher ratio means fewer, rarer imbalances — but also more cases where genuinely one-sided activity is missed.

Setting 2: the minimum volume

The **minimum volume** is the smallest size the aggressive side must have before the row is checked at all. Without it, a row with three aggressive buys against zero aggressive sells on the diagonal counts as an imbalance — "three to nothing".

The divide-by-zero problem

That "three to nothing" case is a real mathematical problem. If the opposite diagonal cell is **zero**, the ratio is $3 \div 0$, which is undefined or infinite. Platforms must choose: either they treat any non-zero against zero as an automatic imbalance, or they skip the row. Either way, the minimum volume filter is what stops a few stray contracts from lighting up your chart. Find out what your platform does.

Watching the settings change the picture

Here is a hypothetical bar and three settings. All three settings use diagonal comparison.

Price (GC)	Bid	Ask	Ask(p) ÷ Bid(p-1)
4210.9	0	4	$4 \div 0$ (divide by zero)
4210.8	0	26	$26 \div 11 = 2.4$
4210.7	11	58	$58 \div 17 = 3.4$
4210.6	17	64	$64 \div 30 = 2.1$
4210.5	30	40	$40 \div 45 = 0.9$
4210.4	45	18	$18 \div 9 = 2.0$
4210.3	9	12	—

Settings	Buy imbalances found
Ratio 1.5:1, minimum 10	4210.8, 4210.7, 4210.6, 4210.4 → 4 rows (4210.9 excluded by the minimum)
Ratio 1.5:1, no minimum (platform counts anything against zero as an imbalance)	the same four plus 4210.9 (4 against 0) → 5 rows
Ratio 3:1, minimum 20	4210.7 only → 1 row
Ratio 4:1, minimum 50	none

The market was identical in every line. Only the settings changed. Between the loosest and the strictest settings, the "buy imbalance" count goes from five to zero.

Why gold needs its own settings

GC and MGC show the same price, but their volumes are very different: GC is the deep, heavily traded contract; micro gold (MGC) trades far fewer contracts' worth of gold. A minimum volume that is sensible on GC during New York may filter out almost nothing on MGC, or almost everything, depending on the period. Settings copied from equity index futures (ES, NQ) are calibrated to completely different volumes and should not be pasted onto gold.

Volume also changes through the day. A typical GC row in the quiet Asian session holds a fraction of what a row holds in the first hour of New York trading. A fixed setting across the day means that an "imbalance" in Asia and an "imbalance" in New York are not the same thing.

The honesty rule: choose settings before you look

Pick your ratio and minimum volume **before** you look at the results, write them down, and keep them fixed. If, after watching a move, you nudge the ratio until an imbalance appears in the "right" place, you are fitting your settings to the past. Every chart will then seem to confirm your method, and none of that confirmation will mean anything going forward. Chapter 12 returns to this trap in detail.

Key idea: Ratio and minimum volume define every imbalance. Lower ratio = more imbalances and more noise. The minimum volume protects you from "3 against 0". Set them per contract and per session, before looking, and keep them fixed.

Common mistake: "300% is the best setting." That is a popular claim, not proven. It is a reasonable common default, nothing more. Changing settings after every trade "to see it better" is the same mistake in another form.

Stacked Imbalances

A **stacked imbalance** is a run of imbalances **in the same direction on consecutive rows** within one bar — for example, buy imbalances at 4210.6, 4210.7 and 4210.8. The number of rows needed is a convention, not a law of nature; **three consecutive rows** is the most common definition, and some tools start counting from two. A row with no trades, or a row that is not an imbalance, breaks the chain.

The logic

A single imbalance can be an accident — one large market order sweeping a price, for instance. But when one side was clearly more aggressive across several ticks in a row, it is more likely that you are looking at **sustained initiative pressure**: buyers who were willing to keep paying up, price after price.

Here is a hypothetical bar with a three-row buy stack, checked diagonally at a 3:1 ratio with a 20-contract minimum:

Price (GC)	Bid	Ask	Ask(p) ÷ Bid(p-1)	Imbalance?
4210.9	6	19	$19 \div 14 = 1.4$	no (also below minimum)
4210.8	14	66	$66 \div 19 = 3.5$	buy
4210.7	19	72	$72 \div 21 = 3.4$	buy
4210.6	21	85	$85 \div 24 = 3.5$	buy
4210.5	24	37	$37 \div 52 = 0.7$	no
4210.4	52	30	—	—

Buy imbalances at 4210.6, 4210.7 and 4210.8: a **three-row buy stack**. The zone 4210.6–4210.8 GC (very roughly 4181–4184 on XAUUSD, using the approximate gap) is where buyers pressed hardest inside this bar.

Two opposite readings — and how location separates them

The same stack can mean very different things depending on **where** it appears.

Reading 1: genuine initiative. The stack appears at the *start* of a move — for example, as price breaks above the initial balance high or leaves a multi-day bracket (Chapter 7) — and afterwards price stays above the stack zone. The market accepted the higher prices. Traders often mark such a zone as "the place where buyers showed up forcefully" and watch how price behaves if it returns there.

Reading 2: late buyers. The stack appears at the *end* of a long rally, right into a major reference level such as the previous day's high, and afterwards price falls back below the stack zone. Now

the same aggressive buyers may be holding losing positions. This is the "trapped traders" idea in Chapter 9.

At the moment the stack prints, the two cases can look identical. What separates them is location in the larger structure and what the following bars do.

The frequency check

With a low ratio and a small minimum volume, stacks become very common. Before giving them any weight, count how many your settings produce on an ordinary GC day. If the answer is dozens per session, each individual stack carries very little information.

Key idea: A stack is several same-direction imbalances on consecutive rows. It describes sustained aggression in that bar. Location decides the reading: at the start of a move with acceptance, it suggests initiative; at the end of a long move with rejection, it may mark trapped traders.

Common mistake: "A stacked imbalance zone will act as strong support or resistance later." This is a popular claim, not proven. Sometimes it does; often it does not. Do not treat it as a rule without testing it on your own data. Also avoid counting stacks in thin Asian-session bars as if they meant the same as stacks in the New York open.

Footprint on Gold: Tick Size, Volume and Choosing the Bar

Everything so far applies to any futures market. This section makes it specific to gold.

The basic numbers

Gold futures move in **ticks of \$0.10**. On the full-size contract (**GC, 100 troy ounces**) one tick is worth **\$10** per contract; on the micro contract (**MGC, 10 troy ounces**) one tick is worth **\$1** per contract. (These were the CME specifications as of October 2026; always re-check the current contract specifications on the CME website.)

With one tick per row, each footprint row is one \$0.10 price step. A bar that covers a \$5 range therefore has **50 rows**. During a major data release gold can cover \$10 or more in a single minute, which means hundreds of rows — and a footprint you cannot meaningfully read in real time.

Bar range (GC dollars)	Rows at 1 tick per row	Rows at 3 ticks per row
\$1.00	10	about 4
\$3.00	30	10
\$5.00	50	about 17
\$20.00	200	about 67

GC or MGC?

GC and MGC track the same gold price, but the GC order flow is far deeper and reflects much more of the large, professional participation in the market. MGC is useful for trading smaller size, but its footprint is thinner and noisier. A sensible approach is to **read the footprint on GC even if you trade MGC** (or the XAUUSD CFD). The information is in the deep contract.

Gold note: If you trade XAUUSD, you can watch a GC footprint alongside your CFD chart and convert levels using the live gap. A stack at 4210.6–4210.8 on GC, with a gap of about \$28 at the time, sits around 4182.6–4182.8 on XAUUSD. The gap moves with interest rates and time to expiry, so measure it each session rather than memorising a number.

Bar type

Time bars (for example 1-minute or 5-minute). Simple, and they line up naturally with time-based references such as the initial balance, session opens and data releases. Their weakness is uneven volume: a one-minute bar in the middle of the Asian session might hold a small fraction of the volume of the one-minute bar at the New York open. Imbalances and POCs in those two bars are not comparable.

Volume bars (one bar every N contracts). Each bar contains the same volume, so footprints are much more comparable with each other. The cost is that the link to clock time is lost: in busy periods bars form very quickly, in quiet periods very slowly, and lining them up with session times becomes awkward.

Range bars and tick bars. Range bars close when price has moved a fixed distance; tick bars close after a fixed number of trades. Some styles find them useful, but the footprint inside them means something slightly different, because the bar's own construction shapes where it starts and ends.

Common mistake: "One special bar type gives more accurate footprint signals." This is a popular claim, not proven. Each bar type trades one weakness for another.

Ticks per row

Many platforms let you merge several ticks into one row (for example 2–5 ticks per row). That makes the footprint far easier to read during fast moves. But it changes the meaning of the diagonal imbalance: the comparison is no longer "the ask against the bid one tick below at the same moment", but a comparison of aggregated blocks. The logic from the diagonal-imbalance section is weakened. If you merge rows, be aware that your imbalances now mean something coarser.

A sensible starting set-up (not a rule)

As a starting point for learning, many traders use:

- **Contract:** GC.
- **Main footprint:** 1-minute time bars during the New York session, one tick per row.

- **Context chart:** a 5-minute chart (footprint or plain candles) for the bigger picture.
- **Imbalances:** diagonal, with a common ratio such as 3:1 and a minimum volume chosen from your own measurements.

Before you rely on any of this, **measure the average volume per bar in each session** on your own data. A small table like this, filled in with your own numbers, tells you where a sensible minimum volume sits and how differently a bar should be read in each session:

Session (New York time)	Average volume per 1-min GC bar	Typical rows per bar	Your notes
Asia (18:00–02:00)	measure it	measure it	
London (02:00–08:00)	measure it	measure it	
New York (08:00–13:30)	measure it	measure it	

Then keep your settings fixed so that today's footprint can be compared with last week's.

Try it: Take one 30-minute stretch of GC during the New York morning. View it four ways: 1-minute time bars, 5-minute time bars, volume bars, and 1-minute bars with three ticks per row. Count the buy imbalances in each version. Write down how much the count changes — that difference is entirely created by your settings, not by the market.

Common mistake: Reading a footprint on MGC and drawing conclusions about "institutions". And reading one-minute Asian-session footprints with the same expectations as the New York open.

Chapter summary

- A footprint chart opens a candle and shows, for every price in the bar, how many contracts traded with an aggressive seller (bid) and how many with an aggressive buyer (ask).
- It needs real tick data with the aggressor side; footprints "estimated" from candle data are not footprints.
- It shows executed trades only — not resting orders, cancelled orders, identity or intent.
- In Bid × Ask, left is aggressive selling into the bid and right is aggressive buying from the ask. Row delta = Ask – Bid; bar delta = $\Sigma \text{Ask} - \Sigma \text{Bid}$.
- Every buy has a seller; delta measures which side was aggressive, not "more buying".
- Layouts (Bid × Ask, Delta, Volume, Profile) show the same data; each hides something. Learn on Bid × Ask.
- The bar POC is the busiest row. Its position adds context; its migration describes where short-term agreement moves; in thin bars it means little.

- Delta and close can disagree because of absorption, pulled liquidity, or timing inside the bar. Delta high/low helps.
- Diagonal imbalance compares the two sides that met at the same moment ($\text{Ask}(p)$ vs $\text{Bid}(p-1)$; $\text{Bid}(p)$ vs $\text{Ask}(p+1)$); horizontal comparison mixes two moments and can give the opposite answer.
- Ratio and minimum volume define imbalances; changing them changes what you see. Choose them before looking, per contract and per session.
- A stack is several same-direction imbalances on consecutive rows; location and follow-through decide whether it reads as initiative or trapped traders.
- On gold: tick \$0.10 (\$10 on GC, \$1 on MGC). Read the footprint on GC. Every bar type and row size is a trade-off. Measure volume per session first.

Checklist

- My platform builds the footprint from tick data with aggressor side, not from candles.
- I have confirmed which column is Bid and which is Ask on my software.
- I can calculate row delta, bar delta and bar volume by hand.
- I know what my current layout hides.
- I find the bar POC and ask whether the bar has enough volume for it to mean anything.
- When delta and candle colour disagree, I check delta high/low and the location before interpreting.
- My imbalances are set to diagonal, and I know how my platform treats zero on the opposite side.
- My ratio and minimum volume are written down, chosen before looking, and separate for GC and MGC.
- I have counted how many imbalances and stacks my settings produce on an ordinary day.
- I read the footprint on GC, convert levels to XAUUSD with the live gap if needed, and measure average bar volume per session.

Quiz

1. Which of these can a footprint chart **not** show you? (a) Volume at each price (b) Aggressive buys versus aggressive sells (c) Limit orders that were cancelled before filling (d) Bar delta
2. A footprint row reads 40×15 (bid $\times$ ask). What is the row delta, and which side was more aggressive at that price?
3. A one-minute GC bar closes red with bar delta of +120. Give the most likely mechanical explanation, and name one other possible explanation.
4. Using diagonal comparison, a **buy** imbalance at price p compares which two cells? Using the rows 4210.5: 10×60 and 4210.4: 40×15 , is 4210.5 a buy imbalance at 3:1?
5. Why is it sensible to read the footprint on GC even if you trade MGC or the XAUUSD CFD?

Quiz answers

1. **(c)** The footprint records executed trades only. Cancelled limit orders (and resting orders that never traded) leave no trace on it; you need order-book tools for those.
2. Row delta = Ask – Bid = 15 – 40 = **-25**. Aggressive sellers were more active at that price.
3. Most likely: **passive sellers absorbed the aggressive buying** — buyers kept lifting the offer, but resting sell orders kept being refilled, so price did not rise. Other possibilities: the timing inside the bar (strong buying early, selling late pushed the close down) or liquidity being pulled on the bid side.
4. Diagonal buy imbalance compares **Ask(p) with Bid(p – 1 tick)**. Here: $\text{Ask}(4210.5) = 60 \div \text{Bid}(4210.4) = 40 = 1.5:1$ — **not** an imbalance at 3:1, even though the horizontal comparison ($60 \div 10 = 6:1$) would have flagged it.
5. Because GC has **much deeper order flow** and reflects far more of the market's large participation. The price is essentially the same, but the GC footprint contains much more information; you can translate its levels to MGC directly and to XAUUSD with the approximate live gap.

9

Footprint Patterns, Honestly

In this chapter: You will meet the best-known footprint patterns — exhaustion, finished and unfinished auctions, absorption, effort versus result, trapped traders, delta-POC shift and CVD divergence — and you will learn to define each one precisely, to notice when different sources use the same word for different things, and to treat every pattern as a way of reading context rather than as a standalone signal. The most important skill in this chapter is a question: *how often does this pattern appear, and when it appears, how often does it actually mean anything more than an ordinary bar would?* By the end you will be able to read a full gold footprint bar step by step, and you will see how a single bar can trigger two "signals" pointing in opposite directions.

A word before we start. Footprint patterns are usually taught with carefully chosen screenshots: the exhaustion top that marked the high of the day, the absorption at the low that held. Those pictures are real, but they are a selection. The bars where the same pattern appeared and nothing happened are rarely shown. In our own testing of footprint rules on gold, the honest conclusion was the same one you will meet throughout this book: patterns on their own did not reliably predict direction after costs. That does not make the footprint useless. It makes it a reading tool. This chapter teaches the patterns as vocabulary for describing what happened at important places — not as a list of triggers.

Gold note: All examples in this chapter use hypothetical GC (gold futures) prices around 4210. XAUUSD (the spot CFD) has typically traded roughly \$27–30 below front-month GC, so 4210 GC is very roughly 4180–4183 on XAUUSD. The gap changes over time; always measure it live and treat it as approximate.

Exhaustion: Thin Volume at the Extreme

Exhaustion describes the idea that the aggressive side has "run out" at the end of a move. Price makes a new high, but at the very top rows of the bar almost nobody trades. The buyers who were pushing price up have stopped arriving. On a footprint, this shows up as **very little volume in the extreme rows** compared with the rest of the bar — for example, compared with the bar's POC row. Some vendors describe it as a drop-off in follow-through volume.

In auction terms (Chapter 7), the auction has reached a price where there are no more aggressive buyers willing to transact. It is a one-bar cousin of *excess* at the end of a session profile.




Figure 23. Exhaustion: volume dwindles at the bar's extreme. It can show a push running out of fuel, but on its own it is not a reliable reversal signal.

A hypothetical exhaustion top

Price (GC)	Bid × Ask	Row volume	Note
4214.0	0 × 2	2	new high of the recent bars
4213.9	3 × 6	9	
4213.8	18 × 24	42	
4213.7	41 × 57	98	
4213.6	63 × 70	133	bar POC
4213.5	52 × 39	91	close
4213.4	20 × 15	35	open

The bar pushed to a new high at 4214.0, but the top two rows hold only 11 contracts out of 410, and the top row holds 2 contracts against 133 at the POC. Price then closed five ticks below the high.

Making it precise

To study a pattern honestly, you have to define it precisely enough that a computer could find every instance. Here is a **made-up teaching definition** of an exhaustion top (the numbers are illustrative, not a recommended setting):

1. The bar makes the highest high of the last several bars.
2. The top row's total volume is less than one-fifth of the POC row's volume.
3. The top two rows contain less than 15% of the bar's aggressive buying (Ask volume).
4. The bar closes at least a third of its range below the high.

The bar above passes all four: top row 2 versus POC 133; top two rows' Ask = 8 out of 213 total Ask (under 4%); close 4213.5 is five ticks below a high that sits six ticks above the low.

The bottom version is the mirror image: a new low, very thin volume in the lowest rows, little aggressive selling there, and a close well above the low.

The honest problem: most bars look like this

Here is the uncomfortable fact. **Thin volume at the extremes is the normal shape of almost any bar.** Volume gathers in the middle of a bar because price spends most of its time there; price usually touches the high and low only briefly, so the extreme rows naturally hold few contracts. Draw the volume histogram of a typical one-minute bar and it bulges in the middle and thins at both ends.

That means "low volume at the extreme" is mostly a description of an ordinary bar, not of a special event. When we tested a reasonable-looking exhaustion rule on gold, it fired on a large share of all one-minute bars — far too often to say anything meaningful on its own about the end of a move. A signal that is lit up on a large fraction of bars cannot, by itself, tell you that *this* bar is special.

When exhaustion is worth noticing

Exhaustion becomes interesting only when **context** makes the bar unusual:

- it happens **at an important reference level** — the previous day's high (PDH), the overnight high (ONH), the value area high (VAH);
- it comes **after an extended move**, not in the middle of rotation;
- the bar carries **meaningful total volume** — a New York bar, not a 40-contract Asian-session bar;
- and ideally the next bars show **no follow-through**: price does not come back to test the high.

Even then, exhaustion raises the possibility of a **pause**, not the certainty of a reversal. Aggressive buyers stopping is not the same thing as aggressive sellers arriving.

Key idea: Thin volume at a bar's extreme is normal. Exhaustion is worth attention only when the location, the prior move, the bar's volume and the following bars make it unusual — and even then it suggests a pause, not a reversal.

Common mistake: "Exhaustion on the footprint means price will reverse." This is a popular claim, not proven. A related mistake is confusing exhaustion with absorption: exhaustion is aggression *stopping*; absorption is aggression *continuing but being soaked up* by passive orders. They look different on the footprint and mean different things.

Finished vs Unfinished Auction at Bar Extremes

In Chapter 7 you met poor highs and poor lows on a session profile. The same idea exists at the scale of **one footprint bar**, and it relies on the edge rows that had no diagonal partner in Chapter 8.

Finished auction

At the top row of a bar, look at the Bid side. If it is **zero** — for example 0×5 — then at the high, only aggressive buyers traded; no aggressive seller hit a bid at that price. The market pushed up until buying stopped, and then it turned. That high is called a **finished auction**. At the low, the mirror case is an Ask of zero, such as 7×0 .

Unfinished auction

If **both sides traded** at the extreme — say the top row is 12×9 — then even at the highest price, some aggressive sellers were hitting bids there. For a trade to happen *on the bid* at the high, the offer must at that moment have been one tick higher, and nobody bought it. The auction in that direction did not reach a clean end. That is called an **unfinished auction**, and the common interpretation (found in several vendors' educational material) is that price may return later to "finish" it.

	Top row (Bid $\times$ Ask)	Label
Bar A	0×5	Finished high
Bar B	12×9	Unfinished high
Bar C	1×2	Unfinished under the simple rule — but is it meaningful?

Two definitions, two answers

Bar C shows the problem. The **simple definition** says an extreme is unfinished if both sides are greater than zero. Under that rule, 1×2 — three contracts in total — counts as unfinished. In practice, that is noise.

A **threshold definition** requires both sides to be meaningful. For example (illustrative only): the extreme is unfinished only if the *smaller* of its two numbers is at least 10 contracts. Under that rule, Bar B (12×9) is finished, because 9 is below 10, and Bar C is clearly finished. Raise or lower the threshold and labels change.

Neither rule is "correct". The point is that "unfinished auction" means different things in different tools, and the same row can be labelled finished by one and unfinished by another. You will see an example of exactly this in the walkthrough at the end of the chapter.

"They always get revisited" — and the baseline

The popular claim is that **unfinished auctions almost always get revisited**. This is a popular claim, not proven. Even vendors who teach the concept note that there is no guarantee, and that in strong

trends such levels can be left behind.

There is also a statistical trap. An unfinished high on a one-minute bar is usually only a few ticks from the current price. Price wanders a few ticks back and forth constantly. So "the unfinished high got tested in the next few bars" will happen very often **by chance alone**. To claim the pattern adds anything, you would have to compare it with a baseline: how often does price return to *any* price the same distance away, in the same time? If the two rates are similar, the unfinished label added nothing.

Key idea: Finished = one side is zero at the extreme. Unfinished = both sides traded there. An unfinished extreme is a possible reference level, not a target. Definitions differ, and "it got revisited" must be compared with how often a random nearby price gets revisited.

Common mistake: Labelling a 1 × 2 row as an unfinished auction, and counting the "revisit" of an unfinished high that was only one tick away as a success for the pattern.

Absorption: Three Definitions, One Word

Absorption at its core means: **passive limit orders soak up the other side's aggression, and price does not move despite the aggression**. A classic description: lots of aggressive selling hits the bid, but the bid keeps getting refilled at the same price and price does not go lower. Delta becomes more and more negative, yet price stays put. Sometimes an iceberg order — an order that shows only a small part of its size and refills as it is consumed — sits behind this, but not necessarily.




Figure 24. Absorption: large volume trades at one area but price makes little progress, because passive orders soak up the aggressive ones.

The trouble is that the word "absorption" is used for at least three different things, with very different quality of evidence.

Definition 1: the order-book definition (most direct)

You watch the order book (DOM or liquidity heatmap, or full order-by-order data called MBO, "market by order") and see a resting limit order at one price being consumed again and again and refilled, while price fails to pass through it. This is the most direct evidence, because you can see the passive orders themselves. **A footprint alone cannot show this**, because a footprint does not show resting limit orders — only the trades that hit them. Chapter 10 covers these tools.

Definition 2: the footprint proxy (indirect)

On the footprint, you see **heavy volume and strongly one-sided delta in one zone, with little price progress**. For example, large Bid numbers (aggressive selling) in the bottom rows of a bar, but the bar closes well above its low. Here is a hypothetical case:

Price (GC)	Bid × Ask	Row delta	Note
4205.6	14 × 22	+8	close
4205.5	26 × 31	+5	
4205.4	38 × 29	−9	
4205.3	71 × 24	−47	open
4205.2	148 × 41	−107	bar POC
4205.1	119 × 18	−101	low

Aggressive sellers traded 267 contracts in the bottom two rows, yet price never got below 4205.1 and the bar closed three ticks above its open, five ticks above its low. Something absorbed that selling — very possibly resting bids. But the footprint cannot confirm it. The same picture could also be the end of a selling wave that simply ran out, with buyers stepping in afterwards.

This is a **proxy**: indirect evidence that is consistent with absorption, not proof of it.

Definition 3: candle shape (weakest)

"A long wick plus high volume equals absorption." This definition calls almost every high-volume pin bar absorption. When we looked at candle-shape proxies of this kind, they fired on many bars and were rarely precise. They are the weakest of the three.

Evidence level	What you observe	What it can confirm
Order book / MBO	Resting orders consumed and refilled at a price	Passive orders actually present and refilling
Footprint proxy	Heavy one-sided aggression, little price progress	Effort without result — consistent with absorption
Candle shape	Long wick on high volume	Very little on its own

Presence is not outcome

Even with the best evidence, absorption shows the **presence** of a large passive participant. It does not tell you the **outcome**. Whether that level holds or breaks depends on how much size the passive trader still has left and how long they want to defend the price — and that is invisible. Absorption can, and regularly does, break when the passive size runs out. A common teaching guideline is not to act on absorption alone, but to wait and see whether the other side actually starts participating (for example, aggressive buyers appearing after absorbed selling).

Key idea: Absorption = passive orders absorbing aggression while price stalls. The order book confirms it most directly; the footprint offers an indirect proxy; candle shape alone is weak. Always ask which definition a source is using. Absorption shows presence, not outcome.

Common mistake: "Absorption at the low means a safe buy." This is a popular claim, not proven. Also avoid calling every high-volume pin bar "absorption", and assuming the footprint can directly show iceberg orders.

Effort vs Result: Wyckoff on the Footprint

Richard Wyckoff, an early twentieth-century market analyst, described several "laws" of market behaviour. One of them is the law of **effort versus result: volume is effort; price movement is result**.

When they agree — heavy volume and a large move in the same direction — the move is "in harmony". When they disagree — very heavy volume or delta but little price movement — someone on the other side was neutralising the effort. Often that is the absorption you just met.

A sharper measure on the footprint

A candle chart only has total volume. The footprint gives a sharper measure of effort: **delta**, the net aggression of one side. A simple ratio compares effort with result:

Efficiency = |close – open| (in ticks) ÷ |bar delta|

That is, how many ticks of movement each contract of net aggression bought.

	Bar delta	Open → close (GC)	Ticks moved	Efficiency (ticks per contract)
Bar X	+200	4210.0 → 4211.2	12	$12 \div 200 = \mathbf{0.060}$
Bar Y	+210	4214.8 → 4214.9	1	$1 \div 210 = \mathbf{0.005}$

Both bars had almost the same net aggressive buying. Bar X turned it into twelve ticks (\$1.20 on the price, \$120 per GC contract). Bar Y turned it into one tick. If Bar Y printed right at the overnight

high, you would describe it as **big effort, small result** — aggressive buyers did not move price, which suggests passive sellers were present. Some traders call this an **aggression failure**.

What counts as "low" efficiency? There is no universal number. You would need to know what is typical for that session: measure the efficiency of ordinary bars in the same session on your own data, and look for bars far below that norm.

Context decides

At a major resistance level, effort without result plausibly reflects passive sellers defending the level. In the middle of a quiet range, the same numbers may be nothing more than normal back-and-forth trading where buyers and sellers took turns.

Limitations

1. **The ratio is sensitive to close minus open.** A bar with a large range that closes near its open has an efficiency near zero, even if it travelled a long way in both directions.
2. **Normalise by session.** A delta of 100 in the Asian session and a delta of 100 at the New York open are different amounts of effort relative to what is normal.
3. **It is a reading rule, not a formula that guarantees anything.** The same idea underlies Volume Spread Analysis (VSA), where it is also a way of reading bars, not a mechanical edge.

Key idea: Effort (volume, delta) should produce result (price movement). When it does not, someone was on the other side. Measure "low" against the same session's norm, and interpret it only at meaningful locations.

Common mistake: "High volume and a small bar always means reversal." Not true. Also avoid comparing raw delta across sessions, and ignoring that a bar can have a big range but a small body.

Trapped Traders

Trapped traders are traders who entered aggressively at a price and then watched price move straight against them. On the footprint, the picture for **trapped buyers** is: strong buying aggression — a buy imbalance or a buy stack, and strong positive delta — in the **upper part** of a bar, but the bar **closes below that zone**, near its own low. Anyone who bought aggressively in that zone is now holding a loss.

Here is a hypothetical example (diagonal imbalances at 3:1):

Price (GC)	Bid × Ask	Ask(p) ÷ Bid(p-1)	Note
4218.4	8 × 31	31 ÷ 17 = 1.8	high
4218.3	17 × 92	92 ÷ 22 = 4.2 → buy	stack
4218.2	22 × 104	104 ÷ 25 = 4.2 → buy	stack
4218.1	25 × 88	88 ÷ 29 = 3.0 → buy	stack
4218.0	29 × 40	40 ÷ 44 = 0.9	open
4217.9	44 × 21	21 ÷ 51 = 0.4	
4217.8	51 × 18	18 ÷ 47 = 0.4	close
4217.7	47 × 9	—	low

There is a three-row buy stack at 4218.1–4218.3 and a bar delta of +160 (403 Ask against 243 Bid), yet the bar closed at 4217.8, three ticks below the bottom of the stack and in the lower part of its range. The buyers who lifted the offer at 4218.1–4218.3 are underwater.

The possible logic

If price keeps moving lower, some of those buyers may be forced out — by choice or by stop orders — and closing a long position means **selling**, which can add to the downward pressure. And if price comes back up to their entry zone, some may sell just to get out at break-even, which can make that zone act as resistance.

A made-up teaching definition

For research, you need precision. An illustrative definition of trapped buyers (numbers are for teaching only):

- a buy stack whose **lowest row is above the bar's close**;
- **positive bar delta**;
- the close sits in the **lower third** of the bar's range.

The example above passes: the stack's lowest row (4218.1) is above the close (4217.8); delta is +160; the range is 4217.7–4218.4 (seven ticks), and the close is one tick above the low. The mirror image describes trapped sellers.

Important limitations

1. **"Trapped" is an assumption about intent.** The footprint cannot tell whether those aggressive buys were new long positions or short sellers closing out. If they were shorts covering, nobody is trapped — they are now flat.
2. **Trapped traders may be hedged**, or may simply be willing to hold through a larger loss.

3. **The pattern is only known when the bar closes**, and the very next bar can reverse straight back up.
4. **The same footprint shape is common in unimportant places**, such as the middle of an Asian-session range. The idea carries most weight at a **failed breakout of a reference level** — for example, a push above the previous day's high that is rejected back below it.

Key idea: Trapped traders = strong aggression in a zone, then a close on the other side of it. The logic is that forced exits may add fuel and the entry zone may become a reference. But intent is invisible, so "trapped" is always an assumption.

Common mistake: "Trapped traders always cause a strong move the other way." This is a popular claim, not proven. Also avoid assuming every aggressive buy was a new position, and seeing a "trap" in every bar with an upper wick.

Delta-POC and Its Shift

In Chapter 8 you met the **volume POC** of a bar: the row with the most **total** volume, the place of greatest agreement. A second, related measure is the **delta-POC**: the row with the largest **absolute** row delta, $|\text{Ask} - \text{Bid}|$. It marks where **one-sided aggression** was most concentrated.

The two can sit at very different prices. The volume POC may be a row where both sides traded heavily and delta was close to zero; the delta-POC is the row where one side clearly dominated.

Price (GC)	Bid × Ask	Row volume	Row delta	
4212.5	12 × 20	32	+8	
4212.4	18 × 71	89	+53	← delta-POC
4212.3	64 × 70	134	+6	← volume POC
4212.2	45 × 38	83	−7	
4212.1	22 × 15	37	−7	

At 4212.3, both sides traded heavily — agreement. At 4212.4, aggressive buyers clearly dominated — aggression.

Delta-POC shift

The **delta-POC shift** compares this bar's delta-POC with the previous bar's, in ticks. Read the **sign** of the delta at the delta-POC as well as its position.

Bar (1-min)	Delta-POC (GC)	Delta at that row	Shift
10:02	4211.8	+41	—
10:03	4212.1	+47	+3 ticks
10:04	4212.4	+53	+3 ticks
10:05	4212.6	−38	+2 ticks

From 10:02 to 10:04, the concentration of aggressive buying moved higher: traders were willing to buy urgently at rising prices. At 10:05 the delta-POC moved up again but its sign turned negative — the strongest aggression in that bar was now **selling**, at higher prices. That is a description of where aggression concentrated and how it moved. It can sit alongside POC migration (Chapter 8) and cumulative delta.

Honest limitations

1. **In thin bars the delta-POC jumps between rows**; large shifts may be pure noise.
2. **Ties decide the output.** When two rows have almost the same $|\text{delta}|$, the tie-break rule your software uses (for example "pick the lower price") determines the delta-POC — the market did not.
3. **It is a derived metric.** It does not come from the classic auction or Wyckoff literature, and there is no public evidence that it predicts anything. Treat it as a descriptive column.

Key idea: Volume POC = where most trading happened (agreement). Delta-POC = where one side was most aggressive (aggression). Its shift from bar to bar, with its sign, describes how aggression is moving.

Common mistake: Confusing delta-POC with volume POC; reading meaning into one- or two-tick shifts in thin bars. "Delta-POC shift predicts the next bar's direction" is a popular claim, not proven.

CVD Divergence — and Why It Often Fails

Cumulative volume delta (CVD), from Chapter 3, is the running total of delta from a starting point called the **anchor** — for example the start of the session. **CVD divergence** means price and CVD disagree: for example, price makes a **higher high** while CVD makes a **lower high**.

Swing	Price high (GC)	CVD at that high
First high	4215.0	+1,850
Second high	4215.6	+1,420

Price went six ticks higher on the second push, but the running total of aggressive buying minus aggressive selling was 430 contracts lower. The classic interpretation: "price rose, but fewer aggressive buyers pushed it there, so the move is weak."

That sounds compelling. On a chart in hindsight, the examples that "worked" look beautiful. In practice it fails often, for at least five mechanical reasons.

Reason 1: passive orders move price too

CVD counts only aggression. If sellers with resting limit orders raise their offers or pull them, price can rise on very little aggressive buying. The higher high is real; the lower CVD just means it took less aggression to get there.

Reason 2: absorption can last a long time

There is no law that says a divergence must "resolve". A large passive participant can keep absorbing aggression for a long time, and the market can stay divergent while price continues in the same direction.

Reason 3: anchor dependence

CVD depends on where you start counting. If both swing points come after the anchor, moving the anchor earlier or later shifts the whole CVD line up or down by the same amount, and the comparison between the two highs does not change. But traders very often compare points that sit on **different sides** of an anchor — a New York high against an overnight high, today against yesterday — and many platforms reset CVD each session. Then the anchor decides whether a divergence exists at all. A divergence on a session-anchored CVD may not exist on a weekly-anchored one.

Reason 4: frequency

On short timeframes, small divergences form all the time. The ones that "worked" and the ones that did not **look identical at the moment they form**. You only know which kind it was afterwards.

Reason 5: data quality accumulates

Bid/ask classification errors in fast moments (Chapter 8) are small for any one bar, but CVD **adds them all up**. Over a long session, those small errors accumulate into the running total.

A better use

Some of the best-known order-flow educators state plainly that CVD is not a tool for rule-based signals. A more useful approach is to treat a divergence **at a reference level** as one piece of the puzzle — for example, a second test of the overnight high on weaker CVD — and to read it together with acceptance or rejection at that level and with the footprint at that moment.

Key idea: CVD divergence is price and running delta disagreeing. Passive orders can move price without aggression, divergences need not resolve, anchors matter when comparing across them, and divergences that fail look identical to ones that work when they form. Use it at levels, alongside other evidence.

Common mistake: "CVD divergence is a reliable reversal signal." This is a popular claim, not proven. Also avoid comparing CVD highs and lows with different anchors on different days, and remembering only the divergences that worked — a form of survivorship bias when reviewing charts.

Base Rates and False Positives: How to Judge Any Footprint Pattern

Nearly every footprint pattern "works" when you scroll back through old charts. That is not because the pattern works; it is because our memory and our eyes are drawn to the examples that fit. To judge a pattern honestly, you need three numbers.

Number 1: fire rate

On what share of bars does the pattern appear? The higher the fire rate, the more the pattern is just a description of an ordinary bar rather than a special event. Remember the exhaustion lesson: a rule that is true for a large fraction of bars cannot, by itself, single out a special moment.

Number 2: precision

When the pattern appears, how often does the claimed outcome actually follow? The outcome must be defined precisely and in advance — for example, "price moves 5 ticks against the prior direction before it moves 5 ticks further in the same direction, within the next 10 bars."

Number 3: baseline

How often does that same outcome follow any bar — or a random bar? This is the number most people never compute, and it is the one that matters most. Precision without a baseline is meaningless.

A worked example with invented numbers

These numbers are made up purely to illustrate the arithmetic. Suppose you study 1,000 one-minute bars:

	Outcome happened	Outcome did not happen	Total
Pattern fired	143	117	260
Pattern did not fire	392	348	740
Total	535	465	1,000

- **Fire rate** = $260 \div 1,000 = 26\%$ of bars.
- **Precision** = $143 \div 260 = 55\%$. "The pattern is right 55% of the time!"
- **Baseline** = $535 \div 1,000 = 53.5\%$ for all bars (and $392 \div 740 \approx 53\%$ for bars where the pattern did *not* fire).

The pattern adds about one or two percentage points over doing nothing at all. That is very little information — and with a sample like this it may not even be distinguishable from chance. Then add the spread and commissions: a tiny edge before costs can easily become nothing after costs. A screenshot of the 143 successes would look convincing. The table shows what it is really worth.

Common traps

- **Conflicting detectors.** One bar can trigger a bearish pattern and a bullish pattern at the same time — you will see exactly this in the walkthrough at the end of this chapter. If you treat both as signals, any outcome "confirms" one of them.
- **A loose outcome definition.** "Price reacted" is true almost every time. If the outcome is not defined precisely before you look, the pattern cannot fail.
- **Choosing settings after seeing results.** Adjusting the ratio or threshold until the pattern fires on the nice examples is overfitting (Chapter 12).
- **Ignoring costs.** Even precision slightly above the baseline may leave nothing after spread and commissions.

The practical conclusion

Read the footprint with the question "**at this reference level, what did aggression do?**" — not with the instruction "**pattern X fired, act on it**".

Key idea: Judge any pattern with three numbers: fire rate, precision and baseline. A high fire rate means little information; precision without a baseline means nothing; define the outcome first and never tune settings after seeing results.

Common mistake: Believing claims like "this pattern has an 80% win rate" when nobody has defined the outcome, the baseline or the data. That kind of claim is a popular one, not proven. Reviewing old charts and counting only the successes, or judging a pattern from a handful of screenshots, are the same mistake.

Try it: Pick one pattern from this chapter and write a precise definition and a precise outcome on paper. Before you look at any results, also write how you will measure the baseline. Then go through one week of GC one-minute bars and fill in the 2×2 table yourself.

Common Footprint Mistakes

Most footprint mistakes come from a single root: **reading detail without context**. Here are ten of the most common, each with a practical check that replaces it.

#	Mistake	The check that replaces it
1	Reading the footprint with no map: interpreting imbalances in the middle of a range, far from any reference level	Ask: which level from my pre-session map (Chapter 7) is this bar next to?
2	Reading thin bars with the same seriousness as New York bars	Compare the bar's volume with the average for that session
3	Reading the footprint on MGC instead of GC	Read the footprint on GC, even if you trade MGC or XAUUSD
4	Using horizontal imbalances without realising it	Confirm your platform is set to diagonal
5	Changing settings after seeing the outcome	Write your settings down and keep them fixed
6	One word, several definitions (absorption, unfinished auction)	Write down the exact definition you are using
7	Treating delta as direction: "positive delta means price must rise"	Check effort against result
8	Treating divergence as reversal	Check the anchor and the location
9	Ignoring fire rate: anything that lights up on a large share of bars is not news by itself	Count how often it fires
10	Staring at one bar in isolation	Always read a sequence of three to five bars and the market's response in the next bar

There is also a behavioural mistake: **overtrading** because the footprint seems to show a "signal" in every bar. A busy gold footprint shows dozens of small events per minute. Most of them are noise. The footprint is most useful when you bring a **small number of specific questions to a small number of important places** in the day — the touch of yesterday's high, the test of the value area low, the break of the initial balance.

Key idea: Context first, footprint second. Compare bar volume with the session, read GC with diagonal imbalances, fix your settings and definitions, separate delta from direction and divergence from reversal, count fire rates, and read sequences rather than single bars.

Common mistake: "The more I watch the footprint, the better I trade." More watching tends to produce more imagined signals. Another: "The footprint alone is enough to trade." That is a popular claim, not proven. A third: copying a video creator's settings, built for a different market, onto gold.

Try it: Take one 30-minute stretch of GC. First look at the footprint with all its colours and count how many "signals" you think you see. Then mark only five reference levels on the chart and look at the footprint *only* at the moments price touches one of them. Compare how many questions you actually needed to answer.

Full Walkthrough: Reading One Gold Footprint Bar

Now let us read a complete bar in a fixed order, using everything from Chapters 7, 8 and 9. The bar is **hypothetical** — a one-minute GC bar with symbolic prices — and the numbers are designed to teach, not to represent a typical session.

- Open = 4210.4, High = 4211.0, Low = 4210.2, Close = 4210.6 (roughly 4180–4184 on XAUUSD, using the approximate gap).
- Teaching settings: diagonal imbalances, ratio 3:1, minimum volume 20.

Price (GC)	Bid	Ask	Total	Row Δ	Note
4211.0	0	5	5	+5	High: Bid = 0 → finished
4210.9	12	22	34	+10	
4210.8	18	64	82	+46	Buy imbalance (64 ÷ 20 = 3.2)
4210.7	20	71	91	+51	Buy imbalance (71 ÷ 22 = 3.2)
4210.6	22	81	103	+59	Buy imbalance (81 ÷ 26 = 3.1) · delta-POC · close
4210.5	26	30	56	+4	
4210.4	110	52	162	-58	Volume POC · Sell imbalance (110 ÷ 30 = 3.7) · open
4210.3	60	24	84	-36	
4210.2	9	4	13	-5	Low: both sides > 0

Step 1: context first

Where is this bar? Is it sitting on the previous day's low, testing the overnight high, at the value area edge — or in the middle of nowhere? Without this, everything that follows is only half meaningful. Write the location down before reading a single number.

Step 2: totals

- Volume = 630 contracts.
- Ask total = 353; Bid total = 277; bar delta = 353 – 277 = **+76**.
- The bar is green: open 4210.4, close 4210.6, two ticks higher.

Step 3: POC and delta-POC

- **Volume POC** = 4210.4 (162 contracts), in the **lower** part of the bar. The main agreement happened low.
- **Delta-POC** = 4210.6 (+59). But look at 4210.4: -58. The two largest one-sided rows are almost a tie in size (59 versus 58), with **opposite signs**. Had one contract been classified differently, the delta-POC would have been the selling row instead. This is the tie problem from earlier in the chapter in action.

Step 4: imbalances

Check every row diagonally:

- **Sell imbalance at 4210.4:** $\text{Bid}(4210.4) = 110 \div \text{Ask}(4210.5) = 30 = 3.7$. Heavy aggressive selling at the bar's busiest price.
- **Buy imbalances at 4210.6, 4210.7 and 4210.8:** $81 \div 26$, $71 \div 22$ and $64 \div 20$, all just above 3. That is a **three-row buy stack**.
- No other row qualifies: for example 4210.9 is $22 \div 18 = 1.2$, and 4210.3 is $24 \div 9 = 2.7$ (below 3).

Step 5: extremes

- **High** 0×5 : no aggressive seller traded at 4211.0 → **finished** under any definition.
- **Low** 9×4 : both sides traded. Under the simple rule (both > 0) it is **unfinished**. Under the illustrative threshold rule from earlier (the smaller side must be at least 10), it is **finished**, because 4 is below 10. **The same row, two labels.**

Step 6: exhaustion at the top?

Apply the made-up exhaustion definition from the start of this chapter, and suppose this bar did make the highest high of the last several bars:

- Top row volume 5 versus POC 162: far below one-fifth. ✓
- Top two rows' Ask = $5 + 22 = 27$, out of 353 total Ask: about 7.6%, below 15%. ✓
- Close 4210.6 is four ticks below the high, in an eight-tick range: half the range, more than a third. ✓

So the **bearish exhaustion** pattern fires.

Step 7: absorption at the bottom?

Now apply an equally made-up footprint absorption proxy: "at least a fifth of the bar's total volume is aggressive selling in the bottom third of its rows, and the bar closes in the upper half of its range."

- The bottom third of the nine rows is 4210.2–4210.4. Aggressive selling there = $9 + 60 + 110 = 179$, out of 630: about 28%. ✓
- Close 4210.6 is exactly at the midpoint of 4210.2–4211.0 — just inside the "upper half" if the midpoint counts. ✓

So the **bullish absorption proxy** fires too — on the same bar.

Step 8: trap, and effort versus result

- **Trapped buyers?** The buy stack's lowest row is 4210.6, which is the close itself, not above it. The illustrative trap definition requires the stack to sit above the close, so **no trap**.
- **Effort versus result:** two ticks of movement on +76 delta gives an efficiency of $2 \div 76 \approx 0.026$ ticks per contract. Whether that is low depends on what is normal for this session; with a

modest delta and a modest move, nothing here stands out.

The honest story

One bar, **two patterns pointing in opposite directions**: bearish exhaustion at the top and bullish absorption at the bottom. If you pick whichever one suits your existing opinion, you will always find "confirmation".

The honest reading puts the numbers into a story without forcing a direction:

Sellers were aggressive at 4210.4 — the busiest price in the bar — but price did not go lower. Buyers answered with a three-row stack from 4210.6 to 4210.8, but at 4211.0 nobody followed them. This was two-sided rotation inside one minute.

Anything more than that has to come from **outside the bar**: the reference level it sits on (Step 1), and what the **next bars** do. Does price hold above the stack and come back to retest 4211.0, or does it slide back to the 4210.4 selling and through it? The footprint has told you who was aggressive, where and how hard. It has not told you what happens next.

Key idea: Read every bar in the same order: context → totals → POC and delta-POC → imbalances → extremes → effort and result → story. Expect conflicting patterns. The story comes from the bar plus its location plus the bars that follow.

Common mistake: Starting from the imbalances instead of the context; picking the one "signal" that matches your bias; and forgetting that a symbolic teaching bar has different scale from real GC bars in different sessions.

Chapter summary

- Footprint patterns are vocabulary for describing what happened, not standalone triggers. On their own, they did not reliably predict direction after costs in our own testing.
- **Exhaustion** (thin volume at the extreme) describes the normal shape of most bars; it is worth noticing only with context — level, prior move, volume, follow-through — and suggests a pause, not a reversal.
- **Finished** extremes have zero on one side; **unfinished** extremes have both sides trading. Definitions differ (simple vs threshold), and "revisits" must be compared with a baseline.
- **Absorption** has three meanings: order-book refills (direct), footprint proxy (indirect), candle shape (weak). It shows presence, not outcome, and it can break.
- **Effort vs result** compares delta with movement; low efficiency must be judged against the same session's norm and read at meaningful locations.
- **Trapped traders** = strong aggression in a zone, then a close on the other side. Intent is invisible, so "trapped" is an assumption; it matters most at failed breakouts.

- **Delta-POC** marks the most one-sided row; its shift describes how aggression moves. It is a derived, descriptive metric with tie and noise problems.
- **CVD divergence** fails often: passive orders move price, absorption lasts, anchors matter when comparing across them, failing and working divergences look identical when they form, and classification errors accumulate.
- Judge any pattern with **fire rate, precision and baseline**, with outcomes defined in advance and costs included.
- Most mistakes come from reading detail without context. Read sequences, at a few important places.
- A single bar can trigger opposite patterns at once. The honest reading is a story, completed by the level and the next bars.

Checklist

- Before reading a bar, I write down which reference level it is next to.
- I compare the bar's volume with the average for its session before taking any pattern seriously.
- I have written, precise definitions for each pattern I use, and I know which definition a source uses before I trust its claims.
- I do not call a 1×2 extreme "unfinished", and I compare revisits of nearby prices with a baseline.
- I separate exhaustion (aggression stopping) from absorption (aggression soaked up).
- I measure effort versus result against the same session's typical efficiency.
- I treat "trapped traders" as an assumption about intent, not a fact.
- I check delta-POC ties and ignore one- or two-tick shifts in thin bars.
- I check the CVD anchor before calling a divergence, and I look at divergences that failed as well as those that worked.
- For any pattern, I know (or plan to measure) its fire rate, precision and baseline, with costs.
- I read bars in a fixed order and expect conflicting patterns.

Quiz

1. A footprint "exhaustion" rule fires on a large fraction of all one-minute bars. How is it best used, and why?
2. The top row of a bar reads 0×7 (bid $\times$ ask). Is this high finished or unfinished? What about a low row reading 9×4 under the simple "both sides > 0 " rule?
3. Which data source can most directly confirm absorption — orders being repeatedly refilled at one price — and why can the footprint alone not confirm it?
4. Bar delta is +300, but the bar closes only one tick above its open, right at a major resistance level. What would Wyckoff's law call this, and what is one honest caution about interpreting it?

5. A pattern is followed by a 5-tick reversal 55% of the time. Random bars are followed by the same reversal 53% of the time. How much does the pattern add, and what else must you consider?

Quiz answers

1. As **context at important levels**, not as a standalone signal. Thin volume at a bar's extremes is the normal shape of most bars, so a rule that fires that often describes ordinary bars and carries little information on its own. It becomes interesting only with location, a prior extended move, meaningful volume and lack of follow-through.
2. The 0×7 high is **finished** (no aggressive selling at the high). The 9×4 low is **unfinished** under the simple rule, because both sides traded — although under a stricter threshold definition it could be labelled finished.
3. **Order-book data** (the DOM, a liquidity heatmap, or full market-by-order data). Absorption is about resting limit orders being consumed and refilled; the footprint shows only executed trades, not the resting orders themselves, so it can offer only an indirect proxy.
4. **Effort without result** (high effort, small result — sometimes called an aggression failure). Cautions: judge "low efficiency" against the same session's norm; the ratio is sensitive to close minus open; and it suggests the presence of passive sellers, not a guaranteed reversal.
5. **Very little information** — about two percentage points over the baseline, which may not even be distinguishable from chance with a modest sample. You must also consider how the outcome was defined, how many bars were tested, whether settings were tuned after seeing results, and whether anything is left after spread and commissions.

PART V • THE ORDER BOOK

10

Liquidity, the Order Book and the Tape

In this chapter: We go one layer beneath candles and footprints, into the machinery that produces every print: the order book of waiting orders, the queue that decides who is filled first, and the tape that records every trade. You will learn the three levels of market data (L1, L2, L3), how to read a DOM and a liquidity heatmap, what walls, pulling and stacking are, why visible liquidity can mislead in both directions, how one large order becomes many small prints, what sweeps, stop runs and icebergs really are, why spoofing is a crime and not a strategy, and what the research on order flow imbalance does and does not show. The running theme: displayed liquidity is a promise, not a commitment, and much of this is harder to detect than social media suggests.

The Order Book: L1 vs L2 vs L3 Explained

Every trade you have seen in this book so far, every candle, every delta box, every footprint cell, was born at the same place: the moment an aggressive order met a waiting order. This chapter is about those waiting orders and about the record of their meetings.

What the order book is

The **order book** is the exchange's list of all **limit orders** that have been submitted but not yet executed. A limit order, as you saw in earlier chapters, is an instruction to buy or sell at a specific price *or better*, and it waits until someone agrees to trade with it. Buy limit orders sit on the **bid** side of the book. Sell limit orders sit on the **ask** side (also called the **offer**). Orders that wait are called **passive**.

A **market order** does not wait. It is an instruction to trade *now* at the best price available, and it does so by hitting the passive orders already in the book. The trader who sends it is the **aggressor**. A trade happens only when an aggressor meets a passive order. If nobody is aggressive, nothing trades, no matter how many orders are waiting.




Figure 25. A DOM ladder shows the resting size at each price. It changes every millisecond as orders are added, cancelled and filled.

Market data about the order book is usually described at three levels of detail.

Level 1: the top of the book

Level 1 (L1), also called **top of book**, shows only four things: the best (highest) bid, the best (lowest) ask, the quantity resting at each of those two prices, and the last trade. A typical retail quote screen, and the raw material for an ordinary candle chart, is L1 plus the stream of trades.

On gold futures, an L1 snapshot might read: best bid 4,240.0 for 14 contracts, best ask 4,240.1 for 9 contracts, last trade 4,240.1. (On an XAUUSD chart the same moment would sit roughly \$27–30 lower, around 4,210–4,213; the gap varies and is always approximate.)

Level 2: market by price

Level 2 (L2), also called **market by price (MBP)** or **depth**, shows the total quantity waiting at *several* prices above and below the market. This is the data behind the DOM ladder of the next section.

The crucial word is *total*. If L2 shows 60 contracts on the bid at 4,239.5, you do not know whether that is one order for 60 contracts, sixty orders for one contract each, or anything in between. CME's classic market-by-price feed publishes up to ten price levels on each side, together with the *number* of orders at each level, but not the size of each individual order.

Level 3: market by order

Level 3 (L3), known on CME as **market by order (MBO)**, shows every individual order: its anonymous identifier, its size, its price and its place in the queue. CME's MBO feed covers the full

depth of the book, not just ten levels, and lets you follow each order as it is added, changed, cancelled or filled.

Why does the distinction matter? Because many ideas in this chapter, such as your position in the queue, native icebergs, and telling a cancellation apart from a fill, can only be examined properly with L3. With L2 you are inferring; with L3 you are reading the event log.

Level	What it shows	Question it answers
L1 (top of book)	Best bid, best ask, their sizes, last trade	"What is the price right now?"
L2 (market by price)	Total size at each of several prices	"How much is waiting near the market?"
L3 (market by order)	Every order: anonymous ID, size, price, queue order	"How many orders, how big, in what sequence, and what happened to each?"

What no level shows you

Even full MBO data has blind spots, and beginners often assume otherwise.

- **Stop orders** are invisible until they trigger. A stop is an instruction held by the exchange (or your broker) that becomes an active order only when price reaches a trigger. Until then it is not in the book at all.
- **The hidden part of an iceberg order** is not displayed (more on this later in the chapter).
- **Implied orders**, which the exchange synthesises from spread markets between contract months, appear in the price-level data but not in the order-by-order feed.
- **Intent changes instantly.** The book is a snapshot of what people are willing to do this millisecond. It can be completely different a millisecond later.

Key idea: The order book is intention. A trade is fact. L1 tells you the price, L2 tells you *how much* is waiting, L3 tells you *how many orders, how big, and in what sequence*, but always anonymously.

Common mistake: "With Level 2 I can see where the institutions are sitting." L2 gives totals at each price, not identities, and not even the size of individual orders. Even L3 IDs are anonymous: you never learn *who* placed an order.

Common mistake: "I can see the stops on the DOM." You cannot. Stops do not exist in the visible book until they are triggered.

L3 data is also larger and more expensive, and not every trader needs it. What matters is knowing which level of data a claim actually requires.

How to Read a DOM Ladder

The ladder

DOM stands for **depth of market**. It is L2 data drawn as a vertical **price ladder**. Each row is one price. On gold futures (GC) and micro gold (MGC), each row is one **tick**, the minimum price increment of \$0.10. Recall the money value of a tick: \$10 per contract on GC (100 troy ounces) and \$1 per contract on MGC (10 troy ounces). Always re-check the current specifications on CME's website, because exchanges do change them.

The ladder has a bid column on one side and an ask column on the other. A number in the bid column is the total quantity of buy limit orders resting at that price. A number in the ask column is the total of resting sell orders.

Bid size	Price	Ask size
	4,240.4	21
	4,240.3	18
	4,240.2	25
	4,240.1	37
12	4,240.0	
19	4,239.9	
24	4,239.8	
31	4,239.7	

Approximate XAUUSD equivalent of this ladder: about 4,210–4,213; the gap between futures and spot CFD prices moves through the day.

Most platforms add more columns: **volume at price** (how much has traded at each price today or this session), sometimes **traded at bid / traded at ask** (how much was hit by aggressive sellers or lifted by aggressive buyers at that price), and **pulling/stacking** columns that show recent changes in resting size (covered two sections from now).

Reading it in three steps

Step 1: find the inside market. The **inside market** is the highest bid and the lowest ask; here, 4,240.0 bid and 4,240.1 ask. The distance between them is the **spread**. In the example it is one tick, worth \$10 per GC contract or \$1 per MGC contract. During active hours gold futures usually trade with a one-tick spread; during thin hours, or in the seconds before a major economic release, the spread often widens to two ticks or more.

Step 2: look at the shape of the depth. Is one side unusually heavy compared with the other? Is there one price with a much larger number than its neighbours? Hold that observation lightly: the

next sections explain why.

Step 3: watch how the inside numbers change when trading happens. When the 37 contracts at 4,240.1 shrink to 30, did they shrink because 7 contracts traded there (a fill), or because someone cancelled 7 contracts without any trade (a pull)? This question, *filled or pulled?*, is one of the most important in this chapter.

What it takes to move one tick

Here is the mechanism that makes the DOM useful. In the ladder above, the price can only print 4,240.2 after the entire 37 contracts on the ask at 4,240.1 are gone, either consumed by aggressive buyers or withdrawn by their owners. The DOM therefore shows you *how much work* it takes to move price. A thin ask side means little work; a thick one means more.

Key idea: Price moves up one tick only when the whole quantity on the best ask is used up or pulled. Price moves down one tick only when the whole quantity on the best bid is used up or pulled.

Limits of the DOM

The DOM shows only *now*. It has no memory of what was there a minute ago, which is exactly why the heatmap was invented. Numbers far from the current price are especially unstable, because orders there can be cancelled at no cost long before price arrives. And watching a fast gold DOM by eye for hours is tiring and invites perception errors: you start seeing patterns in noise.

Common mistake: "The side with more size on the DOM will win." This is a popular claim, not proven. Resting size can be cancelled in an instant, and a heavy side may simply attract aggressors who want to trade against it.

Common mistake: Confusing the DOM with the footprint. The DOM shows orders *waiting* to trade; the footprint shows contracts that *did* trade. They answer different questions.

Try it: Open a GC or MGC DOM in a simulator around the New York stock market open (09:30 New York time). For two minutes, watch only the inside market. Count how often the spread widens from one tick to two, and try to notice at least once when the best ask shrinks without any print on the tape.

Without a level you chose in advance, the DOM is mostly noise. Its value appears when price arrives at a place you already care about.

Market by Order (MBO): Seeing Individual Orders

Every order has a life story

With market-by-order data, each order has a lifecycle you can follow:

- **Add:** the order enters the book.
- **Modify:** its size or price is changed.
- **Cancel:** it is withdrawn.
- **Fill (execution):** it trades, fully or partially.

CME gives each order an anonymous **OrderID**, which stays the same for the order's whole life, and a **priority** value, which determines its place in the queue at its price and can change when the order is modified in certain ways (the next section explains which ways).

Filled or pulled: the key question

Suppose the L2 bid at 4,239.5 drops from 80 contracts to 20. Two very different stories fit that single fact:

1. Aggressive sellers sold 60 contracts at 4,239.5, and the resting buyers **absorbed** them. Buyers were genuinely present and did trade.
2. The resting buyers **pulled** (cancelled) 60 contracts. Nobody traded; buyers stepped back.

The first is buyers standing their ground; the second is buyers retreating. With L2 you can only guess which, from the timing of trades and book changes. With MBO you see exactly which orders were filled and which were cancelled. Imagine the same moment split by MBO into 35 contracts filled and 25 contracts cancelled: the reality was a mixture, which L2 alone could never reveal.

What else MBO tells you

MBO shows the *number and sizes* of orders at each price: one 60-lot or forty small orders. It lets you estimate the size of the queue in front of a specific order. And it carries the fingerprints of native iceberg orders, which you will meet later in this chapter. Commercial data vendors distribute CME's market data feed (called MDP 3.0) in MBO form with precise timestamps. Our own order-book indicator, AK FlowBook, is built on this kind of order-by-order data; in this book we use it only to illustrate concepts.

What MBO still cannot show

- **It is anonymous.** One large order is not necessarily one institution, and several orders may belong to the same trader.
- **Parent orders are hidden.** Large traders often use algorithms that slice one big intention (the **parent order**) into many small **child orders** over minutes or hours. MBO shows the children, never the parent.
- **It is heavy.** The volume of messages is enormous, and correctly rebuilding the book from them (handling event order, gaps and recoveries) is a technical skill in its own right. A bug in

the rebuild produces confident but wrong conclusions.

- **Implied liquidity is missing**, as noted above.

Common mistake: "MBO lets me see the institutions' hands." This is a popular claim, not proven. The data is anonymous and orders are sliced; better data does not reveal identities.

Common mistake: "Every drop in DOM size means trading happened." It may have been a cancellation.

Key idea: The most practical gift of MBO is simple: it separates *executed* from *cancelled*. Better data does not fix bad analysis, but it removes a whole class of guessing.

Queue Position and FIFO: Why Being First Matters

Who fills first?

When many limit orders wait at the same price and an aggressor arrives, the exchange's **matching engine** must decide which waiting orders get filled first. CME uses several **matching algorithms**. The most widely used is **FIFO (first in, first out)**: at a given price, the order that arrived first is filled first. Published data indicates FIFO accounts for the largest share of volume on CME Globex. Some other products use **pro-rata** matching (fills split in proportion to order size) or a hybrid. Check the official CME specification for the algorithm on any contract you trade; in this book we work with FIFO for gold.

Think of a queue at a ticket window: first to join, first served, however large your order.

Rules that keep or lose your place

On CME, priority follows a simple logic:

- **Reducing** the size of your order **keeps** your place in the queue.
- **Increasing** the size **loses** your place: the order goes to the back.
- **Changing the price loses** your place: at the new price you join the back of that queue.

Touched is not filled

Here is the consequence that surprises almost every new trader. Suppose your buy limit order is at 4,240.0 (roughly 4,210–4,213 in XAUUSD terms) and there are 300 contracts in the queue ahead of you. Price trades down to 4,240.0, 200 contracts trade there, and then price bounces. You were never filled, because 100 contracts were still ahead of you when the selling stopped.

A worked example: your order sits 140th in a 220-contract queue. Price touches your level once and 90 contracts trade; it touches again later and 30 more trade. In total 120 contracts traded at

your price, and you were still 20 contracts from the front. Nobody cheated you. You were simply not first in line.

Easy fills can be bad fills

The opposite case teaches something deeper. When your limit order fills instantly and completely, sometimes it is because the entire queue was consumed and price went straight through your level. You were filled precisely because the market was moving against you. This is called **adverse selection**: the fills that come easiest are often the worst ones, because they happen when the other side has a stronger reason to trade.

This is also why backtests that assume "price touched my limit, so I was filled" are optimistic. Chapter 12 returns to this in detail.

Key idea: At one price, first come, first served. Price touching your order is not the same as your order being filled.

Common mistake: "Price reached my limit and I wasn't filled, so the broker cheated." Usually the FIFO queue explains it completely.

Common mistake: Repeatedly nudging an order up and down by a tick "to be ready". Every price change sends you to the back of a queue.

Without MBO data, nobody knows your exact queue position. Platforms that display "queue ahead" are *estimating*, using assumptions such as "cancellations happen at the back of the queue" or "cancellations are spread proportionally". These are models, not measurements.

The Liquidity Heatmap: The Order Book With Memory

The DOM, recorded over time

The DOM shows only the present instant. A **liquidity heatmap** is the DOM recorded over time and drawn as a picture:

- The **horizontal axis** is time.
- The **vertical axis** is price.
- The **colour** of each point is the quantity of resting limit orders at that price at that moment. Warm or bright colours mean a lot of resting liquidity; cool or dark colours mean little.

On top of this background, most heatmaps draw the price line and **trade bubbles**: circles whose size shows the traded quantity and whose colour shows which side was the aggressor.




Figure 26. A heatmap paints the order book's history: lasting walls show as bright bands, pulled orders as bands that suddenly end.

Three basic patterns

Learn to recognise three things a bright band can do:

1. **Held.** A bright horizontal line persists for a long time. Someone, apparently, is willing to trade a lot at that price.
2. **Eaten.** Price reaches the band, trade bubbles appear, and the band fades. The liquidity really traded: it was filled.
3. **Pulled.** Price approaches and the band fades *without* trade bubbles. The liquidity was withdrawn.

The difference between patterns 2 and 3 is the same filled-or-pulled question from the MBO section, now visible as a picture.

Why the heatmap helps

The heatmap gives the order book a memory. You can see where liquidity keeps returning, where the book is thin (dark areas, where price may travel quickly because little is in the way), and what happened the last time price met a particular band.

Why the heatmap can mislead

- **It is usually built from L2**, so it inherits L2's ambiguity about fills versus cancels when both happen at the same price.
- **Colour settings change the story.** Contrast and threshold controls decide what looks "huge". A band that looks enormous on one screen may look ordinary on another simply because of the colour scale. Comparing two screenshots made with different settings is meaningless.

- **It is seductive.** Heatmaps are beautiful, and beautiful pictures invite stories. The risk of seeing meaning in randomness (storytelling bias) is higher here than on almost any other chart.

Common mistake: "Bright bands act like magnets and pull price toward them." This is a popular claim, not proven.

Common mistake: Treating a bright band as certain support or resistance. Liquidity can be pulled at the last moment.

Try it: Choose one colour setting for your heatmap, write it down, and never change it during a study period. Then, over a full New York session (around 08:20–13:30 New York time), mark one example each of a band that held, one that was eaten and one that was pulled.

Liquidity Walls, Pulling and Stacking

Definitions

A **liquidity wall** is a price in the book where the resting quantity is unusually large compared with the nearby prices, for example several times the typical size of the surrounding levels. On a heatmap it is a bright band; on the DOM it is one conspicuously large number.

"Unusually large" is always **relative**. Forty contracts might be a wall at 02:00 New York time during the quiet Asian session and completely ordinary at 10:00 during New York. Defining a wall by a fixed number for every hour of the day is a mistake.

Stacking means limit orders being *added* at one or more prices: resting size grows without any trade to explain it. **Pulling** means limit orders being *cancelled*: resting size shrinks without any trade. Many platforms show a pulling/stacking column next to the DOM that displays the recent net change at each price. With L2 data, that column sees only the *net* change: if trades and cancellations happen at the same price at the same time, separating them is an estimate.

A simple example: at 4,236.0 the bid shows 30 contracts. Over the next few seconds it rises to 90 while nothing trades there. That is +60 of stacking. A minute later it drops from 90 to 35, again with no trades printed at that price: −55 of pulling.

A reasonable (not certain) interpretation

If price is moving down toward a bid level and bids at and below that level are being **stacked**, waiting buyers are showing up. That *may* mean temporary support. If those same bids are being **pulled** as price approaches, buyers are stepping back, and price may pass through more easily. This is ordinary auction logic: price travels more easily in the direction that offers less resistance.

Why can this be meaningful at all? **Market makers**, firms that continuously quote both a bid and an ask and earn from the spread, and other short-term traders cancel their orders when they sense

risk. Pulling can therefore reflect real information or real fear.

Why most of it means little

These changes happen in milliseconds, and algorithms react far faster than any human eye. Much stacking and pulling is routine quote management that carries no "message" at all.

Common mistake: "A big wall means an institution intends to defend that price." This is a popular claim, not proven.

Common mistake: "When a wall is pulled, price must move in that direction." This is a popular claim, not proven.

Key idea: Stacking and pulling describe intention in the moment, not a forecast. They are worth watching only at levels you chose before the session, such as yesterday's value area edges or the overnight low, and even then as context rather than as a trigger.

Why Visible Liquidity Can Be Fake

A promise, not a contract

A limit order is, until the moment it executes, a **cancellable offer**. No rule obliges someone who has posted a 200-contract bid to leave it there, provided they did not place it with the intention of cancelling it (that would be spoofing, covered later). So treat displayed liquidity as a *promise*, not a *commitment*. There are at least five reasons the displayed book differs from what you could actually trade against.

- 1. Normal, legitimate cancellations.** Market makers adjust orders constantly according to risk. When price rushes toward them, the rational move is to pull. The uncomfortable result: liquidity tends to disappear exactly when it is needed most.
- 2. Flickering.** Some orders appear and disappear within fractions of a second. On your screen they look big; in practice you never had a real chance to trade with them.
- 3. Hidden liquidity.** The error runs the other way too. Iceberg orders display less than their true size, so what you see can be *smaller* than what is really there, not only larger.
- 4. Manipulative orders.** Spoofing and layering are orders placed with the intent to cancel before execution. They are illegal in US markets, but from the outside proving intent is practically impossible.
- 5. Feed delay.** What reaches your screen is already late, by milliseconds or more, compared with the state of the matching engine. In fast moments, the picture you are looking at is already old.

What the minutes before a release look like

A classic gold example: in the two minutes before a scheduled 08:30 New York time data release, such as the consumer price index (CPI) or the US jobs report, liquidity on a GC heatmap often thins out on both sides and the spread widens. Market makers do not want to be caught holding resting orders when the number comes out. When the release hits, price can jump through the empty space, because there is little resting liquidity in the way.

What to trust

The practical conclusion is simple: take liquidity seriously only when **trading confirms it**. That means price reached the level, volume actually traded there, and the level held. That is absorption you can see in the footprint (Chapter 9). Liquidity that has never been touched is weak information.

Key idea: Displayed size can be larger than the real tradable size (pulls, flicker, spoofing) or smaller (icebergs). Only trades at a level tell you what was real.

Common mistake: "Every wall that disappeared was a spoof." Most cancellations are legitimate risk management. Calling something spoofing without regulatory-grade data is meaningless.

Common mistake: "Then the DOM is worthless." It is not worthless; it is probabilistic information, not certain information.

Time & Sales: Reading the Tape

What the tape is

Time & Sales (T&S), also called **the tape**, is the chronological list of every completed trade. Each line is a **print** and shows:

- the **time**,
- the **price**,
- the **size** (number of contracts),
- and, on CME futures, the **aggressor side**: whether a buyer lifted the ask with a market order (a **buy aggressor**, usually shown in green) or a seller hit the bid (a **sell aggressor**, usually red).

The name comes from the paper ticker tapes that telegraph machines printed in old stock exchanges.

A short slice of gold tape might look like this:

Time (NY)	Price	Size	Aggressor
10:14:07.312	4,240.1	3	Buy
10:14:07.312	4,240.1	1	Buy
10:14:07.890	4,240.0	2	Sell
10:14:08.105	4,240.0	7	Sell
10:14:08.106	4,239.9	1	Sell

Approximate XAUUSD equivalent: around 4,210–4,213.

Intention versus action

The DOM shows **intention**: orders waiting. The tape shows **action**: what actually traded. The delta and footprint charts from earlier chapters are in fact **statistical summaries of the tape**. The footprint adds up these same prints by price and by aggressor side; delta is the difference between aggressive buying and aggressive selling.

So why look at the raw tape at all? Because it preserves **sequence and rhythm**, which summaries throw away. Did the large prints arrive back to back? Did price rise on many small prints or a few big ones? Did trading keep hitting one price while price refused to move, a possible sign of absorption?

Limits of the tape

- **A print is not necessarily an order.** One large market order that matches against several resting orders may appear as several prints, depending on the data feed and platform. The section on big trades below explains this in detail.
- **In busy moments the tape runs faster than a human can read.** Filters and grouping become necessary.
- **Some trades are reported under different rules**, for example legs of calendar spreads or privately negotiated block trades. They should not be mixed with ordinary screen trades without care.

Common mistake: "A green print means someone bought and nobody sold." Every trade has one buyer and one seller. The colour only tells you who was the aggressor. A green print of 5 at the ask means a buyer used a market order and a resting seller was filled for 5 contracts.

Common mistake: "Old-school tape reading still gives a reliable edge." This is a popular claim, not proven. Competing by eye against algorithms has become very hard.

Tape Speed: Bursts, Quiet, and What Pace Tells You

Measuring pace

Tape speed, or **pace**, is how many trades or how many contracts execute per unit of time, for example per second or per ten seconds. Old floor traders felt it as the sound of the ticker; modern platforms show it as a number or a bar.

Speed only means something relative to the hour

Absolute speed means nothing on its own. Gold at 03:00 New York time (London), at 10:00 (New York) and at 08:30 on a release day runs at completely different rhythms. Fifty trades in ten seconds could be a frenzy at 03:00 and an ordinary moment at 10:00. Always measure speed **relative to what is normal for that same time of day** on previous days. A simple way is to compare the current pace with the typical pace for the same clock time over the last few weeks, as a z-score (how many standard deviations above or below normal) or as a percentile.

A burst has three possible meanings

A **burst** is a sudden spike in the number of trades in a short time. It can mean:

1. **The start of a move.** New information or a break of a level activates aggressors.
2. **The end of a move (a climax).** The last buyers or sellers, often stop orders, rush in, and the other side absorbs them.
3. **A mechanical event.** Stops being triggered behind a level, or a scheduled news release.

The burst itself does not say which. Two things decide: **location** (are we at a level that matters?) and **outcome** (did price continue after the burst, or stall?).

Quiet is information too

When price arrives at an important level and the tape is calm, nobody is in a hurry to push through. Many traders believe that slow breakouts are less trustworthy than fast ones. That is a popular claim, not proven, and it is exactly the kind of idea you should test on your own data before relying on it.

Key idea: Pace measures *activity*, not direction. Combine it with delta or the footprint, and always judge it against the same hour of the day.

Common mistake: "Fast tape means get in." A burst at a climax can be the worst possible moment to act.

Common mistake: "High speed means institutions have arrived." This is a popular claim, not proven. It may just be a cluster of small traders' stops being triggered.

Gold note: On a typical gold day you can expect at least three bursts that have nothing to do with each other: the London open, any 08:30 New York data release, and often one at the high or low of the day. Same spike on the speed panel, very different meaning.

Big Trades: Why One Order Prints as Many Small Fills

How matching splits an order

Imagine someone sends a market order to buy 50 GC contracts. At the best ask there are 12 small limit orders from different traders, queued in FIFO order. The matching engine fills the 50-lot against each of them in turn, and if they are not enough, it moves on to the next price. The result, depending on the feed and the platform, may be **dozens of small prints** on the tape (3, 1, 5, 2, 4...) all with almost the same timestamp, instead of one print of 50.

CME's market data feed reports the trades caused by **one** aggressive order as a single match event, summarised by price level and optionally with order-level detail. Many platforms and data vendors, however, display trades per resting order or per price, which breaks the event into pieces. That is why tools exist that **group** near-simultaneous prints in the same direction back together, to estimate the true size of the aggressive order.

A real example: 278 contracts in one go

On 2 October 2026, the day of the US jobs report, gold futures reacted at 08:30 New York time as the number was released. In the first moments, one aggressive buy order for **278 GC contracts** hit the book. It was too large to be filled at the best ask, so it consumed the whole ask queue at one price, then the next, then the next, sweeping several consecutive price levels in a fraction of a second. On a raw tape this single decision appeared as a long column of small prints at rising prices. Only by grouping the fills of the same aggressive event does the true picture emerge: one participant, one order, 278 contracts. That is 27,800 troy ounces (278×100), which at a gold price above \$4,200 is well over \$100 million of notional value, and every \$1 move in gold changes its value by \$27,800.




Figure 27. Oct 2, 8:30 New York time: one aggressive buy order of 278 contracts swept five ask levels. A footprint shows five prints; order-level data shows one order.

Note what this one order tells you and what it does not. It tells you that one participant was impatient enough to pay through several prices at the moment of the release. It does not tell you who they were, whether this was their whole position or one slice of it, or what price would do next.

Two problems that remain

1. **Grouping is an estimate.** Grouping prints by time and direction can merge two independent orders that happened to arrive almost together, and can split a single order if the timestamps are imperfect.
2. **Truly large traders often do not send one huge order.** Execution algorithms such as **TWAP** (time-weighted average price) and **VWAP** algorithms, or icebergs, slice a parent order into dozens of child orders spread over minutes or hours, specifically to avoid looking big. So a "big trade" on the tape is better described as **an impatient aggressor** than as "smart money".

What our own testing showed

In our own testing, simple "big trade" detection from raw prints was unreliable: most of the prints it flagged were not genuinely large orders at all, precisely because real orders print as many small fills and unrelated small fills can look like one. Grouping prints improved the picture. But even correctly identified big trades were not, on their own, a signal that predicted direction after costs.

Key idea: A raw print is not an order size. One large market order matches against many smaller resting orders and can print as many small fills.

Common mistake: "Big bubble = institution = follow it." This is a popular claim, not proven.

Common mistake: Filtering the tape for prints above a fixed size without grouping. You get both false alarms and missed events. And comparing big-trade statistics between two platforms that group prints differently compares apples with oranges.

Liquidity Sweeps and Stop Runs

Sweeps

A **sweep** happens when one or more aggressive orders consume **several price levels** in rapid succession: the entire ask at the first price is taken, then the second, then the third. On a heatmap it looks like a steep staircase with trade bubbles one after another. The 278-lot order from the previous section is a textbook sweep.

A useful objective definition: a sweep is the consumption of at least a few consecutive price levels within a very short time (fractions of a second to a few seconds), by aggressors in one direction.

Stop runs

A **stop run** is a particular kind of sweep. Above an obvious **swing high** (a recent peak that everyone can see), and below an obvious swing low, the previous day's high or low, or the overnight high or low, **stop orders** tend to accumulate. Above a high, these are the protective buy stops of traders who are short, plus breakout buy orders from traders who want to join a move up. When price reaches that zone, the stops become market orders and create new aggressors in one burst. The result is a jump in tape speed and a short, sharp price spike.

CME's feed also marks trades that came from triggered stop orders, and some platforms use MBO data to highlight stop executions on the chart, with limitations.

The famous pattern: break and close back inside

The pattern most traders talk about goes like this: price breaks the high, stops are triggered, but there is no **follow-through**, and price closes back **inside the previous range**. This is often called a **failed breakout**. In the language of auction market theory (Chapter 7), the auction above the level was **not accepted**. The logic: if no new buyers arrived and only stops were executed, the fuel for the move ran out.

What to check, after the fact:

1. Did the sweep happen at a level you had marked in advance?
2. What did delta do above the high? Heavy buying with no further progress in price suggests absorption.
3. In the next few minutes or bars, was price **accepted** back inside the range, or did it build value above the high?

The limitation that matters

Many sweeps **continue**. A genuine breakout starts with exactly the same sweep. You can only see "it closed back inside" afterwards, never at the moment of the sweep itself.

Common mistake: "Market makers deliberately hunt my stops." This is a popular claim, not proven. Price goes where liquidity is, and stops *are* liquidity. No conspiracy is needed to explain why price visits them.

Common mistake: "After every sweep, price reverses." This is a popular claim, not proven.

Common mistake: Acting in the middle of a sweep, before any evidence of acceptance or rejection exists.

Iceberg Orders: Native vs Synthetic

What an iceberg is

An **iceberg order** is a large limit order of which only a small part, the **display quantity**, is visible in the book. When the visible part is filled, the next part is released from a hidden reserve. This is called a **refill**. The name is the obvious image: only the tip of an iceberg is above the water.

Why anyone uses one

Showing a 500-contract order tells the whole market about it. Others may queue in front of it, step away from it, or move price away from it. An iceberg lets a large order be worked with less market impact.

Two kinds on CME

1. Native icebergs. The exchange itself holds the hidden reserve. On CME the order is sent with a display quantity, and when the visible tip is filled, Globex releases the next tip automatically. Two technical clues follow from this. In MBO data, the refreshed tip keeps the **same OrderID** as before. And the exchange's trade reporting can show a traded quantity at that order that is **larger than the quantity that was ever displayed**. Together these make native icebergs reasonably detectable with MBO data. Whether native icebergs are permitted depends on the product; check CME's documentation for the contract you trade.

2. Synthetic icebergs. The trader's platform or algorithm builds the iceberg itself. Each time the small visible order is filled, the software sends a **brand-new order with a new OrderID**. To the market these are just ordinary limit orders, indistinguishable from anyone else's. One more consequence: every synthetic refill is a new order, so under FIFO it joins the **back** of the queue behind everything already waiting at that price.

	Native iceberg	Synthetic iceberg
Who holds the hidden size	The exchange	The trader's platform or algorithm
OrderID on refill	Same	New each time
Queue position on refill	Managed by the exchange	Back of the queue each time
Detectable with MBO	Reasonably, from the clues above	Only as a statistical pattern, often not at all

A real example: shows 2, trades 176

Back to 2 October 2026. A few minutes after the 08:30 jobs-report burst, with gold trading around 4,241.5 on the GC contract (roughly 4,212–4,215 in XAUUSD terms, using an approximate \$27–30 gap), a hidden buyer appeared on the bid at 4,241.5. Its visible part was only **2 contracts**. Yet **176 contracts** traded into that order as aggressive sellers kept hitting it, and the visible 2 kept reappearing. Price held at that level for about 25 minutes.



Figure 28. An iceberg shows a small visible size that refills after each fill. On Oct 2 a hidden buyer absorbed 176 lots at 4241.5 while showing 2.

Then the level broke, and gold fell about \$40 (on GC that is \$4,000 per contract; on MGC, \$400 per contract).

This example is valuable precisely because it does *not* end the way the folklore says it should. A buying iceberg was detected, it really absorbed a lot of selling, it really held price for a while, and price then fell sharply anyway. An iceberg tells you that someone had hidden size at a price. It does not tell you their overall direction (the buying could have been closing a short position, or one leg of a hedge or spread), how long they intend to stay, or where price will go next. The order can also be cancelled at any moment.

Key idea: An iceberg reveals hidden size at a price, not a prediction of direction.

Common mistake: "A buy iceberg means price will go up." This is a popular claim, not proven, and the 2 October example shows the opposite can happen.

Common mistake: "Icebergs always mean banks or institutions." This is a popular claim, not proven. Anyone with a suitable platform can place one.

Detecting Icebergs: Why It's Harder Than It Looks

How detection generally works

Commercial tools and academic research on CME data (for example a well-known 2019 study by Dmitry Zotikov on iceberg detection in CME futures) use roughly the same logic.

For native icebergs:

1. Compare the quantity traded against an order with the quantity it displayed. If more traded than was ever shown, a hidden reserve existed.
2. Watch for size reappearing under the **same OrderID** immediately after a fill.

For synthetic icebergs: look for new limit orders at the **same price**, appearing **immediately after** the previous one is filled, with similar sizes. This is only a **statistical pattern**, never a proof.

Why it is hard

1. The synthetic ambiguity. In a busy gold market, dozens of traders may be placing orders at the same price at the same time. One trader's "refill" is indistinguishable from another trader's fresh order. You get many false positives (detections of icebergs that were not there) and many false negatives (real icebergs that went unnoticed). At least one major commercial vendor states in its documentation that its detection targets native icebergs for this reason.

2. Data and timing. Vendor documentation also notes that timestamp differences between systems and data aggregation under heavy load can shift detections or create false ones. A feed interruption means an incomplete history.

3. Detection is not meaning. Even a correct detection says only "hidden size was here". It does not say whether the owner is a net buyer or seller overall. A sell iceberg might be hedging a larger long position elsewhere. It does not say where price goes next.

4. Detection only comes afterwards. An iceberg is detected after it has refilled several times, which means part of the "news" is already in the price by the time you see the label.

A useful way to hold all this in mind is a simple two-by-two table:

	Iceberg really present	Iceberg not present
Detector says yes	True positive	False positive
Detector says no	False negative	True negative

And beside it, one more line: *even a true positive is not a direction.*

What our own testing showed

In our own testing, the side of a detected iceberg did not, on its own, predict where the market went next. Icebergs can be *seen*, and they show where someone wanted size, but that is information about intent, not a directional signal. The 2 October example is one vivid case of a general result.

Common mistake: "Tool X shows all icebergs." No tool reliably detects synthetic icebergs.

Common mistake: "A sell iceberg at the high means go short." This is a popular claim, not proven.

Common mistake: Judging a detection tool by a handful of hand-picked screenshots. That is **selection bias**: you only remember the examples that worked. Judge tools on all events, including the boring and the wrong ones.

Spoofing and Layering: What They Are and Why They're Illegal

Definitions

Spoofing means placing a bid or offer **with the intent to cancel it before it executes**. The usual purpose is to create a false impression of supply or demand so that other participants react, while the spoofer trades on the opposite side at a better price.

The classic sequence has three steps:

1. A large (fake) sell order appears above the market, creating the impression of selling pressure.
2. Meanwhile, the spoofer's real buy order lower down gets filled by sellers who reacted to that impression.
3. The large sell order is cancelled.

Layering is the multi-level version: many orders placed at several prices on one side of the book to build false depth.

The law

In the United States, the Dodd-Frank Act of 2010 added section 4c(a)(5)(C) to the Commodity Exchange Act. It explicitly prohibits conduct that "is, is of the character of, or is commonly known to the trade as, 'spoofing' (bidding or offering with the intent to cancel the bid or offer before execution)". Exchange rules forbid it too, notably CME Rule 575 on disruptive practices. The key element is **intent**; carelessness alone is not enough. And the law does not depend on size: a small trader placing orders to scare others and then cancelling them is breaking the same rule.

Real cases exist, including in gold. In 2020 the US Commodity Futures Trading Commission (CFTC) resolved a case against JPMorgan that included spoofing in precious metals futures, gold among them, on COMEX between 2008 and 2016. Total payments across the related US cases came to roughly \$920 million. Another well-known case is Navinder Sarao, who used a layering algorithm in E-mini S&P 500 futures and whose activity was linked by authorities to the events of the May 2010 "flash crash".

What it might look like on a heatmap

Possibly: large bands that keep retreating as price approaches, or vanish just before being touched, while trading happens on the other side of the book. **But** legitimate market-maker cancellations produce exactly the same shape. Only a regulator, with identity data, full order records and communications, can establish intent. A screenshot cannot.

Key idea: We teach spoofing so that you are not fooled by it, never as something to do or to "ride". Any course or video that presents spoofing as a strategy is inviting you to break the law.

Common mistake: "Every wall that got pulled was a spoof." No. Most cancellations are legitimate.

Common mistake: "I'm small, so placing orders to scare others and cancelling them is fine." The law is about intent, not size.

Common mistake: "I can detect spoofs and trade against them." This is a popular claim, not proven.

Order Flow Imbalance (OFI): What the Research Actually Says

The question

In a widely cited paper, Rama Cont, Arseniy Kukanov and Sasha Stoikov ("The Price Impact of Order Book Events", *Journal of Financial Econometrics*, 2014) asked a simple question: over short

intervals, what explains price changes? Their answer was **order flow imbalance (OFI)**: the imbalance between supply pressure and demand pressure *at the best bid and best ask*.

How OFI is built

Every event at the inside market adds either buying pressure or selling pressure:

Event at the inside	Contribution to OFI
Size on the best bid grows, or the best bid moves up	+ (more demand)
Size on the best bid shrinks (trade or cancel), or the best bid moves down	– (less demand)
Size on the best ask grows, or the best ask moves down	– (more supply)
Size on the best ask shrinks (trade or cancel), or the best ask moves up	+ (less supply)

OFI over an interval is the sum of these contributions. The main finding is an approximately linear relationship:

$$\text{price change} \approx \beta \times \text{OFI}$$

where β (beta) is a sensitivity coefficient. And **β is inversely related to market depth**: in a thin book, the same imbalance moves price more; in a deep book, it moves price less. This is the same intuition as the DOM section, now stated as an equation: the less resting size there is to get through, the further a given push travels.

A small worked example. In a ten-second window on gold, 25 contracts are added to the best bid (+25), 10 contracts on the best ask are lifted by buyers (+10), and 15 contracts are added to the best ask (–15). $\text{OFI} = +25 + 10 - 15 = +20$. If, in that market at that time, β were about one tick per 20 contracts of imbalance, the model would associate this window with a price change of about +1 tick (\$10 on GC, \$1 on MGC). The numbers are illustrative; β has to be estimated from data and changes with depth and time of day.

How OFI differs from delta

Delta counts only aggressive trades. **OFI** also counts the **adding and cancelling** of limit orders at the inside. So a mass cancellation of bids (pulling) produces negative OFI even if not a single contract trades. OFI is, in that sense, closer to a measure of the whole contest at the inside, not just the part that ended in trades.

The honesty point that matters most

The paper's core finding is about **contemporaneous explanation**: in the *same* interval in which OFI was positive, price went up. That is very different from "OFI is positive now, so price will go up later". The power to predict *future* price changes is much weaker, concentrated at very short horizons (seconds), and fought over by the fastest participants, high-frequency trading firms with

infrastructure no retail trader can match. The study was also carried out on US stocks; any application to gold futures has to be tested separately.

Key idea: Explains is not predicts. OFI describes the mechanism of price change in the moment; it is not, on its own, a forecasting tool for a human trader.

Common mistake: "Academic research has proved that OFI predicts price." The main relationship is contemporaneous.

Common mistake: "OFI is just delta." It is not: it includes additions and cancellations of resting orders.

Common mistake: Applying results from US equities to gold futures without testing them on gold data.

For our purposes, OFI is valuable as a way to *understand* why price moves when it does, and as a piece of context at levels chosen in advance. It is not an independent signal.

Chapter summary

- The **order book** (resting limit orders) is intention; a **trade** (aggressor meets passive order) is fact.
- **L1** = best bid, best ask, last trade; **L2** = total size at several prices; **L3/MBO** = every order, anonymously, with its lifecycle. None shows untriggered **stops** or hidden **iceberg** size.
- The **DOM** is L2 as a ladder with no memory. Price moves one tick only when the whole inside quantity is consumed or pulled. **MBO** separates **filled** from **pulled**.
- Gold matches **FIFO**: reducing size keeps priority, increasing size or changing price loses it. **Touched is not filled**; easy fills can be bad fills.
- The **heatmap** is the book with memory: bands hold, get eaten, or get pulled. **Walls** are relative; **stacking** and **pulling** are momentary intention, mostly routine risk management.
- Displayed liquidity can be **larger** than real (pulls, flicker, spoofing, feed delay) or **smaller** (icebergs). Trust liquidity that **trades and holds**.
- The **tape** records actual trades with aggressor side and keeps **sequence and rhythm**. **Tape speed** is judged against the same hour; a burst's meaning depends on location and outcome.
- **One large order prints as many small fills**. The 278-contract buy on 2 October 2026 swept several prices in a single event. Big trades are impatient aggressors, not necessarily smart money.
- **Sweeps** consume several levels fast; **stop runs** are sweeps into clusters of stops. Failed and genuine breakouts start the same way.

- **Icebergs** hide size. **Native** ones are reasonably detectable with MBO; **synthetic** ones look like ordinary orders. Detection is hard, comes late, and is not direction: on 2 October a hidden buyer absorbed 176 contracts while showing 2, and price later fell about \$40.
- **Spoofing and layering** are illegal in the US (CEA 4c(a)(5)(C), CME Rule 575). Intent is the key element, and a heatmap cannot prove it.
- **OFI** explains price changes in the same interval, with impact inversely related to depth. Explaining is not predicting.

Checklist

- I can say which data level (L1, L2 or L3) a claim actually requires before I believe it.
- I can find the inside market and the spread in dollars for GC and MGC, and when resting size drops I ask "filled or pulled?".
- I remember that touched is not filled, and I do not assume fills at the first touch in any test.
- My heatmap colour settings are fixed and written down.
- I judge walls relative to the session and hour, not by a fixed number.
- I only take liquidity seriously when trades confirm it at a level I chose in advance.
- I judge tape speed against the same time of day on previous days.
- I know my platform's rule for grouping or splitting prints before I use big-trade statistics.
- I treat a sweep as unresolved until price shows acceptance or rejection.
- I treat icebergs as information about hidden size, not as direction.
- I never place orders I intend to cancel in order to influence others.
- I remember that OFI explains, it does not predict, and that research on stocks must be retested on gold.

Quiz

1. Level 2 shows 60 contracts on the bid at 4,239.5. What do you know for certain?
2. On CME's FIFO matching, which change to a resting order keeps its place in the queue: increasing the size, moving the price by one tick, or decreasing the size?
3. On a heatmap, a bright band fades as price approaches, with no trade bubbles at that price. What most likely happened?
4. Why can a single 50-lot market buy order appear on the tape as many small prints?
5. Under US law, what makes an order "spoofing"?

Quiz answers

1. Only that 60 contracts are resting at that price right now, coming from orders you cannot see individually. You do not know whether it is one order or many, who placed them, or whether

they will hold price. They may be cancelled at any moment.

2. Decreasing the size. Increasing the size or changing the price sends the order to the back of the queue.
3. The liquidity was most likely **pulled** (cancelled). Fading *with* trade bubbles would indicate it was filled.
4. Because it matches against many smaller resting orders, one after another in FIFO order and possibly across several prices, and many feeds and platforms display each of those fills as a separate print.
5. It was placed with the **intent to cancel it before execution**. Size alone, or the mere fact that an order was cancelled, does not make it spoofing; intent is the key element.

PART VI • PRACTICE

11

Putting It Together on Gold

In this chapter: You will turn the separate tools of the previous chapters (VWAP, volume profile, auction theory, delta, footprint, the order book and the tape) into one repeatable decision process: **Location first, then Confirmation from flow, then Risk**. You will learn how to build a short map of gold levels before the session, how to write a confirmation-or-veto checklist, how to write a one-page pre-market plan, how to size a position per tick on GC and MGC and measure results in R, why daily loss limits matter and why trailing drawdowns quietly shrink your room, how to keep a journal that can answer real questions, how a full gold day looks when you follow the process, and how gold behaves around CPI, jobs reports and Fed decisions. Everything here is education about process. Nothing in this chapter is a trade signal.

Location First: The Framework

By now you have a large toolbox: delta and cumulative delta, profiles, VWAP, footprint cells, the DOM, the heatmap, the tape, sweeps and icebergs. The danger of a large toolbox is that you start using every tool all the time.

Gold futures produce dozens of "events" every minute: tape bursts, stacked imbalances, large prints, walls of offers that appear and vanish. If you react to all of them, you are reacting to noise, and every reaction costs commission and slippage. A trader who takes forty small decisions a day because "the flow looked strong" has multiplied their costs by forty and cannot tell which decisions were good.

The framework in this book has three steps, always in the same order.

Step 1 — Location: where?

Before the session opens, you choose a small number of price levels that have an auction logic behind them. These are places where many participants are watching and where the previous auction showed balance or imbalance:

- Yesterday's value area high and low (VAH and VAL) and yesterday's point of control (POC).
- A naked POC: an earlier day's POC that price has not traded back to yet.
- VWAP for the session, and sometimes a weekly VWAP or a VWAP anchored to an event.
- Session highs and lows: Asia, London, the prior day and the overnight range.

At these levels the market is more likely to make a **decision**: accept the price (trade there with time and volume) or reject it (move away quickly). Between them, most of what happens is rotation and noise.

Step 2 — Confirmation: is the flow agreeing?

Only when price reaches one of your pre-chosen levels do you look closely at the order flow. Are the aggressive buyers or sellers making progress, or are they being absorbed? Is liquidity stacking

in front of price or being pulled away? The flow has three possible jobs at this point:

1. **Confirm:** the evidence is consistent with your hypothesis about the level.
2. **Time:** it helps you decide *when*, not just *where*.
3. **Veto:** the evidence contradicts your hypothesis, so you do not act.

Of these three, the veto is probably the most valuable. A veto removes a decision, and a removed decision costs nothing.

Step 3 — Risk: where am I wrong, and how big?

Before any position exists, you know where your idea is proven wrong (usually beyond the level), how large the position must be so that being wrong costs a fixed amount, and where the next reference level is.



Figure 29. The pre-London map for Oct 5, 2026: yesterday's levels in futures prices with their XAUUSD equivalents (gap -27.8).

Why this order?

Why not let the flow find opportunities for you? Because, in published research and in our own testing, order-flow signals on their own did not reliably predict direction for a non-high-frequency trader once realistic costs were included. Many ideas that looked like "something" before costs became nothing after them (Chapter 12).

What remains useful is order flow as a **filter in context**. At levels chosen in advance, it reduces the number of decisions you take and improves the quality of how you evaluate them. That is a

process, not a guarantee of profit.

Key idea: Order flow is not where you look. It is what you check once you are already at the right place.

Common mistake: Adding a new level during the day because "price reacted there". A level drawn after the reaction is hindsight, not location. If it was not on your map before the open, it is not a level for today.

Common mistake: Fifteen lines on the chart. When everything is a level, nothing is a level. A map with too many lines guarantees that price is always "at a level", and the location step stops filtering anything.

"Order flow alone tells you the direction" is a popular claim, not proven. The honest version: order flow tells you more about *who* is active, *where* and *how hard*, and that information is most useful at a place you chose before the session.

Try it: Take a gold five-minute chart from a recent day. Before looking at that day's action, mark only five levels from the prior day (VAH, VAL, POC, the London high and the session VWAP). Shade everything between them grey and label it "no-trade zone". Then scroll through the day. Notice how much of the activity happened inside the grey areas.

Building Your Gold Level Map

A level map is a short list of prices where you will "listen" today. It is not a prediction of where price will turn; it is a list of places where a market decision is more likely, written down in advance so you cannot rewrite it later.

The standard sources

All of these were defined in earlier chapters; here is how each one earns a place on the map.

Yesterday's value area and POC. The value area is the price range where roughly 70% of yesterday's volume traded, with VAH and VAL as its edges, and the POC is the single price with the most volume (Chapter 5). In auction market theory (Chapter 7), opening outside yesterday's value and then returning inside, or being rejected from one of the edges, are among the most important decisions of the day.

Naked POC. A POC from an earlier session that price has not traded at since (Chapter 6). It marks a price where a previous auction found strong agreement and that the market has not revisited.

VWAP. The volume-weighted average price of today's session (Chapter 4): the average price at which today's participants have actually traded. A weekly VWAP or an anchored VWAP from an

event (for example, the moment of a big data release) can be added when relevant.

Session highs and lows. Asia, London, the prior day and the overnight session. In gold, the London high and low deserve special attention, because a large part of price discovery happens during London hours, before New York is fully active.

Large round numbers. Every \$50 or \$100 in gold, for example. The usual explanation is human behaviour and clustered orders. This is a popular claim, not proven; treat round numbers as secondary levels at most.

Ranking: confluence first

Not all levels are equal. Give the highest priority to levels where two or more independent sources line up. If yesterday's VAL sits within a few ticks of today's London low, that zone is more interesting than either alone, because two different groups of participants have a reason to care about it.

Merge levels that are within two or three ticks of each other. Four lines in a \$0.30 band are one zone, not four levels.

Write it down, with "what I want to see"

For each level, record four things: the price, the source, the confluence, and what you would want to see from the flow if price arrives. This last field is what turns a line on a chart into a testable hypothesis. It also blocks hindsight bias: after the session you cannot say "I had that level too" if it was not on the list.

Here is an example map in table form. The numbers are hypothetical, chosen to illustrate the format. XAUUSD equivalents use an illustrative gap of about \$28 (GC minus ~\$28), and the real gap must be measured on the day, as explained in the appendix.

Rank	Level	GC price	XAUUSD $\approx$	Source	Confluence	What I want to see
1	Yesterday's VAL / London low	4,171.5	$\approx$ 4,143.5	Profile + session	Two sources	Aggressive selling that fails to make progress below
2	Yesterday's VAH	4,192.0	$\approx$ 4,164.0	Profile	One source	Buyers absorbed, no acceptance above
3	Naked POC from two days ago	4,203.4	$\approx$ 4,175.4	Profile	One source	Reaction or clean acceptance (both informative)
4	Yesterday's POC	4,182.6	$\approx$ 4,154.6	Profile	Near VWAP early in the day	Rotation; likely a "middle" of the day
5	Session VWAP	Moves through the day	GC – gap	Volume	Varies	Position of price relative to VWAP, and its slope

Example — not a signal. Prices are illustrative; the XAUUSD gap is approximate and varies by broker and by day.

Five to seven levels is enough. If you find yourself with twelve, rank them and cut the bottom half.

What a level is, and what it is not

No level "works" in a certain sense. A level is a place where the probability of a decision is higher. A claim like "level X holds every time" needs to be tested on your own data with the methods of Chapter 12.

On strong trend days, price can move through several levels without pausing. That is not the map failing. That is information: the market is **accepting** new prices, and the other-timeframe participants are in control.

Two pieces of folklore deserve labels. "A naked POC always gets filled, like a magnet" is a popular claim, not proven. The so-called "80% rule" (if price spends two 30-minute periods inside yesterday's value, there is an 80% chance it travels to the other side of value) is also a popular claim, not proven. Both may be useful as ideas to test; neither is a law.

Gold note: "Yesterday" must match your session setting (Chapters 4 and 6). With full Globex sessions (18:00–17:00 ET), yesterday's value comes from the whole session; with the day session only (about 08:20–13:30 ET), the overnight hours form their own profile. Do not mix the two.

Try it: Build tomorrow's map tonight. Limit yourself to six rows. For each row, fill in all four columns, including "what I want to see". Tomorrow, do not edit it after the open; only add notes in a separate column.

Confirmation From Flow: Yes, No, or Veto

When price reaches a level on your map, the question is not "is the flow bullish or bearish?" The question is narrower: **is the flow evidence consistent with my hypothesis about this level, or against it?**

Let us work through one hypothesis: "Price is rejected at yesterday's VAH and rotates back down into value." (Example only.) What would count as evidence for, and against?

Evidence consistent with the hypothesis (confirm)

- **Footprint and delta: absorption.** Aggressive buying at the level is heavy, but price does not make progress. Strong positive delta with no movement is the classic sign of effort without result (Chapter 9): someone passive is taking the other side of all that buying.
- **Footprint and delta: exhaustion.** Positive delta that fades quickly as price pushes into the level; buyers run out of urgency at the extreme.
- **DOM and heatmap.** Sell-side liquidity above the level stays in place or builds, and it is **confirmed by trades:** orders actually get filled there rather than only being displayed.
- **Tape.** A burst of buying at the level, followed immediately by quiet, and price drifting back.
- **Sweep.** The high of the level is swept (several prices taken in quick succession) and price closes back inside.

Evidence against the hypothesis (veto)

- Aggressive buying **with** progress; price trades above the level with time and volume. That is **acceptance** above the level.
- Offers above the level are pulled before price reaches them.
- Cumulative delta moves in the direction of the breakout with no divergence.

A traffic-light way to think about it

Light	Meaning	Typical evidence
Green — Confirm	Flow agrees with the hypothesis	Absorption, exhaustion, liquidity confirmed by fills, sweep and return
Amber — Wait	Mixed or not enough evidence yet	Price at level, flow inconclusive, no acceptance either way
Red — Veto	Flow contradicts the hypothesis	Acceptance beyond the level, liquidity pulled, delta and price agreeing on the break

Amber is a legitimate state. Many level tests end in amber, and the correct response is to keep waiting or to let the opportunity pass.

Write the checklist before the day

The most important rule of this lesson: the checklist is written **before** the session, in **objective** terms, not improvised in the moment. "Flow looks strong" is not a checklist item. "At least one one-minute bar at the level with strongly positive delta that does not close above the level" is a checklist item, because two people looking at the same bar would agree on whether it happened.

Keep it short. Three to five items is plenty. More conditions mean fewer opportunities, which is not necessarily bad, but there is a trap here: **mental overfitting**. After every loss, it is tempting to add a new condition that would have filtered out that specific trade. Ten losses later, your checklist describes your last ten losses perfectly and the next ten market situations not at all. Chapter 12 explains the statistical version of this trap.

If you use an order-book indicator for the DOM and heatmap side of this checklist (ours is called AK FlowBook, but any reliable heatmap and DOM tool will do), the same principle applies: decide in advance which of its readings count as confirm, wait or veto.

Honest limits

No combination of evidence is certain. Absorption can fail later: a resting iceberg that was absorbing buyers can be cancelled in a millisecond, and then price moves through the level. All the checklist does is improve your odds and define your risk. Two popular claims need their labels: "absorption always leads to a reversal" is a popular claim, not proven; "three stacked imbalances make a certain level" is a popular claim, not proven.

Your checklist is a **hypothesis**. It needs to be recorded, counted and evaluated like any other hypothesis.

Key idea: The most useful thing order flow may tell you on a given day is: do not take this trade.

Common mistake: Editing the checklist during the session to justify an entry you already want. If the checklist changes while price is moving, it is not a checklist; it is a story.

Try it: Write a three-item confirmation list and a three-item veto list for one level type (for example, "yesterday's VAL"). Then go through ten historical tests of that level type and, without looking at what happened next, record only what the checklist said: confirm, wait or veto.

The Pre-Market Plan for Gold

A pre-market plan is a short document written **before** the session, so that important decisions are made calmly instead of under the pressure of moving prices. In a "cold" state you weigh evidence; in a "hot" state, with a position open and price moving, you react. The plan moves decisions from hot to cold.

It fits on one page and has six parts.

1. Context: observe, do not predict

What kind of day was yesterday: trending, balanced, rotational? Where is gold opening relative to yesterday's value: inside or outside? What did the overnight sessions (Asia and London) do? Did anything notable happen in related markets, such as the US dollar index or Treasury yields? Write observations, not forecasts. "London made a new high above yesterday's VAH and then returned inside value" is an observation. "Gold will rally today" is a forecast and does not belong in the plan.

2. News calendar

Which data are released today: CPI, the jobs report (NFP), an FOMC decision, a scheduled speech? Write the exact time in New York time (ET). Then write your rule for before and after the release (see the last section of this chapter).

3. Level map

Your five to seven levels, ranked, each with "what I want to see".

4. Scenarios (if-then), each with an invalidation

Two or three scenarios are enough. For example:

- *If price reaches yesterday's VAL and I see evidence of sellers being absorbed, then my hypothesis is a rotation back toward yesterday's POC.*
- *If price is accepted below VAL (time and volume trading below it), then that hypothesis is invalid; I do not act on it and I wait for the next level.*

Every scenario needs its invalidation written next to it. A scenario without an invalidation is a hope.

5. Today's risk rules

Risk per trade (in dollars and in R), the maximum number of trades, the **daily loss limit**, and a stop condition such as "after two losses in a row, I stop for the day". These are the subject of the next two sections.

6. Personal state

Sleep, focus, stress. This is not a soft extra. If your state is poor, the plan should say "reduced size" or "no trading today". Most professionals treat their own condition as an input, the same way they treat the news calendar.

A filled example

Example — not a signal. All numbers are hypothetical.

Section	Content
Context	Opening inside yesterday's value. London made the overnight high, then rotated back. Yesterday was a balanced, D-shaped profile. Dollar index flat overnight.
News	CPI at 08:30 ET. Rule: no new positions from 08:25 until 08:45 ET.
Levels	1) VAL $\approx$ London low 4,171.5 GC ($\approx$ 4,143.5 XAUUSD) · 2) VAH 4,192.0 ($\approx$ 4,164.0) · 3) Naked POC 4,203.4 ($\approx$ 4,175.4) · 4) POC 4,182.6 ($\approx$ 4,154.6) · 5) VWAP (live)
Scenarios	A) At level 1, absorption of selling → hypothesis: rotation toward POC; invalid if accepted below. B) At level 2, absorption of buying → hypothesis: rotation toward POC; invalid if accepted above.
Risk	Fixed 1R per trade, maximum 3 trades, daily limit 2R, stop after 2 consecutive losses.
State	Slept well; focused. Normal size.

XAUUSD values use an illustrative gap of about \$28; measure the real gap on the day.

Why it works, and its limit

As a process, the plan works because it moves decisions into a calm state and creates something to compare against in your journal. Did you follow it? Did the day look like any of your scenarios?

Its limit: the plan is a discipline tool. Its quality depends on the quality of the hypotheses in it. A beautifully formatted plan built on levels without logic only looks professional.

Common mistake: A plan built around a fixed bias ("today gold goes up"). Bias makes you see confirmation everywhere and vetoes nowhere.

Common mistake: Writing the plan and then ignoring it. Your journal should measure "plan adherence" so you can see how often this happens.

Key idea: Almost every decision made after the open is worse than the version of it you could have made before the open.

Position Sizing Per Tick: GC vs MGC and R-Multiples

Position sizing is where all the abstract talk about risk becomes arithmetic. On gold futures, the arithmetic is pleasantly simple because tick values are fixed.

Tick values, again

Both GC and MGC move in ticks of \$0.10 per troy ounce.

- **GC** is 100 troy ounces, so one tick is worth **\$10** per contract.
- **MGC** is 10 troy ounces, so one tick is worth **\$1** per contract.

A \$1.00 move in gold is 10 ticks: **\$100 on one GC** and **\$10 on one MGC**. On XAUUSD, profit and loss per \$1 depends on how many ounces your lot represents, which varies by broker; check your broker's specification.

SIZE FROM RISK, NOT FROM CONVICTION

	GC	MGC
Ounces per contract	100	10
Tick size	\$0.10	\$0.10
Value of one tick	\$10	\$1
Value of a \$1 move	\$100	\$10
Risk \$200, \$4 stop	$\frac{1}{2}$ (not possible)	5 contracts

$1R = \text{the amount you decided to risk. Size} = \text{risk} \div (\text{stop distance} \times \text{value per point}).$

Figure 30. Position size comes from the risk you accept (1R) and the stop distance, not from how sure you feel. MGC makes small accounts sizeable.

The sizing formula

Number of contracts = Allowed risk in dollars ÷ (Stop distance in ticks × Tick value + Estimated costs per contract)

"Estimated costs" means commission and fees for the round trip plus an allowance for slippage. Always round **down**.

Worked example

Hypothetical, for learning the arithmetic. Not a trade idea.

Suppose your rules allow you to risk \$150 on one trade, and the logical point where your idea is proven wrong is \$3.00 away from where you would enter. \$3.00 is 30 ticks.

On GC: 30 ticks × \$10 = \$300 per contract, before costs. Even one contract is twice your allowed risk. **GC does not fit this trade.**

On MGC: 30 ticks × \$1 = \$30 per contract. Add an estimate of a few dollars for commission, fees and one tick of slippage, say about \$3. That gives about \$33 per contract. \$150 ÷ \$33 ≈ 4.5. Round down: **4 MGC contracts**, risking about \$132 including costs.

Allowed risk	Stop distance	GC (risk per contract ≈ \$300 + costs)	MGC (risk per contract ≈ \$33)
\$100	30 ticks (\$3.00)	0 contracts	3 contracts
\$250	30 ticks (\$3.00)	0 contracts	7 contracts
\$500	30 ticks (\$3.00)	1 contract (about \$300 + costs)	15 contracts

Costs are illustrative; use your own broker's commission and fees.

Notice the last row: at \$500 of allowed risk, one GC (about \$300 plus costs) and 15 MGC (about \$495) are both possible, and the MGC size uses the allowed risk more precisely. This is exactly why the micro contract exists.

Stop first, size second

The single most important sentence in this section: **the market decides where your stop belongs, not your wallet**. First find the logical place where your hypothesis is invalid (on the other side of the level, beyond where acceptance would prove you wrong). Then compute the size that makes that distance affordable.

The reverse, shrinking the stop until "one GC fits", produces a stop with no logic. It sits at a price that means nothing to the market, so it gets hit by ordinary noise. MGC gives you the flexibility to keep a logical stop at a small size.

Measuring in R

R is the amount you lose if your stop is hit: your initial risk on that trade. Express every result as a multiple of R:

- A full loss at the stop = **-1R**.
- A gain of twice the initial risk = **+2R**.
- A loss with extra slippage might be **-1.2R**.

R makes trades of different sizes comparable and summable. Two +\$200 wins are very different achievements if the risks were \$100 and \$400: in R, they are +2R and +0.5R. Once results are in R, you can compute **expectancy**: the average R per trade over many trades. That number, not the win rate, tells you whether a process has any edge at all.

The R-multiple approach was popularised by Van Tharp. The rule "risk 1% of the account per trade" is a common conservative convention, not a scientific law. "Risking 1% guarantees survival" is a popular claim, not proven: survival also depends on your real expectancy and on whether losses come in clusters.

Slippage breaks the neat -1R

A stop-market order becomes a market order when it is triggered. In a fast market, especially around news, it can fill several ticks beyond your stop price. Your real loss can then be more than 1R. A 30-tick stop that fills 9 ticks worse on MGC is a 39-tick loss: about -1.3R. Always include a slippage allowance in the sizing calculation, and record actual slippage in your journal.

Key idea: One tick of gold is \$10 on GC and \$1 on MGC. That difference decides whether your logical stop fits your risk.

Try it: Risk \$200, stop distance 20 ticks, ignore costs. How many MGC? (20 ticks × \$1 = \$20 per contract; \$200 ÷ \$20 = 10 contracts.) Now add \$3 of costs per contract and recompute. (Answer: \$23 per contract; 8.7, rounded down to 8.)

Common mistake: Reporting results only in dollars. Dollar results from different sizes cannot be compared or added meaningfully. Record R.

Daily Loss Limits & Why Trailing Drawdowns Are Dangerous

The daily loss limit

A **daily loss limit** is an amount, expressed in R (for example 2R or 3R), that ends your trading for the day once you have lost it. You decide it in the pre-market plan, not after the second loss.

Losses tend to come in **clusters**: some days the market does not behave the way your hypotheses assume, and some days your own state is poor. Decisions right after a loss are, on average, worse; the urge to "win it back" brings larger size, looser criteria and trades outside the plan. This is **revenge trading**. A daily limit stops it and puts a ceiling on your worst day.

Drawdown and the three types of floor

Drawdown is the fall from the highest point of your account. Prop firms (companies that give traders a "funded" account after an evaluation) usually set a **maximum drawdown**: a floor below which the account is closed. The rules differ by firm; what follows is a general description of the mechanics, not an endorsement or review of any firm. Read each firm's rules at the source.

There are three common types:

Type	How the floor moves
Static	The floor is fixed, for example always \$X below the starting balance.
End-of-day (EOD) trailing	The floor rises with your highest end-of-day balance and never comes down.
Intraday trailing	The floor rises with your highest account value at any moment, including open (unrealised) profit , and never comes down.

Why trailing floors are dangerous: a worked example

Hypothetical numbers to illustrate the mechanics.

You have \$2,000 of drawdown room on an account with an intraday trailing floor. You open a trade on gold. It moves in your favour until the open profit is \$600. At that moment, with intraday trailing, your floor has just risen by \$600: the room you have is still \$2,000, but it is now measured from the higher peak.

Then gold turns. The trade goes back through your entry and you close it at a loss of \$400.

- Your net result for the trade: **-\$400**.
- Your remaining drawdown room: \$2,000 – \$600 (open profit given back) – \$400 (loss) = **\$1,000**.

You lost \$400 of money but **\$1,000 of room**. On an EOD trailing account, the same trade would have cost only \$400 of room, because the open profit never became an end-of-day peak. On a static account, also \$400.

So an account can be closed without any single large loss: ordinary swings of open profit, given back, consume the room step by step. In gold, which can move several dollars in seconds around news, the effect is stronger.

Practical consequences

- Set your daily loss limit **smaller** than your remaining drawdown room, with a margin to spare.
- Compute risk per trade as a fraction of the **distance to the floor**, not of the account balance. If the floor is \$1,000 away, risking \$250 per trade means four full losses end the account.
- Track the floor's position in your journal every day, so you always know your real room.

Three myths: "if I don't take a loss, my drawdown isn't used" (on intraday trailing, giving back open profit uses it); "one good day will make it all back" (usually followed by bigger size after losses); "a prop account means zero risk" (evaluation fees are real money).

Key idea: You can close the day down only \$400 and still lose \$1,000 of room on an intraday trailing account. Know which kind of floor you are trading against.

Common mistake: Sizing from the account balance instead of the distance to the floor. A \$50,000 account with the floor \$1,500 away is, for risk purposes, a \$1,500 account.

The Trading Journal That Actually Teaches You

A journal is not a diary of feelings. It is a **database** you will ask questions of later: "Do my trades at yesterday's VAL do better than at VWAP?" "Is my slippage worse just after 08:30 ET?" "Do I break my plan more often after a loss?" It can only answer questions it was built for.

The fields

Every trade, and every trade you **vetoed**, should have these fields:

Group	Fields
Context	Date and time (ET), session, contract (GC/MGC)
Location	Which level from today's map, and its source
Scenario	Which if-then scenario from the plan
Evidence	Which checklist items were present, as ticks (yes/no), not free text
Execution	Entry, stop distance, size, initial R in dollars
Outcome	Result in R; actual slippage in ticks
Excursions	MAE and MFE in R
Process	Plan adherence (yes/no + short note); mental state (1-5)
Evidence images	Screenshot before entry and after exit

MAE (maximum adverse excursion) is the furthest the trade went against you; **MFE** (maximum favourable excursion) is the furthest it went in your favour, both measured in R. Over many trades, they show whether your stops are too tight (many trades with MAE just beyond $-1R$ that later went your way) or whether you leave a lot on the table (large MFE, small results).

Why log the trades you did not take?

Because without them you cannot know whether your filter removes bad trades or good ones. If your veto rule removed twenty situations, and those twenty would mostly have worked, your filter is costing you. If they would mostly have failed, it is helping. This is the logic of a **control group** in research: you need to see what happens to the cases you did not treat.

Example rows

Educational, hypothetical. Not a record of real trades.

Time ET	Contract	Level	Scenario	Evidence ticks	Decision	Stop dist.	Size	Result	Slippage	MAE / MFE	Plan kept	State
09:52	MGC	Y-VAL ≈ London low	A	Absorption ✓ Liquidity confirmed ✓ Acceptance below X	Taken	30 ticks	4	$-1.1R$	3 ticks	$-1.1R$ / $+0.6R$	Yes	4
10:40	MGC	Y-VAH	B	Absorption X Acceptance above ✓	Vetoed	—	—	(would-be: logged after)	—	—	Yes	4
11:15	MGC	Naked POC	—	Price accepted through	Vetoed	—	—	(would-be: logged after)	—	—	Yes	3

Note the first row: a loss, recorded honestly, with the plan followed. In process terms, that is a good day.

Analysing the journal

After at least several dozen trades, look at expectancy (average R) broken down by level type, scenario and time of day. But be careful. With 40 trades split into 10 categories, each category has about 4 trades. Any difference between them is almost entirely chance. This is the multiple-testing problem of Chapter 12 on a small scale: the more ways you slice a small sample, the more "patterns" you find that are not there.

A related myth: "My journal says setup X wins 80% of the time" based on ten trades. That is a popular claim, not proven, and with ten trades it cannot be anything else.

Fill it in before you know the outcome

At least the entry part of the journal (level, scenario, evidence, stop, size) should be filled in **before** you know the result. Otherwise memory quietly rewrites your reasons to match the outcome. Winners get "strong evidence" and losers get "I wasn't sure about that one". That is hindsight bias, and it destroys the value of the database.

Key idea: A good journal can answer a question you have not asked yet.

Common mistake: Logging only the "interesting" trades, or changing the whole system after five bad trades. Both make the journal useless as evidence.

Try it: Build the table above in a spreadsheet. Add two calculated columns: cumulative R and a running count of trades per level type. Keep it for twenty sessions before drawing any conclusion.

A Complete Gold Trading Day, Walked Through

This section follows one full day with the framework, phase by phase, in New York time (ET). It is a walkthrough of process, not a record of profit. The most instructive day is usually an ordinary one, including the long stretches where the correct action was to do nothing.

Pre-market (about 07:00–08:15 ET)

Review the overnight sessions. Asia trades from the 18:00 ET reopen the previous evening; London becomes active from about 03:00 ET. Where is gold relative to yesterday's value? Where are the London high and low? Check the economic calendar. Write the one-page plan: context, news rule, level map, two or three scenarios with invalidations, risk rules, personal state.

The main gold session opens (about 08:20 ET) and 08:30 data

The traditional COMEX day session begins around 08:20 ET. Many US data releases arrive at 08:30 ET. On a news day, follow your pre-written rule: for example, no new positions from a few minutes before the release until a set time after it. The New York stock market opens at 09:30 ET, which usually adds another wave of activity.

The first test of a level

Suppose price reaches the London low, which on today's map overlaps with yesterday's VAL. Now, and only now, you read the flow in detail:

- Is aggressive selling hitting the level without progress (absorption)?
- Are bids stacking and being confirmed by trades?
- Or are bids being pulled, and is price being accepted below the level (veto)?

Decision and execution

If the checklist says **confirm**: the position is opened with the stop already defined (beyond the level or beyond the sweep), the size taken from the formula, and the next reference level (for example yesterday's POC or VWAP) noted as the logical objective. If the checklist says **veto**: nothing happens, and the veto is logged in the journal with the same care as a trade.

Management

Management follows rules written in advance, for example "the stop is moved only after +1R" or "the stop is never moved". Emotional management, such as widening the stop "to give it room" or closing at the first discomfort, is usually the largest single source of errors.

Midday and settlement (about 13:30 ET)

Activity usually slows through the middle of the New York day. Around the GC settlement window (13:29–13:30 ET), behaviour can differ, because settlement prices are computed from trading in that window. Many day traders finish before or around then.

Review (after the session)

Complete the journal, compare the plan with what happened, save the screenshots and write one short lesson. One sentence is enough.

The day as a timeline

Time (ET)	Phase	Mode
07:00–08:15	Overnight review, calendar, write plan	Think
08:20–08:45	Session open, data release; wait by rule	Wait
Variable	Level test 1: flow says veto, logged	Wait
Variable	Level test 2: flow says confirm; position per plan	Act
Variable	Management by pre-set rules; exit	Act
~13:30	Settlement; stop for the day	Wait
After close	Journal, plan vs reality, one lesson	Review

Two pieces of folklore to label here: "The first 15 minutes always tell you the direction of the day" is a popular claim, not proven. And showing only great days, the ones that make the process look brilliant, is cherry-picking. A process is judged over many days, never on one.

Common mistake: Trading in the middle of the day out of boredom. A quiet market between levels is not an opportunity; it is the grey "no-trade zone" on your map.

Key idea: The process is the same on every day. The outcome of a single day tells you very little about whether the process is good.

Try it: Pick a past gold session and replay it bar by bar from 07:00 ET. Write the plan using only information available at 08:15. Then step forward and record every decision point: level reached, checklist result, action. Do not skip the quiet hours.

Gold on News Days: CPI, NFP, FOMC

Gold is sensitive to major US data because it changes expectations for interest rates and the dollar. Three releases matter most.

Release	What it is	Usual time (ET)
CPI	Consumer Price Index: monthly inflation	08:30
NFP / Employment Situation	Non-farm payrolls: the monthly jobs report, usually the first Friday of the month	08:30
FOMC	The Federal Reserve's interest-rate decision	Statement 14:00; press conference usually 14:30

Around FOMC, the reaction can change direction more than once during the press conference.

What the order book and tape usually look like

What follows is a typical pattern as an observation, not a guaranteed rule.

Before the release. Market makers reduce their risk ahead of an unknown number. The book **thins out**: there is less resting size at each price. On a liquidity heatmap the area around price gets visibly "darker" (emptier). The **spread widens**. The tape slows down.

At the release. The tape **bursts**: a huge number of trades in a very short time. Price can "jump" through several prices that have almost no resting liquidity, which looks like a small gap within the day. Stop orders are triggered and multi-price sweeps occur. CME has protective mechanisms such as **Velocity Logic**, which can pause matching for a few seconds if price moves too far in too short a time.

After the release. Liquidity gradually returns. Sometimes the first move continues; sometimes it reverses completely. Both are common.



Figure 31. The Oct 2 jobs report: the iceberg buyer held price for about 25 minutes; when the level broke, gold fell about \$24 in the next 20 minutes.

Practical consequences

1. **Slippage.** A stop-market order triggered at the moment of the release may fill many ticks beyond its price. Your real loss can be larger than 1R. A stop placed tightly just before the release, on the assumption that it will fill exactly at its price, is built on a false assumption.
2. **Reading the flow in the first seconds is almost meaningless.** Everything is a burst. Absorption, imbalances and delta signatures in that window mostly describe the chaos of the release itself.
3. **Trailing drawdowns do the most damage here,** because open profit can appear and vanish within seconds.
4. **Many traders have a rule** such as "no new positions from X minutes before a major release until Y minutes after". Some prop firms have their own news rules. Whatever your rule is, it belongs in the pre-market plan, written before the day.

The number versus the expectation

Gold does not react to the number itself; it reacts to the difference between the number and what the market expected, and to the broader situation at the time. "Higher-than-expected CPI always means gold goes down" is a popular claim, not proven. "The first move after NFP is always fake" is also a popular claim, not proven. Neither should be part of a plan.

Gold note: Release times are in New York time. If you trade XAUUSD on a broker platform, the platform clock is often set to the broker's server time, not New York time. Convert the release times before the session, and remember that the US and Europe change daylight saving time on different dates, which shifts the gap between London and New York by an hour for a few weeks each year.

Common mistake: Placing a tight stop right before the release and assuming it will fill exactly at the stop price. When liquidity has been pulled, price skips through empty levels.

Try it: For the next CPI or jobs-report day, watch (without trading) the heatmap and tape from 08:25 to 08:40 ET. Write down three observations: when the book started thinning, what the spread did, and how long it took for resting liquidity to come back.

Chapter summary

- The framework is always **Location** → **Confirmation** → **Risk**: levels before the session, flow only at those levels, risk before any position.
- Order flow has three jobs at a level: **confirm**, **time**, **veto**. The veto (not trading) is a valid and often valuable outcome.
- A level map has five to seven ranked levels (prior value, naked POCs, VWAP, session highs/lows), each with "what I want to see"; confluence raises priority.
- The confirmation checklist is written before the day, in objective terms, and treated as a hypothesis to be tested.
- The one-page pre-market plan covers context, news, levels, if-then scenarios with invalidation, risk rules and personal state.
- Position size = allowed risk ÷ (stop ticks × tick value + costs). GC is \$10 per tick, MGC \$1 per tick. Stop first, size second. Results are recorded in R.
- Daily loss limits cap the worst day and block revenge trading. Intraday trailing drawdowns count open profit, so giving back open profit consumes room.
- A journal is a database: structured fields, results in R, actual slippage, MAE/MFE, plan adherence, and vetoed trades as a control group. Small samples prove little.
- On CPI, NFP and FOMC days the book thins, the spread widens, the tape bursts and slippage grows. Gold reacts to the surprise versus expectation, not to the raw number.

Checklist

- My level map has no more than seven levels, each with a source and "what I want to see".
- Levels with confluence are ranked highest; levels within two or three ticks are merged.

- XAUUSD equivalents are written next to GC levels using today's measured gap.
- My confirmation and veto checklist is written before the session in objective terms.
- My pre-market plan contains context (observations only), news times in ET, scenarios with invalidations, risk rules and my personal state.
- I find the logical stop first, then compute the size, and I round down.
- I know my daily loss limit in R and my remaining distance to any drawdown floor.
- I know whether my account's floor is static, end-of-day trailing or intraday trailing.
- I log every trade and every veto, with the entry part filled in before the outcome is known.
- I have a written rule for major news releases.

Quiz

1. In the framework of this chapter, when should you start reading the footprint and DOM closely?
2. Which level deserves higher priority: a round number on its own, or yesterday's VAL that lines up with today's London low? Why?
3. Your allowed risk is \$200 and your logical stop is 20 ticks away. Ignoring costs, how many MGC contracts can you trade, and why is GC unsuitable here?
4. On an intraday trailing drawdown account, what raises the floor?
5. Why does a stop-market order often fill worse at a CPI release?

Quiz answers

1. Only when price reaches a level you chose before the session. Between levels, the flow is mostly noise for this framework.
2. Yesterday's VAL that lines up with the London low. Two independent sources (the prior auction's value and the overnight session's extreme) give more participants a reason to care about the zone. A round number alone is a secondary level at best.
3. Ten MGC contracts: $20 \text{ ticks} \times \$1 = \$20$ per contract, and $\$200 \div \$20 = 10$. One GC would risk $20 \times \$10 = \$200$ before costs, so with any costs at all it exceeds the allowed risk, and it leaves no flexibility.
4. Any new peak in account value, including open (unrealised) profit. That is why giving back open profit consumes drawdown room.
5. Liquidity is pulled before the release, so the book is thin and the spread is wide. When the stop triggers and becomes a market order, price skips through nearly empty levels, and the fill can be many ticks away from the stop price.

12

**Honest Research: Testing Ideas
Without Fooling Yourself**

In this chapter: You will learn to judge any claim that "this strategy works", including your own claims and ours, using the standards of honest research. We cover the classic ways a backtest lies (look-ahead bias, survivorship and data bias), why choosing the best of many tests produces brilliant-looking results from pure luck, and how the Deflated Sharpe Ratio and the Probability of Backtest Overfitting correct for it. We then look at the silent killers of order-flow strategies (impossible fills, slippage, commissions and fake volume), lay out a complete testing protocol (pre-registration, a time-based split with a locked out-of-sample period, walk-forward analysis and realistic costs), and explain forward testing as the first truly honest evidence. Finally, we describe in general terms what failed in our own testing and why we tell you, and we give you a checklist for evaluating any trading educator, including us.

Why Most Backtests Lie: Look-Ahead & Survivorship

A **backtest** is the process of running a strategy's rules over historical data to see what would have happened "if we had traded this way back then". It is the most common tool in trading research and one of the easiest to get wrong. Small mistakes in how a backtest is built can produce a result that looks excellent and is completely unreal.

The worst thing about these mistakes is that they produce no error message. The code runs. The equity curve rises smoothly. Nothing on the screen tells you that the result is impossible.



Figure 32. Three ways a backtest flatters an idea. Each one makes the past look better than the future will be.

Look-ahead bias: reading tomorrow's newspaper

Look-ahead bias means using information in a decision that was not yet available at the moment the decision was made. Imagine betting on yesterday's horse races while holding today's newspaper: you would look like a genius, and you would learn nothing about betting.

In order-flow and volume work, look-ahead hides in many places:

- **Using today's final profile during today.** A rule says "act when price touches today's VAL". If the backtest uses the VAL of the *completed* day, then at 10:00 ET it is using a value area that will only exist at the end of the session. At 10:00 the profile was still developing (Chapter 5), and its VAL was somewhere else.
- **Using the day's high or low, or a candle's close, for a decision made inside that candle.** If a rule checks "did this bar close above the level?" and then enters at a price from the middle of that same bar, it has used the close before the close existed.
- **Detections dated too early.** An algorithm can only label something an "iceberg" after it has seen several refills, and it can only label a "big trade" after it has grouped many small prints together (Chapter 10). If the backtest attaches that label to the moment of the first print, it has given the strategy knowledge it could not have had.
- **Normalising with the whole period.** Scaling a feature using the mean and standard deviation of the entire dataset means every early data point "knows" something about the future.
- **Tuning on all the data, then "testing" on the same data.** If parameters were chosen by looking at the full history, any test on that history is not a test.

The defence is procedural. Every number a strategy uses should carry a timestamp of when it became **knowable**, not when it happened. Profiles, VWAP, highs and lows must be computed "up to that moment". Detections must be recorded at the moment they could really have been detected. Normalisation must use only past data. And it helps a great deal to have someone else review the code with these questions in mind.

A gold illustration

Illustrative, not a real result. Imagine two versions of the same simple rule on GC: "buy when price touches the day's VAL". Version A uses the completed day's VAL. Version B uses yesterday's VAL, or the developing VAL as it stood at that moment. The difference between them is one line of code. Version A will very often look far better, because the completed VAL is, by construction, a price where the day's trading later proved to be important. Version A has quietly looked at the answer.

Survivorship and data bias

Survivorship bias means your data is an unrealistic sample of the past because something has been removed. The classic stock-market example is testing on today's index members only: the companies that went bankrupt and dropped out are missing, so the past looks better than it was.

Gold futures do not have bankrupt companies, but they have their own versions of the problem:

- **Continuous contracts built by back-adjustment.** To join successive contract months into one long chart, a back-adjusted series shifts all past prices by the gaps at each roll (Chapter 1). Past price levels are then no longer the numbers traders actually saw on those days. A strategy based on round numbers or on exact historical levels is being tested on prices that never existed.

- **Ignoring the roll.** Volume migrates from one contract to the next around each roll. A backtest that reads volume from the wrong contract will see "volume drying up" that is really just traders moving to the next month.
- **Testing only the days you remember.** "I tested it on the days I remember the pattern working" is not a sample; it is a highlight reel.

Key idea: A backtest can be wrong without any visible error. Only process saves you: timestamps of when each piece of information was knowable, correct contract handling, and independent review.

Common mistake: "My backtest covers five years of data, so it's reliable." The length of the data does not cure look-ahead. Five years of tomorrow's newspaper is still tomorrow's newspaper.

A related claim, "my indicator doesn't repaint", is a popular claim, not proven, unless you have checked how the indicator calculates each value. An indicator "repaints" if its past values change after new data arrives, which is a form of look-ahead built into the chart itself.

Try it: Take one rule you believe in. Write down every input it uses. Next to each input, write the exact moment that input becomes knowable in real time. If any input becomes knowable after the decision moment, your test of that rule has look-ahead.

Overfitting & Multiple Testing: The "Best of 100" Trap

Overfitting: memorising noise

Overfitting means a strategy has been tuned so closely to the details of past data that it has memorised the noise rather than learned a pattern. It fits the past perfectly and the future poorly.

Order-flow strategies are especially exposed, because the space of choices is enormous: the delta threshold, the size that counts as a "big trade", the time window, the type of level, the time of day, the speed filter, the imbalance ratio, the bar type. Every combination of choices is, in effect, a separate experiment.

The multiple-testing problem

Here is the core problem. If you test 100 strategies that are **completely random**, with no edge at all, some of them will produce very good results by chance. If you report only the best one, you look like a genius. And that best one will almost always perform much worse in the future, because its "edge" was luck.

A coin example makes it concrete. Flip a fair coin ten times. The chance of getting nine or more heads is 11 in 1,024, about 1.1%. Now give 100 people a coin each and have them all flip ten times.

The chance that **at least one** of them gets nine or more heads is about $1 - (1 - 0.011)^{100}$, roughly two in three. The "winning" coin gets a crown, and on its next ten flips it is a fair coin again.

The same thing happens with Sharpe ratios. The **Sharpe ratio** is average return divided by the variability of returns, usually annualised: a measure of return per unit of risk. With two years of daily results, a strategy with zero true edge has a measured annual Sharpe that wanders around zero with a typical spread of about 0.7. The best of 100 such strategies will usually show a Sharpe around 1.7 or higher, purely from noise. On a chart, that looks like a very attractive strategy.

The key discipline is to count **every** experiment, including the ones you "just tried quickly" and the parameter tweaks you made in your head. The number of trials is part of the result.

Tools that correct for it

Deflated Sharpe Ratio (DSR). Developed by David Bailey and Marcos López de Prado. It compares the Sharpe ratio you observed with the highest Sharpe ratio you would *expect* to see by luck alone, given how many trials you ran, how spread out their results were, how skewed and fat-tailed the returns are, and how long the test period was. Its output is, roughly, the probability that the true Sharpe is above zero after all these corrections. A strategy that looks impressive after one trial can look ordinary after 300.

Probability of Backtest Overfitting (PBO). Developed by Bailey, Borwein, López de Prado and Zhu. The data is split into many pieces, which are recombined into many in-sample and out-of-sample pairs. For each pair, the procedure asks: "Did the configuration that was best in-sample end up below the median out-of-sample?" The fraction of pairs where the answer is yes is an estimate of how likely your *selection process* is to pick an overfit winner.

Statistical multiple-testing corrections. Methods such as Bonferroni and Holm, or controlling the false discovery rate with Benjamini–Hochberg, make the significance threshold stricter as the number of tests grows. In finance, Campbell Harvey, Yan Liu and Heqing Zhu argued that, given how many factors researchers have tried, a new "discovery" should need a t-statistic above about 3.0 rather than the traditional 2.0.

None of these tools is magic. Every one of them depends on an honest count of how many things were tried. If you tried 300 combinations and tell the formula you tried 3, the formula will happily give you a wrong answer.

Key idea: Flip 100 coins ten times each and one of them will look like a genius. That is what most backtests are: the luckiest coin, presented as skill.

Common mistake: "I only have three parameters, so I can't be overfitting." What matters is how many *versions* you tried, not just how many parameters the final version has. Three parameters with ten values each is 1,000 versions.

"A backtest Sharpe above 2 means the strategy is real" is a popular claim, not proven. Without knowing how many trials produced it, a Sharpe number on its own tells you very little.

Try it: Look back at the last idea you tested. Count honestly how many variations you tried: different thresholds, time windows, filters, markets, bar types. Write the number at the top of your notes. That number belongs in every report of the result.

Impossible Fills, Costs & Fake Volume

Order-flow strategies are usually **short-term**: small objectives measured in a few ticks, and many trades. In this world, costs and the quality of fills decide almost everything. A small error in how fills are modelled can be larger than the entire edge being measured.

1. Impossible fills

"My limit order filled when price touched it." On CME, orders at the same price are generally filled in time priority: first in, first out (FIFO). A touch is not a fill (Chapter 10). If price touches your limit price and then turns away, the orders ahead of you in the queue may have absorbed all the trading at that price. A conservative backtest only fills a limit order when price trades *through* it by at least one tick, or uses a model of queue position.

The unknown path inside the bar. A bar has an open, high, low and close, but many different paths can produce the same four numbers (Chapter 1). If both your stop and your objective sit inside one bar's range, the bar does not tell you which was hit first. A backtest that assumes the favourable one has invented a fill.

More volume than really traded. A backtest that fills 50 contracts at a price where only 12 traded has filled orders that the market could not have absorbed.

Zero latency. A signal that depends on a bar closing, or on data processing, is only visible *after* that happens. Your real reaction time, and your order's travel time, are never zero. A backtest that enters at the exact price of the signal moment assumes instant reaction.

2. Slippage

Market orders and stop-market orders fill worse than the expected price in fast markets, especially around news (Chapter 11). A conservative assumption is at least one tick of slippage per side for market and stop orders, and more around major releases.

3. Fixed costs

Commission, exchange and clearing fees, regulatory fees, and data and platform costs. Express them **in ticks**, because that makes them comparable with the size of your objective.

Illustrative arithmetic. Check your own broker's costs.

Suppose all-in commission and fees for a round trip on one MGC are about \$2. On MGC a tick is \$1, so that is **2 ticks** per round trip. Add one tick of slippage on exits by stop.

Now take a simple strategy with a 6-tick objective and a 6-tick stop:

- Gross: a win is +6 ticks, a loss is -6 ticks. You break even if you are right half the time.
- Net: a win is $6 - 2 = +4$ **ticks**; a loss is $-6 - 2 - 1 = -9$ **ticks**. To break even you now need to be right about $9 \div (4 + 9) \approx 69\%$ of the time.

Nothing about the market changed. Only the accounting did. On GC, where a tick is worth \$10, the same dollar commission is a much smaller fraction of a tick, which is one reason why costs per tick differ so much between the two contracts even though the price is the same.

Think of it as a waterfall: gross result, minus commission and fees, minus slippage, minus the fills you assumed but would not have received, equals the net result. For many short-term ideas, the waterfall drains most of the gross.

4. Tick volume versus real volume

In decentralised markets, such as spot forex or gold CFDs (XAUUSD), the "volume" on most charts is **tick volume**: the number of price changes in your broker's feed, not the number of contracts traded. A trade of 1,000 units and a trade of 1 unit each count as one tick, if they count at all. A backtest of "delta" or "footprint" patterns on such data does not measure what those words mean in futures (Chapter 2). That is why this book reads GC and MGC data: a central market with real volume and an aggressor side for every trade.

The recurring pattern

In published research and in our own testing, one pattern repeats: many order-flow signals show "something" before costs and then disappear after costs and realistic fills for a non-high-frequency trader. The edges that exist at the scale of milliseconds tend to belong to the fastest participants, whose costs and latencies are far below a retail trader's.

Key idea: For short-term strategies, costs are destiny. Always write the cost in ticks next to the objective in ticks.

Common mistake: "Commission is negligible." On MGC, with objectives of a few ticks, it is not.

"My platform runs backtests realistically" is a popular claim, not proven, unless you have checked its fill assumptions. And "a footprint on a gold CFD is the same as a footprint on gold futures" is simply false: the CFD footprint is built from estimated data from one broker's feed.

Try it: For any short-term idea you have, write three numbers: objective in ticks, stop in ticks, and round-trip cost plus expected slippage in ticks. Compute the break-even hit rate before and after costs. If the gap shocks you, that is the lesson.

How to Test an Order-Flow Idea Properly

Here is a testing protocol suitable for teaching and close to what we use in our own research. It is not the only valid protocol, but every step addresses one of the traps of the previous sections.

Step 1 — Pre-registration

Before you see any results, write a short document and date it. It contains:

- **Hypothesis.** For example: "Absorption of selling at yesterday's VAL, with confirmation X, is followed in the next 30 minutes by a positive average R after costs."
- **Exact rules.** Entry, exit, stop and size, defined so that a computer (or another person) could apply them without judgement.
- **Data.** Contract, period, session definition, how the roll is handled.
- **Success criteria.** What result would count as success: for example, a Deflated Sharpe above a stated level and a minimum number of trades.
- **Maximum number of variants.** How many versions you are allowed to try in total.

Pre-registration stops you from **moving the goalposts**: deciding what counts as success after you have seen what you got.

Step 2 — Split the data by time

- **Train:** the period used to build and tune the idea.
- **Validation:** a later period used to choose between a small number of versions.
- **Locked out-of-sample (OOS):** the most recent period, which you **never look at** until the final version is frozen. It is run **once**. If you change the idea after seeing the OOS result, the OOS is no longer out-of-sample; it has become part of the training data.

The split is always **chronological** (past → future), never random. Randomly assigning days to train and test lets information leak between neighbouring days, because markets have memory over hours and days. It is also good practice to leave a small gap, an **embargo**, between periods so that overlapping effects (for example, a trade that opens at the end of one period and closes in the next) cannot leak.

Period	Purpose	How often you look
Train	Build and tune	As often as needed
Embargo	Gap against leakage	Never used
Validation	Choose among a few versions	A limited number of times
Embargo	Gap against leakage	Never used
Locked OOS	Final, honest test	Once

Step 3 — Walk-forward

Instead of one split, move a window forward through time: tune on period 1, test on period 2; tune on periods 1 + 2, test on period 3; and so on. Then put all the test periods side by side. Walk-forward shows whether the idea is stable across different market regimes: quiet and volatile, trending and balanced, different rate environments.

Step 4 — Execution realism from day one

Fills based on the queue, or on a conservative assumption (filled only when traded through). Slippage. Commission and fees. Latency. These are not refinements to add at the end; they belong in the first version, because an idea that only works without them is not an idea.

Step 5 — Correct statistics

Record the total number of trials. Use the Deflated Sharpe Ratio, PBO or a multiple-testing correction. Report the distribution of R, the maximum drawdown and the number of trades, not just the share of winning trades. A high share of winners with occasional large losers can still have negative expectancy.

Step 6 — A negative result is a result

If the idea fails, document it, and if you publish research, publish the failure too (see "What Failed in Our Testing" below). A literature that contains only successes is itself biased.

Key idea: Lock part of your data in a vault and open it once. If you peek and adjust, the vault is empty.

Common mistake: "Just one small look" at the out-of-sample results, followed by "just one small change". After that, the OOS is contaminated; it is effectively in-sample.

"A good walk-forward result guarantees the future" is a popular claim, not proven. Walk-forward provides stronger evidence than a single split; it does not provide certainty.

Try it: Write a pre-registration form for one idea using the five fields above: Hypothesis, Rules, Data, Success criteria, Maximum variants. Date it. Do not run anything until it is complete.

Forward Testing: The First Honest Evidence

Even the most careful backtest was run on data that the researcher knew something about, directly or indirectly, while building the idea. You may not have looked at the locked out-of-sample period, but you lived through those months; you remember how gold behaved. A **forward test** runs **frozen** rules on data that **did not exist yet** when the idea was built. It is the first evidence that is genuinely about the future.

Two forms

Form	How it works	Strength	Weakness
Paper / simulated	Orders go to a simulator fed by live data	Cheap; no money at risk	Fills are optimistic: a simulated order is not really in the queue
Small live	Real orders at minimum size (for example, one MGC)	Real fills, real slippage, real psychology	Costs real money, even if small

Paper testing is a useful first step; small live testing reveals what the simulator hides.

A protocol for forward testing

- **Frozen rules.** No changes at all during the forward period. Any improvement idea goes into a "next version" list, to be tested later as a new idea.
- **Duration and sample size decided in advance.** For example, a minimum number of trades and a minimum number of months, whichever comes later. Conclusions from ten trades are close to meaningless.
- **An expected band from the backtest.** Before starting, write down the range you expect for average R, trade frequency and drawdown, based on the backtest. The forward results should fall inside this band. If they are clearly below it, either the edge was never there or the live execution is different from the test.
- **Separate tracking of execution.** Compare actual slippage with assumed slippage. Compare the actual number of signals with the expected number. If the live system produces far fewer (or far more) signals than the backtest predicted, the live data or code probably differs from the backtest, or the backtest had look-ahead.
- **Stop criteria decided in advance.** For example: "If drawdown exceeds a stated number of R, or if average R after a stated number of trades is below a stated level, stop and review."

Honest limits

A forward test is also a limited sample, and a good period can be luck. But the combination of a strict backtest and a frozen forward test is far stronger evidence than either alone. The backtest says "this could work"; the forward test is the first time the market itself gets a vote.

Key idea: A backtest is a story about the past. A frozen forward test is the first time the market gets a vote.

Common mistake: Changing the rules after three bad trades in the forward test. That resets the test; the new rules have zero forward evidence.

"Two profitable weeks on paper means you are ready for a large account" is a popular claim, not proven. Two weeks is a tiny sample, and paper fills are optimistic.

Try it: For an idea you are working on, write an "expected band" card: average R range, trades per week range, worst expected drawdown in R. Then write the stop-and-review criteria. Pin both to your screen before the forward test begins.

What Failed in Our Testing (And Why We're Telling You)

This section is framed explicitly as **"in our testing"**: results on our own data, with our own methods. They are not universal truths. They were obtained with a protocol like the one described above (pre-registered hypotheses, locked out-of-sample periods, realistic costs), and we describe them in general terms.

Example 1 — The side of an iceberg as a direction signal

The idea is attractive: "If a buy-side iceberg is detected, price will rise in the following minutes." In our testing, the side of a detected iceberg did not predict the direction of the next move in any useful way when run live. The likely reasons were covered in Chapter 10. First, an iceberg can only be detected after it has refilled several times, so by the time it is labelled, part of its story is over. Second, synthetic icebergs, built by a trading platform rather than the exchange, look just like ordinary orders and create ambiguity. Third, and most fundamentally, an iceberg is evidence of **hidden size** at a price, not evidence of its owner's overall direction. A large passive buyer at one price might be hedging, executing one leg of a spread, or completing an order that has nothing to do with the next five minutes.

Example 2 — Naive big-trade detection

The idea: "Large prints mean large players." In our testing, a naive detector that flagged large prints directly from the raw trade feed was wrong much more often than it was right: a large share of its flags were not single large orders at all. The reason is the one explained in Chapter 10: one real order often prints as many small fills, so the raw sequence of prints is a poor picture of the orders behind it. Grouping prints into the orders behind them improved the quality of detection considerably. But even when large orders were identified correctly, "a big trade happened" was not, on its own, a profitable signal after costs in our tests.

The broader pattern

Across the published literature and our own work, the same pattern appears again and again: many order-flow signals that show "something" before costs vanish after costs and realistic fills for a trader who is not operating at high frequency. Edges measured in seconds or less tend to belong to faster competitors.

The common cause in both examples is worth stating plainly: **what we see on the screen is not the same as the real event behind it.** Prints are not orders. A visible iceberg is not an intention. The

data is a shadow of the decisions that created it.

So what is left?

Order flow as **confirmation, timing and veto** at **levels chosen in advance**, as described in Chapter 11. The veto role, declining trades that the flow contradicts, is especially valuable, because it reduces the number of trades and therefore the costs. This, too, is a hypothesis, and anyone using it should test it on their own data with the protocol of this chapter.

Why publish failures?

Because publishing only successes is survivorship bias at the level of content. If every educator shows only what worked, the overall picture the public sees is a highlight reel of the luckiest coins. An educator's credibility comes, in large part, from telling you what did not work.

Key idea: The right conclusion is not "order flow is useless". It is that order flow changes role: from signal generator to filter in context.

Common mistake: "These ideas failed for you because you didn't use them correctly." That is possible. That is exactly why we describe the method and its limits openly and encourage independent testing.

"Commercial tool X has solved this problem" is a popular claim, not proven, unless it is backed by an independent test with genuine out-of-sample evidence.

Try it: Pick one order-flow idea you find attractive. Write down the hidden assumption behind it (for example, "a large print is one large order" or "an iceberg's side is its owner's direction"). Then write how you would test whether that assumption is true before testing the trading idea itself.

How to Evaluate Any Trading Educator (Including Us)

The market for trading education is full of claims that cannot be checked. The goal of this section is not blind cynicism. It is to give you a standard set of questions, and to ask you to apply them to this book as well.

Official red flags

The US Commodity Futures Trading Commission (CFTC) publishes warnings about trading education and advice. Among the red flags they describe:

- Promises of large profits, or of a method that cannot lose or is "guaranteed".
- "Simple secrets and tricks" for a complex market.
- Claims of exceptional results that you have no way to verify.

- Time pressure: "only a few seats left", "price goes up tonight".
- A free seminar that ends in an aggressive sales pitch for an expensive course.
- Large payments demanded before you can judge the quality of what you are buying.

Simulated results have built-in limits

In the United States, CFTC Rule 4.41 requires that advisers it covers who present hypothetical or simulated performance must accompany it with a specific warning. In substance, the warning says that simulated results have inherent limitations, do not represent actual trading, and are designed with the benefit of hindsight. Whenever you see a "backtest" or "system results" in marketing, remember that warning, whether or not it is printed.

The question checklist

1. Is the claim **falsifiable**: specific and precisely defined, so that it could be shown to be wrong?
2. Are the results **real and verifiable**, or selected screenshots?
3. **How many ideas or parameter sets** were tried to arrive at this result?
4. Is there **out-of-sample or forward** evidence?
5. Are **costs and slippage** included?
6. Are **failures and limits** discussed?
7. Where does the educator's **income** come from: courses, tools, broker or prop-firm commissions, or trading? Is there a conflict of interest?
8. Are terms **defined**, or are they attractive labels ("smart money", "institutional") without a definition?
9. Do the examples include **ordinary and bad days**, or only spectacular ones?
10. Does the educator encourage you to **test things yourself**, or only to trust them?

A comparison exercise

Imagine two anonymous posts about gold (fictional, no real person or brand). The first says: "This GC order-flow setup wins 90% of the time!!" and shows one selected screenshot. The second gives a precise definition of the pattern, the test period, the number of variations tried, the cost assumptions, and one example where it failed. Score both with the checklist. The first scores close to zero; the second may not be right, but it gives you everything you need to find out.

About us

This book makes no profit or win-rate claims. It labels folklore as folklore ("a popular claim, not proven"). It describes what failed in our testing and explains the methods behind our conclusions. If you find a place where we break these principles, use this same checklist to criticise us.

Three myths to finish. "He drives an expensive car, so his strategy works": the income may come from selling courses. "A P&L screenshot is proof": screenshots can be selected, faked, and are not representative. "X years of experience means the method is right" is a popular claim, not proven: experience and evidence are different things.

Key idea: A good claim is specific, falsifiable and verifiable. Ask how many things were tried, whether there is out-of-sample evidence, and whether costs were included.

Common mistake: Treating confidence as evidence. The most certain-sounding educators are often the ones with the least verifiable evidence.

Try it: Score the next three trading posts you see on social media with the ten-question checklist. Then score this book. Write down where each one loses points.

Chapter summary

- A backtest can be badly wrong without any error message. **Look-ahead bias** (using information not yet available) is the most common culprit: final profiles used intraday, detections dated too early, normalisation with future data.
- **Survivorship and data bias** in gold futures come from back-adjusted continuous contracts, mishandled rolls and testing only memorable days.
- **Overfitting** memorises noise. Testing many variants and reporting the best produces impressive results from pure luck; every trial must be counted.
- The **Deflated Sharpe Ratio**, the **Probability of Backtest Overfitting** and multiple-testing corrections adjust for the number of trials, but only if that number is reported honestly.
- **Impossible fills**, slippage, commission and fees can turn a short-term edge negative; express all costs in ticks. **Tick volume** from CFD feeds is not real volume.
- A proper protocol: **pre-registration**, a chronological train / validation / **locked out-of-sample** split with embargoes, **walk-forward**, realistic execution, correct statistics, and publication of negative results.
- A **forward test** with frozen rules, a sample size set in advance, an expected band and stop criteria is the first genuinely prospective evidence.
- In our testing, popular order-flow ideas used as stand-alone direction signals did not survive; what remains is order flow as **confirmation, timing and veto** at pre-chosen levels, itself a hypothesis to test.
- Evaluate any educator, including us, with falsifiability, verifiability, number of trials, out-of-sample evidence, costs, failures shown and conflicts of interest.

Checklist

- Every input to my rule has a timestamp of when it became knowable, and none is later than the decision.
- I test on the real contract with the roll handled correctly, not on back-adjusted prices for level-based ideas.

- I have counted every variant I tried, including quick experiments.
- My limit orders fill only when traded through (or by a queue model), and I do not assume the favourable path inside a bar.
- Commission, fees and slippage are expressed in ticks and included from the first version.
- I have a dated pre-registration with hypothesis, rules, data, success criteria and maximum variants.
- My data split is chronological, with a locked out-of-sample period I will run once.
- I have an expected band and stop criteria written before any forward test.
- I record failures as carefully as successes.
- I apply the ten-question checklist to every educator, including this book.

Quiz

1. At 10:00 ET you use today's final POC as a level in your backtest. What kind of error is this?
2. You tried 300 parameter combinations and report the best one. What else must you report, and why?
3. In spot forex or gold CFD charts, what does "volume" usually count?
4. You peeked at your out-of-sample results and changed one rule. What is the status of your out-of-sample period now?
5. In our testing, why did naive big-trade detection from raw prints perform poorly?

Quiz answers

1. Look-ahead bias. Today's final POC only exists at the end of the session; at 10:00 ET the profile was still developing.
2. The number of combinations tried (and ideally the spread of their results). Without it, nobody can judge how much of the best result is luck; tools like the Deflated Sharpe Ratio need that number.
3. The number of price changes (ticks) in that broker's feed, not the number of contracts traded.
4. It is contaminated. Because you adjusted the idea after seeing it, it has effectively become in-sample data, and it can no longer serve as an honest test.
5. Because one real order often prints as many small fills, so individual prints are a poor picture of the orders behind them, and many of the flags did not correspond to single large orders. Even correctly identified big trades were not, on their own, a profitable signal after costs.

REFERENCE

A

**Glossary, Pre-Session
Checklist, GC→XAUUSD
Conversion and Contract Cheat
Sheet**

In this appendix: A glossary of every important term in the book, from A to Z, with a pointer to the chapter where it is explained in full. A one-page pre-session checklist you can print and use before every gold session. A step-by-step explanation of how to convert gold futures (GC) levels into approximate XAUUSD (CFD) levels, with a worked example. And a cheat sheet of the key facts about the GC and MGC contracts. Contract facts were checked in October 2026; exchanges change specifications, margins and hours from time to time, so always re-check the official CME Group pages before relying on any number here.

Glossary A–Z

Chapter numbers in brackets show where each term is explained in depth. "ET" means New York time throughout.

A

Absorption — Heavy aggressive trading at a price that fails to move the price, because large passive orders on the other side take it all in. "Effort without result." A popular reading is that absorption leads to a reversal; that is a popular claim, not proven. [Ch 9]

Acceptance — The market trading at a price with both **time and volume**, showing that participants agree the price is fair for now. The opposite of rejection. [Ch 7]

Active month (lead month, front month) — The contract month of a futures product where most volume and liquidity are concentrated at a given time. Order-flow analysis should be done on the active month, not on a stale one. [Ch 1]

Aggressor — The side of a trade that crossed the spread with a market order (or a marketable limit order) and "hit" a resting order. Buys at the ask are aggressive buys; sells at the bid are aggressive sells. [Ch 2, Ch 3]

Anchored VWAP — A VWAP that starts its calculation at a chosen event (a major data release, a swing high or low, the start of a week) instead of at the session open. [Ch 4]

Ask (offer) — The lowest price at which someone is currently willing to sell. You buy at the ask with a market order. [Ch 0, Ch 2]

Asia session — The part of the gold trading day from the 18:00 ET reopen until around the London open; usually lower volume, with activity from Tokyo, Singapore and Shanghai. Session boundaries are conventions and vary slightly between sources. [Ch 1, Ch 7]

Auction Market Theory (AMT) — A way of understanding markets as continuous two-sided auctions whose purpose is to facilitate trade: price moves to find buyers and sellers, alternates between balance and imbalance, and "advertises" prices until trade is accepted or rejected. [Ch 7]

B

Back-adjustment — A method of building a continuous futures chart by shifting past prices by the gap at each roll, so the series has no artificial jumps. It changes historical price levels and does not correct volume or delta. [Ch 1, Ch 12]

Backtest — Running a strategy's rules over historical data to see what would have happened. Easy to get wrong without any visible error. [Ch 12]

Balance — A market state where price rotates within a range (a bracket) and both buyers and sellers find the prices acceptable. Compare imbalance. [Ch 7]

Bar delta — Aggressive buy volume minus aggressive sell volume within one bar. A bar with positive delta can still close lower. [Ch 3]

Base rate — How often an outcome happens anyway, without any pattern or signal. A pattern is only informative if what follows it differs from the base rate after costs. [Ch 9]

Basis (GC–XAUUSD gap) — In this book, the difference between the gold futures price (GC, active month) and the spot or CFD price (XAUUSD). In October 2026 it was roughly \$27–30, with GC above XAUUSD. It drifts and must be measured. See the conversion section below. [Ch 0, Appendix]

Bid — The highest price at which someone is currently willing to buy. You sell at the bid with a market order. [Ch 0, Ch 2]

Big trade — A large print on the tape. One real large order often prints as many small fills, so large prints and large orders are not the same thing. [Ch 10, Ch 12]

Bracket — A trading range in which the market is in balance. [Ch 7]

b-shaped profile — A volume profile with most volume at the bottom and a thin upper part, often associated with long liquidation (longs exiting) rather than fresh selling. [Ch 6]

C

CFD (Contract for Difference) — An agreement with your broker to settle the difference between the opening and closing price. You own nothing; the broker (or its partner) is the other side; volume shown is usually tick volume from that broker's feed. XAUUSD is the common gold CFD. [Ch 0]

Clearing house — The central counterparty that guarantees both sides of every exchange-traded futures trade and handles daily mark-to-market. [Ch 0, Ch 1]

COMEX — The CME Group exchange where gold futures (GC, MGC) are listed. [Ch 0, Ch 1]

Composite profile — A volume profile built over several sessions (a week, a month, or a custom range) to show longer-term value. [Ch 5]

Confluence — Two or more independent reasons pointing to the same price zone, such as yesterday's VAL coinciding with the London low. Used to rank levels. [Ch 11]

Continuous contract — A chart that joins successive futures contracts into one series (for example "GC1!" or "@GC"). Useful for long-term price views; unreliable for exact historical levels, volume or delta around rolls. [Ch 1]

Cost of carry — The cost of holding gold until a future date (mainly interest, minus any lease income). It is why the futures price normally sits above the spot price, and why the GC-XAUUSD gap exists. [Appendix]

CPI (Consumer Price Index) — The monthly US inflation report, usually released at 08:30 ET. A major news event for gold. [Ch 11]

Cumulative Volume Delta (CVD) — The running total of delta over a session or another chosen period. Its value depends on where it starts (its anchor). [Ch 3]

D

Daily loss limit — A pre-set loss, usually in R, after which you stop trading for the day. [Ch 11]

Deflated Sharpe Ratio (DSR) — A Sharpe ratio adjusted for the number of trials, the spread of their results, non-normal returns and test length; roughly, the probability that the true Sharpe is above zero. Developed by Bailey and López de Prado. [Ch 12]

Delta — Aggressive buy volume minus aggressive sell volume over a period. It measures who was more urgent, not who was "right". [Ch 3]

Delta-POC — The price inside a footprint bar with the largest absolute delta: where one-sided aggression was most concentrated. A descriptive metric, not a signal. [Ch 9]

Developing profile — A volume profile (and its POC, VAH and VAL) as it stands at a given moment during the session, before the session is complete. Using the final profile during the day is look-ahead bias. [Ch 5, Ch 12]

Diagonal imbalance — A footprint imbalance computed by comparing the ask volume at one price with the bid volume one tick lower (or the bid at one price with the ask one tick higher), because those two sides faced each other at the same moment. [Ch 8]

DOM (Depth of Market) — A ladder display of resting limit orders at each price above and below the market, usually with the ability to place orders directly. [Ch 10]

Double distribution — A profile with two separate areas of value connected by a thin, low-volume section, typical of a day that moved from one value area to another. [Ch 6, Ch 7]

Drawdown — The fall in account value from its highest point. Prop firms use static, end-of-day trailing or intraday trailing drawdown floors. [Ch 11]

E

Embargo — A gap left between training, validation and test periods so that overlapping information cannot leak from one period into another. [Ch 12]

Excess — A decisive end to an auction in one direction: price overshoots, meets strong opposition and leaves quickly, leaving a tail on the profile. [Ch 7]

Exhaustion — Aggressive volume thinning out at the extreme of a move: the side pushing price runs out of urgency. [Ch 9]

Expectancy — The average result per trade, best measured in R over many trades. More informative than the share of winning trades. [Ch 11]

F

False positive — A signal that fires when the event it is supposed to detect did not happen. Every pattern has a false-positive rate that must be measured. [Ch 9]

FIFO (first in, first out) — A matching rule where resting orders at the same price are filled in the order they arrived. Under FIFO, price touching your limit price does not mean your order filled. [Ch 10]

Fill — The execution of an order, in whole or in part. [Ch 10, Ch 12]

FOMC — The Federal Open Market Committee, which sets US interest rates. Statements are usually released at 14:00 ET, followed by a press conference usually at 14:30 ET. [Ch 11]

Footprint chart — A chart that shows, for each price inside a bar, how much traded at the bid and how much at the ask, so you can see where and how aggressively trading happened within the bar. [Ch 8]

Forward test — Running frozen rules on data that did not exist when the idea was built, either on paper (simulated) or live at minimum size. The first genuinely prospective evidence. [Ch 12]

G

GC — The standard COMEX gold futures contract: 100 troy ounces, tick \$0.10 = \$10 per contract. [Ch 1]

Globex — CME Group's electronic trading platform. Gold trades on Globex from Sunday 18:00 to Friday 17:00 ET with a daily halt from 17:00 to 18:00 ET. [Ch 1]

H

Heatmap (liquidity heatmap) — A chart that colours resting order-book size at each price over time: the order book with memory. Shows where liquidity was added, held or pulled. [Ch 10]

High Volume Node (HVN) — A price area on a profile where a lot of volume traded; a zone of agreement where price tends to slow down. [Ch 5]

Hindsight bias — Believing, after the fact, that an outcome was predictable, or rewriting your reasons to match the outcome. Writing plans and journal entries before outcomes are known guards against it. [Ch 11]

I

Iceberg order — A large limit order that shows only a small part (the display quantity) in the book and refills from a hidden reserve as it trades. **Native** icebergs are managed by the exchange; **synthetic** icebergs are managed by a platform or algorithm and look like ordinary new orders. [Ch 10]

Imbalance (footprint) — A footprint row where one side's aggressive volume exceeds the other side's by a chosen ratio (commonly between 1.5 and 4 times) and above a minimum volume. [Ch 8]

Initial balance (IB) — The price range of the first hour of the main session. In this book, for gold, 08:20–09:20 ET (the first hour of the traditional COMEX day session). [Ch 7]

Initial margin — The deposit required to open a futures position; collateral, not the price of the contract. Set by the exchange (with brokers often adding their own requirements) and changed from time to time. [Ch 0, Ch 1]

L

Layering — Placing several orders at multiple prices to create a false impression of supply or demand, with the intent to cancel them. A form of spoofing; illegal. [Ch 10]

Lee-Ready algorithm — A method from 1991 for classifying trades as buyer- or seller-initiated when the aggressor side is not reported: compare the trade with the bid–ask midpoint, and use the tick rule for trades exactly at the midpoint. [Ch 3]

Leverage — Controlling a large position (notional value) with a small deposit. It magnifies both gains and losses; futures losses can exceed the deposit. [Ch 0]

Limit order — An order to buy or sell at a specified price or better. It rests in the book (passive) until filled or cancelled. [Ch 2]

Liquidity — How easily you can trade a meaningful size without moving the price much; in practice, the resting size near the current price and the speed with which it is replenished. [Ch 1, Ch 10]

London session — The part of the gold day from about 03:00 ET, when London, the centre of the physical gold market, is active. The London–New York overlap is usually the most active part of the day. [Ch 1, Ch 7]

Long — A position that profits if the price rises. You buy first and sell later to close. [Ch 0]

Look-ahead bias — Using information in a decision (or in a backtest) that was not yet available at the moment of that decision. [Ch 12]

Low Volume Node (LVN) — A price area on a profile with little volume, where the market moved quickly; price often travels through such areas fast. [Ch 5]

M

MAE / MFE — Maximum adverse excursion (how far a trade went against you) and maximum favourable excursion (how far it went in your favour), usually recorded in R. Used to evaluate stops

and exits. [Ch 11]

Maintenance margin and margin call — The minimum account balance needed to keep a futures position open. Falling below it triggers a margin call or forced liquidation. [Ch 1]

Market by Order (MBO, Level 3) — Market data that shows every individual order in the book (add, modify, cancel, fill), each with an anonymous ID and queue position. [Ch 10]

Market by Price (MBP, Level 2) — Market data that shows the total resting size at each of several prices, without breaking it down into individual orders. [Ch 10]

Market order — An order to buy or sell immediately at the best available price. It crosses the spread and is the aggressor. [Ch 2]

Mark-to-market — Daily settlement of futures profits and losses based on the official settlement price: losers pay winners every day through the clearing house. [Ch 1]

MGC (Micro Gold) — The COMEX micro gold futures contract: 10 troy ounces, tick \$0.10 = \$1 per contract, one tenth of GC. Separate order book and separate volume from GC. [Ch 1]

Multiple testing — Testing many variants of an idea. The more variants you test, the more likely the best one looks good by chance. Every trial must be counted. [Ch 12]

N

Naked POC (virgin POC) — A point of control from an earlier session that price has not traded at since. "Naked POCs always get filled" is a popular claim, not proven. [Ch 6]

New York session — The part of the gold day from about 08:20 ET (the traditional COMEX day session) to settlement, including 08:30 ET US data and the 09:30 ET stock-market open. [Ch 1, Ch 7]

NFP (Non-Farm Payrolls) — The monthly US Employment Situation report, usually released on the first Friday of the month at 08:30 ET. [Ch 11]

Notional value — The full value of a futures position: price × contract size × number of contracts. One GC at \$4,200 has a notional value of \$420,000. [Ch 0, Ch 1]

O

Opening types — James Dalton's classification of how a session starts: Open-Drive, Open-Test-Drive, Open-Rejection-Reverse and Open-Auction, from most to least conviction. [Ch 7]

Order book — The list of all resting limit orders waiting at each price, bids below the market and asks above it. [Ch 2, Ch 10]

Order Flow Imbalance (OFI) — A research measure (Cont, Kukanov and Stoikov, 2014) of the imbalance between supply and demand pressure at the best bid and ask. It explains a large part of contemporaneous short-term price changes; explaining is not predicting. [Ch 10]

Other timeframe (OTF) participant — A longer-term participant (fund, commercial, institution) whose activity moves price out of balance; contrasted with short-term, day-timeframe traders. [Ch 7]

Out-of-sample (OOS) — Data not used to build or choose a strategy. A **locked** OOS period is examined only once, after the final version is frozen. [Ch 12]

Overfitting — Tuning a strategy so closely to past data that it captures noise instead of a real pattern. [Ch 12]

Overnight high and low (ONH/ONL) — The highest and lowest prices of the overnight (Asia and London) part of the session, before the New York day session. [Ch 7]

P

Paper trading — Simulated trading with live data and no real orders. Cheap, but fills are optimistic because simulated orders are not really in the queue. [Ch 12]

PBO (Probability of Backtest Overfitting) — An estimate of how likely a strategy-selection process is to choose a configuration that was best in-sample but below median out-of-sample. Developed by Bailey, Borwein, López de Prado and Zhu. [Ch 12]

PDH / PDL — Prior day's high and prior day's low. Common reference levels. [Ch 7]

POC (Point of Control) — The price with the most traded volume in a profile: the price of greatest agreement. [Ch 5]

Poor high / poor low — A session extreme without excess: the profile is flat at the edge, suggesting the auction in that direction was not completed. Often marked as a future reference; "poor highs always get retested" is a popular claim, not proven. [Ch 7]

Position — A trade that is currently open. Profit or loss is final only when the position is closed. [Ch 0]

Position sizing — Choosing the number of contracts so that being wrong costs a fixed, acceptable amount: $\text{allowed risk} \div (\text{stop distance in ticks} \times \text{tick value} + \text{costs})$. [Ch 11]

Pre-registration — Writing down and dating the hypothesis, rules, data, success criteria and maximum number of variants before seeing any test results. [Ch 12]

Prop firm — A company that gives traders access to a funded account after an evaluation, usually with strict drawdown and daily-loss rules. [Ch 11]

P-shaped profile — A volume profile with most volume at the top and a thin lower part, often associated with short covering (shorts exiting) rather than fresh buying. [Ch 6]

Pulling and stacking — Resting orders being cancelled as price approaches (pulling) or added in front of price (stacking). Visible on the DOM and heatmap; only trades confirm that liquidity was real. [Ch 10]

Q

Queue position — Your place in line among resting orders at the same price. Under FIFO, orders ahead of yours fill first. [Ch 10]

R

R / R-multiple — R is your initial risk on a trade (what you lose if the stop is hit). Results are expressed as multiples of R: $-1R$ is a full loss; $+2R$ is a gain of twice the risk. [Ch 11]

Range extension — Price moving beyond the initial balance range, often a sign that other-timeframe participants have become active. [Ch 7]

Real volume — The actual number of contracts traded on a central exchange, as in futures. Compare tick volume. [Ch 2]

Rejection — Price moving away quickly from a level with little time or volume spent there; the market refused those prices. [Ch 7]

Relative volume — Volume compared with the typical volume at the same time of day on previous days. Gold's volume has a strong daily pattern, so raw volume comparisons across hours are misleading. [Ch 3]

Revenge trading — Trading to "win back" losses, usually with larger size, looser criteria and no plan. Daily loss limits exist largely to prevent it. [Ch 11]

Roll — Moving a futures position from the expiring contract month to the next one. Volume migrates during roll periods. [Ch 1]

RTH (Regular Trading Hours) — For gold in this book, the New York day session of about 08:20–13:30 ET, as opposed to the full 23-hour Globex session. [Ch 4, Ch 7]

S

Settlement price — The official daily price used for mark-to-market. For GC it is determined from trading in the 13:29–13:30 ET window, hours before the 17:00 ET daily close. [Ch 1]

Sharpe ratio — Average return divided by the standard deviation of returns, usually annualised: return per unit of variability. [Ch 12]

Short — A position that profits if the price falls. You sell first and buy back later to close. [Ch 0]

Single prints — In a TPO profile, prices traded during only one time period, showing the market moved through quickly. In the middle of a profile they form a "bridge" between value areas; at an extreme they form a tail. [Ch 7]

Slippage — The difference between the price you expected and the price you actually got, usually on market and stop orders in fast markets. [Ch 1, Ch 11, Ch 12]

Spoofing — Placing orders with the intent to cancel them before they execute, to mislead other participants about supply or demand. Illegal in US futures markets. [Ch 10]

Spread — The difference between the best ask and the best bid. A cost every time you cross it. [Ch 0, Ch 1]

Stacked imbalance — Several imbalances on the same side in consecutive price rows of one footprint bar (commonly three or more). "Three stacked imbalances make a certain level" is a popular claim, not proven. [Ch 8]

Stop order — An order that becomes active when a trigger price is reached. A stop-market order then executes as a market order and can suffer slippage. Stop orders are not visible in the book until triggered. [Ch 10, Ch 11]

Survivorship bias — A distorted sample of the past caused by leaving out what did not "survive": delisted companies, failed strategies, forgotten days, unpublished failures. [Ch 12]

Sweep — One aggressive order (or a rapid burst) taking out liquidity at several consecutive prices in a fraction of a second. Often associated with stop runs. [Ch 10]

T

Tail — Single prints at the extreme of a session profile, showing excess and the end of an auction in that direction. [Ch 7]

Tape (Time & Sales) — The running list of every trade: time, price, size and (in futures) aggressor side. [Ch 10]

Tape speed — How fast trades are printing. Bursts and quiet periods carry information about urgency. [Ch 10]

Tick — The minimum price increment of a contract. For gold futures, \$0.10 per ounce. [Ch 1]

Tick rule — A method of classifying a trade as a buy or sell by comparing its price with the previous trade's price (uptick = buy, downtick = sell). [Ch 3]

Tick size — See Tick. \$0.10 for both GC and MGC. [Ch 1]

Tick value — The dollar value of a one-tick move for one contract: \$10 for GC, \$1 for MGC. [Ch 1]

Tick volume — The number of price changes in a feed over a period, used as a stand-in for volume in decentralised markets such as spot forex and CFDs. Not the same as real volume. [Ch 2, Ch 12]

TPO / Market Profile — A profile built from time (Time Price Opportunities, usually 30-minute periods marked by letters) rather than from volume. [Ch 6, Ch 7]

Trailing drawdown — A drawdown floor that rises with the account's peak and never falls. **End-of-day** trailing follows the highest closing balance; **intraday** trailing follows the highest value at any moment, including open profit. [Ch 11]

Trapped traders — Traders who entered aggressively at a price that the market then rejected, leaving them on the wrong side and potentially forced to exit. [Ch 9]

Trend day — A day of persistent directional movement with a thin, elongated profile and little time spent at any price. [Ch 6, Ch 7]

U

Unfinished auction — An auction that has not completed in one direction, seen as a poor high or low on a session profile, or as non-zero trading on both sides at a footprint bar's extreme. [Ch 7, Ch 9]

V

Value area (VA, VAH, VAL) — The price range containing about 70% of a session's volume, built outward from the POC. Its upper edge is the value area high (VAH) and its lower edge the value area low (VAL). [Ch 5]

Value migration — The movement of value areas from day to day: higher, lower, overlapping or unchanged. [Ch 6]

Velocity Logic — A CME mechanism that can briefly pause matching when price moves too far within a very short time, as can happen at major news releases. [Ch 11]

Veto — In this book's framework, the decision not to act because the order flow at a level contradicts your hypothesis. A valid and often valuable outcome. [Ch 11]

VWAP (Volume-Weighted Average Price) — The average price traded over a period, weighted by volume at each price: $\Sigma(\text{price} \times \text{volume}) \div \Sigma(\text{volume})$. For gold, its value depends on when the session starts (Globex or day session). [Ch 4]

VWAP bands — Lines drawn a number of volume-weighted standard deviations above and below VWAP, describing how stretched price is relative to the session's average. [Ch 4]

W

Walk-forward analysis — Repeatedly tuning on one period and testing on the next, moving through time, to see whether an idea is stable across market regimes. [Ch 12]

Wyckoff effort vs result — The principle that effort (volume) should produce a result (price movement); when large effort produces little result, someone is absorbing it. [Ch 9]

X

XAUUSD — The symbol for spot gold priced in US dollars, used by most CFD and forex brokers. Its price usually sits below the active GC futures price by the basis (roughly \$27–30 in October 2026). [Ch 0, Appendix]

Pre-Session Checklist (One Page)

Print this page, or copy it into your notes, and work through it before every session. Write answers, not ticks alone; the act of writing is what moves decisions from a hot state to a cold one. This is a process checklist, not a trading signal.

1. Contract and clock

- ☐ I am looking at the **active month** of GC (and MGC if I trade it), not an expiring contract or a back-adjusted chart.
- ☐ I know today's session definition (full Globex from 18:00 ET, or day session from about 08:20 ET) and my profile, VWAP and CVD all use the same one.
- ☐ I have checked for daylight-saving differences and for holiday or early-close hours.

2. Context (observations, not predictions)

- ☐ Yesterday's day type (balanced, trend, double distribution, rotational) and profile shape.
- ☐ Where we are now relative to yesterday's value: above, inside or below.
- ☐ What Asia and London did: overnight high and low, London high and low.
- ☐ Is value migrating higher, lower or overlapping over the last few days?
- ☐ Anything notable in related markets (dollar index, Treasury yields), as observation only.

3. News

- ☐ Today's releases and their exact times in ET (CPI, NFP, FOMC, speeches).
- ☐ My written rule for before and after each release.
- ☐ If I trade XAUUSD, the release times converted to my platform's clock.

4. Level map (five to seven levels)

- ☐ Yesterday's VAH, VAL and POC.
- ☐ Any naked POC within reach.
- ☐ Session VWAP (and weekly or anchored VWAP if relevant).
- ☐ Prior day high and low; overnight and London high and low.
- ☐ Levels ranked by confluence; levels within two or three ticks merged.
- ☐ For each level: source, rank and "what I want to see".
- ☐ **XAUUSD equivalents** written next to every GC level, using today's measured gap.

5. Scenarios

- ☐ Two or three if-then scenarios, each with a written **invalidation**.
- ☐ My confirmation and veto checklist for each scenario, in objective terms.

6. Risk

- ☐ Risk per trade in dollars and in R.
- ☐ Contract choice (GC or MGC) checked against the typical logical stop distance today.
- ☐ Maximum number of trades.
- ☐ Daily loss limit in R.
- ☐ Stop condition (for example, after two consecutive losses).
- ☐ If I trade a funded account: type of drawdown floor and today's distance to it.

7. Personal state

- [] Sleep, focus, stress (rate 1–5). If low, reduced size or no trading today.
- [] No outside commitment that will force me to manage a position badly.

8. After the session

- [] Journal completed for every trade and every veto (entry part was filled in before the outcome).
- [] Plan compared with what happened: which scenario, if any, played out?
- [] Screenshots saved.
- [] One lesson written in one sentence.

Converting GC Levels to XAUUSD

Most of the analysis in this book is done on gold futures (GC), because only an exchange with a central order book shows real volume and the aggressor side of every trade. But many readers trade gold as a CFD, usually under the symbol XAUUSD, on a broker platform. The two prices are not the same, so every futures level needs an approximate CFD equivalent.

Why there is a gap

GC is a contract for gold delivered at a future date. XAUUSD tracks spot gold, for immediate settlement. Holding gold until a future date has a cost, mainly the interest you could have earned on the money, partly offset by the income from lending the gold. This **cost of carry** makes the futures price sit above spot in normal conditions. The further away the delivery month, the larger the gap.

A rough formula:

$$\text{Futures} \approx \text{Spot} \times (1 + \text{net carry rate} \times \text{days to expiry} \div 365)$$

Illustrative numbers. With spot gold at 4,168, a net carry rate of about 3% a year and about 84 days to the expiry of the December contract:

$$4,168 \times 0.03 \times 84 \div 365 \approx \$28.8$$

That is why, in October 2026, GC traded roughly **\$27–30 above** XAUUSD. You do not need to compute the carry yourself; you measure the gap directly. The formula only explains why it exists and how it behaves.

How the gap behaves

- **It shrinks slowly** as the contract approaches expiry, because fewer days of carry remain.
- **It jumps at the roll.** When the active month changes (for example from December to February), the new contract is further away, so the gap widens. With the same illustrative numbers, two extra months of carry add roughly \$20.
- **It drifts** a little during the day and from day to day as interest rates, lease rates and supply and demand change.

- **It differs slightly between brokers**, because each CFD broker builds its own XAUUSD price feed and adds its own spread.

How to measure it

1. At one moment, ideally just before you write your level map, read the **active-month GC** bid and ask and compute the mid price: $(\text{bid} + \text{ask}) \div 2$.
2. At the **same moment**, read your own broker's XAUUSD bid and ask and compute the mid price.
3. $\text{Gap} = \text{GC mid} - \text{XAUUSD mid}$.
4. Convert every level: **XAUUSD level $\approx$ GC level – gap**. To go the other way: $\text{GC level} \approx \text{XAUUSD level} + \text{gap}$.
5. Re-measure at least once per session, and always after a roll.

Using mid prices matters: comparing a GC ask with an XAUUSD bid would add half of both spreads to your gap.

Worked example

Illustrative numbers, not a market analysis.

At 07:45 ET:

Instrument	Bid	Ask	Mid
GC (December, active month)	4,196.0	4,196.2	4,196.1
XAUUSD (your broker)	4,168.15	4,168.45	4,168.30

$$\text{Gap} = 4,196.1 - 4,168.3 = \mathbf{27.8}$$

Now convert the level map:

Level	GC	XAUUSD $\approx$ (GC – 27.8)
London high	4,215.3	$\approx$ 4,187.5
Yesterday's VAH	4,210.0	$\approx$ 4,182.2
Yesterday's POC	4,196.6	$\approx$ 4,168.8
Yesterday's VAL	4,181.5	$\approx$ 4,153.7
Naked POC (three days ago)	4,160.4	$\approx$ 4,132.6

Write the CFD equivalents as **zones** rather than exact prices, for example "VAL $\approx$ 4,152–4,155 XAUUSD", because the gap drifts by small amounts and your broker's feed is not the futures market.

What converts and what does not

- **Price levels convert** approximately: value area edges, POC, VWAP, session highs and lows.
- **Volume does not convert.** The volume, delta and footprint information behind those levels comes from the futures market. Your broker's XAUUSD volume is tick volume and does not tell the same story.
- **Order-book events do not convert.** A sweep or an iceberg seen in GC happened in the futures book, not in your broker's feed. The CFD price usually follows, because arbitrage keeps the markets aligned, but the CFD chart will not show the event itself.
- **Times may need converting.** Many CFD platforms display server time, often two or three hours ahead of London, not New York time.

Gold note: Always say "approximately" when you quote a converted level. The gap is an estimate measured at one moment, from one broker. It is close enough for a level map, not precise enough for exact orders.

Contract Cheat Sheet: GC and MGC

Verified October 2026. Always re-check the official CME Group contract specifications, margins and holiday calendar before trading.

The two contracts side by side

Item	GC (Gold futures)	MGC (Micro Gold futures)
Exchange	COMEX (CME Group), traded on CME Globex	COMEX (CME Group), traded on CME Globex
Contract size	100 troy ounces	10 troy ounces
Price quoted in	US dollars per troy ounce	US dollars per troy ounce
Tick size	\$0.10 per ounce	\$0.10 per ounce
Tick value	\$10 per contract	\$1 per contract
Value of a \$1.00 move	\$100 per contract (10 ticks)	\$10 per contract (10 ticks)
Value of a \$10.00 move	\$1,000 per contract	\$100 per contract
Notional value at \$4,200	\$420,000	\$42,000
Equivalence	1 GC	10 MGC = 1 GC in exposure
Order book and volume	Its own (deeper, larger participants)	Its own, separate from GC
Typical use in this book	Reading order flow	Flexible position sizing

Hours and key times (New York time, ET)

Time (ET)	What happens
Sunday 18:00	Weekly open on Globex
18:00 daily	Start of the CME trading day (trading after 18:00 belongs to the next day's date)
18:00 → ~03:00	Asia session (boundaries are conventions)
~03:00	London active
~08:20	Traditional COMEX day session begins; start of the initial balance used in this book
08:30	Many US data releases (CPI, NFP and others)
08:20–09:20	Initial balance as defined in this book
09:30	US stock market opens
13:29–13:30	GC daily settlement window
14:00 / 14:30	FOMC statement / press conference (on FOMC days)
17:00–18:00	Daily halt
Friday 17:00	Weekly close

Remember that the US and Europe change daylight-saving time on different dates, so for a few weeks a year the gap between London and New York is an hour different from usual.

Contract months

Futures month codes: F Jan, G Feb, H Mar, J Apr, K May, M Jun, N Jul, Q Aug, U Sep, V Oct, X Nov, Z Dec. A contract is written as product + month code + year, for example **GCZ26** (GC, December 2026) or **MGCZ26**.

For gold, the main active months are traditionally February, April, June, August and December (G, J, M, Q, Z), and CME settlement notices have also treated October as a lead month. Volume concentrates in one active month at a time and migrates to the next one around each roll. Check which contract is active before every session.

Delivery, margin and limits

- **Delivery.** Both contracts are deliverable. Day traders close or roll positions before the delivery period; check CME's calendar for first notice and last trading dates.
- **Margin.** Margin is collateral, not the price of the contract. Exchange margins are typically a small percentage of notional value and change, sometimes sharply in volatile periods; brokers set their own day-trading margins on top. Never treat a margin figure as fixed.

- **Losses can exceed your deposit.** Futures losses are not limited to the margin you posted.
- **Fast-market protections.** CME uses mechanisms such as Velocity Logic and price-fluctuation limits that can pause or limit trading in extreme moves.

Quick formulas

- **P&L** = (price change ÷ 0.10) × tick value × number of contracts.

Example: gold rises from 4,200.0 to 4,202.5 = 25 ticks. On 1 MGC: $25 \times \$1 = \$25$. On 1 GC: $25 \times \$10 = \$250$.

- **Position size** = allowed risk ÷ (stop distance in ticks × tick value + costs per contract). Round down.
- **Costs in ticks** = round-trip commission and fees in dollars ÷ tick value. The same dollar cost is ten times as many ticks on MGC as on GC.
- **XAUUSD equivalent** ≈ GC level – measured gap (roughly \$27–30 in October 2026; always measure, always say approximately).